



Parallels Remote Application Server

Administrator's Guide

18.3

Parallels International GmbH
Vordergasse 59
8200 Schaffhausen
Switzerland
Tel: + 41 52 672 20 30
www.parallels.com

© 2022 Parallels International GmbH. All rights reserved. Parallels and the Parallels logo are trademarks or registered trademarks of Parallels International GmbH in Canada, the U.S., and/or elsewhere.

Apple, Safari, iPad, iPhone, Mac, macOS, iPadOS are trademarks of Apple Inc. Google, Chrome, Chrome OS, and Chromebook are trademarks of Google LLC.

All other company, product and service names, logos, brands and any registered or unregistered trademarks mentioned are used for identification purposes only and remain the exclusive property of their respective owners. Use of any brands, names, logos or any other information, imagery or materials pertaining to a third party does not imply endorsement. We disclaim any proprietary interest in such third-party information, imagery, materials, marks and names of others. For all notices and information about patents please visit <https://www.parallels.com/about/legal/>

Contents

Introduction	15
Parallels RAS 18 release history	15
About Parallels RAS	15
About this guide	16
Terms and abbreviations used in this guide	17
Installing Parallels RAS	19
System requirements	19
Hardware requirements	19
Software requirements	20
Microsoft license requirements	23
Install Parallels RAS	23
Log in and activate Parallels RAS	24
Getting Started with Parallels RAS	27
The Parallels RAS Console	27
Set up a basic Parallels RAS Farm	29
Add an RD Session Host	30
Publish applications	36
Invite users	37
Azure Virtual Desktop	41
Conclusion	41
Farm and Sites	42
Connecting to a Parallels RAS Farm	42
About Sites	44
Sites in the RAS Console	45
Adding a Site to the Farm	47
Replicating Site settings	48
Managing Licensing Site	49
Managing administrator accounts	50
Adding an administrator account	50
Administrator account permissions	51

Managing administrator accounts.....	53
Configure RAS Console idle sessions.....	55
Using instant messaging.....	55
Joining Customer Experience Program	56
RAS Publishing Agent	57
Configuring RAS Publishing Agents.....	57
Secondary Publishing Agents.....	59
Managing Secondary Publishing Agents	61
Using computer management tools.....	63
RAS Secure Client Gateway.....	64
Overview	64
Adding a RAS Secure Client Gateway	66
Manually adding a RAS Secure Client Gateway.....	66
Checking the RAS Secure Client Gateway status	67
Configuring a RAS Secure Client Gateway	67
Enable or disable a Gateway	67
Set public address.....	68
Set IP addresses for client connections	68
Site defaults (Gateways)	68
Gateway mode and forwarding settings	69
Gateway network options.....	69
SSL/TLS encryption	70
Configure HTML5 Gateway	74
Wyse ThinOS support	76
Gateway security.....	77
Web request load balancing	78
Gateway tunneling policies.....	80
Configure logging.....	81
Viewing Gateway summary and metrics	81
Using computer management tools.....	82
RD Session Hosts.....	83
RD Session Host types	83
Add an RD Session Host.....	84
Installing the agent manually.....	86

Planning for high availability	87
Viewing RD Session Hosts	87
Configuring an RD Session Host	89
Check an RD Session Host Agent status.....	89
Change RD Session Host Site assignment	90
View and modify RD Session Host properties.....	90
User profile	96
Optimization	102
Drive redirection cache	107
Configure logging	107
Grouping and cloning RD Session Hosts.....	108
Using scheduler	112
RD Session Host drain mode examples	115
Maintaining RD Session Hosts based on a template.....	115
Managing RD Session Host sessions	116
Managing logons	116
Using computer management tools.....	117
Publishing from an RD Session Host	118
Publishing a desktop	118
Publishing an application	119
Publishing a web application	121
Publishing a network folder.....	121
Publishing a document	122
Publishing containerized applications	124
Publishing App-V applications	124
Viewing published resources	125
Virtual Desktop Infrastructure (VDI)	126
Supported VDI providers	126
RAS VDI Agent information	127
RAS VDI Agent installation options	128
Add a VDI provider	129
Add a hypervisor VDI provider	129
Add a cloud VDI provider.....	131
Installing RAS VDI Agent using the installer.....	137

Modifying VDI Provider configuration	138
Configure logging.....	142
Enabling high availability for VDI	142
Change VDI provider Site assignment	144
Site defaults (VDI).....	144
Viewing guest VMs on a VDI provider	146
Templates	147
Template types.....	147
Creating a template	150
How guest VMs are created from a template.....	161
Manually adding a guest VM.....	162
Assigning a template to a pool or group	163
Template maintenance	163
Template status.....	166
Managing multi-provider template distribution	168
Managing template-based guest VMs	169
VDI pool management.....	170
Adding and deleting pools.....	170
Adding and deleting pool members.....	170
Using a wildcard to filter VMs	171
Managing guest VMs in pools.....	171
Managing guest VMs	172
Persistent guest VMs	174
Using a VDI provider in multiple farms	175
Using computer management tools.....	176
Publishing from a guest VM.....	176
Publishing a desktop	176
Publishing an application	177
Publishing a web application	178
Publishing a network folder.....	179
Publishing a document	180
Viewing VDI provider summary	181
Managing VDI sessions	181
Remote PC pools.....	181

Adding a VDI provider.....	182
Configuring a VDI provider.....	183
Adding Remote PCs to a pool.....	184
Managing Remote PCs in a pool	185
Persistent Remote PCs	186
RAS Guest Agent installation options	186
Publishing from a pool-based Remote PC.....	187
Azure Virtual Desktop.....	188
Introduction.....	188
Prerequisites	190
Deploy Azure Virtual Desktop	192
Enable Azure Virtual Desktop and add a provider	193
Add a workspace	194
Add a host pool.....	195
Manage Azure Virtual Desktop	197
Manage providers.....	197
Manage workspaces	198
Manage host pools.....	198
Manage templates.....	202
Manage hosts	205
Manage sessions.....	207
Using scheduler.....	207
Site defaults (Azure Virtual Desktop)	208
Site defaults for multi-session hosts	208
Site defaults for single-session hosts.....	211
Verify the deployment.....	213
Publish resources.....	213
Using Parallels Client with Azure Virtual Desktop	214
Remote PCs.....	216
Overview	216
Adding a Remote PC to a Farm	216
Admin-initiated Remote PC enrollment	217
Self-service Remote PC enrollment	218
Configuring a Remote PC	220

Viewing Remote PC summary	222
Using computer management tools.....	222
Publishing from a Remote PC	223
Publishing a desktop	223
Publishing an application	223
Publishing a web application	224
Publishing a network folder.....	225
Publishing a document	225
Session Management.....	227
Overview	227
Session information.....	228
Monitoring settings.....	231
Managing sessions	231
The Resources tab.....	233
Published Resources Management.....	234
Overview	234
General management tasks	235
Manage published applications	236
Manage published desktops	240
Manage published documents	241
Manage folders	243
Site defaults (Publishing)	245
Using filtering rules	247
Configuring preferred routing.....	250
Understanding session prelaunch	252
Checking effective access	253
Specifying client settings	254
Quick keypad.....	256
SSL Certificate Management	258
Generating a self-signed certificate.....	259
Generating a certificate signing request (CSR)	260
Importing a certificate.....	260
Exporting a certificate.....	261

Assigning a certificate to Gateways and HALB	261
Auditing certificates	263
Permissions to manage certificates	263
Upgrading from an older RAS version	264
Connection and Authentication Settings	265
RAS Publishing Agent connection settings	265
Remote session settings	266
Restricting access by Parallels Client type and build number	268
Multi-factor authentication	269
Using RADIUS	269
Using Deepnet DualShield	275
Using SafeNet	282
Using Google Authenticator	283
Configuring MFA exceptions	285
Allowing users to change domain password	286
Load Balancing and HALB	288
Resource based & round robin load balancing	288
Configure CPU optimization	290
High availability load balancing (HALB)	291
Deploying a Parallels HALB appliance	292
Adding a HALB virtual server	293
HALB Device status and version number	296
HALB maintenance	296
HALB connection and session information	297
Changing the HALB appliance password	297
RAS Multi-Tenant Architecture	299
Overview	299
Architecture description	300
Implementation overview	301
User connection flow	302
Deploying Tenant Broker and Tenants	303
Deploying Tenant Broker	304
Deploying a Tenant	305
User authentication	310

Unjoining from Tenant Broker	310
Managing Tenants	311
Tenant configuration.....	311
Deleting a Tenant object.....	312
Opening a Tenant console.....	312
Shared Gateways.....	312
Third-party network load balancers	313
HTML5 Client and Themes.....	314
Monitoring Tenants	315
Tenant Broker compatibility and updates	315
Upgrading from an older RAS version	316
Configuring notifications	316
Communication ports.....	317
SAML SSO Authentication	319
Introduction.....	319
System requirements	321
SAML basics.....	322
SAML configuration.....	323
Prerequisites.....	323
IdP side configuration	324
SP side configuration (RAS side)	325
Active Directory user account configuration.....	327
Configure certificate authority templates.....	329
RAS Enrollment Server configuration	338
RAS Enrollment Server high availability	340
SAML integration examples and tips	340
Parallels Client configuration.....	341
Parallels client policy configuration	342
Test the SAML SSO deployment.....	342
Error messages.....	343
Parallels HTML5 Client.....	345
Configure HTML5 Client	345
Configure Themes.....	346
General Theme settings.....	346

HTML5 Client Theme settings	347
Parallels Client for Windows Theme settings.....	350
General Theme tasks.....	351
Delegating session management permissions	352
Open Parallels HTML5 Client.....	353
Main menu options	355
Running remote applications and desktops.....	357
Using drag and drop functionality	357
Native clipboard experience	358
Other useful features	358
Auto login.....	359
Direct App access.....	360
Using the toolbar.....	361
Using the toolbar on desktop computers.....	361
Using the Toolbar on Mobile Devices.....	363
Using the remote clipboard	364
Hiding toolbar items	365
Universal Printing	367
Managing Universal Printing settings	367
Universal Printing drivers	368
Font management.....	369
Universal Scanning.....	371
Managing Universal Scanning	371
Adding scanning applications.....	372
User Device Management and Client Policies	373
Inviting users to connect to Parallels RAS.....	373
Mass configuring user devices	374
Enabling Help Desk support.....	375
Monitoring devices	376
Getting additional device information.....	377
Windows device groups.....	377
Managing Windows devices.....	379
Windows desktop replacement	383
Scheduling Windows devices & groups power cycles	385

Client Policies.....	386
Add a new client policy.....	387
Configure session settings.....	388
Configure client policy options.....	401
Configure control settings.....	405
Configure Gateway redirection	406
Client policy backward compatibility.....	407
Policy information in Parallels Client.....	407
Configuring remote file transfer.....	408
Configure file transfer for a server	409
Configure file transfer for HTML5 Gateway	410
Configure file transfer for a client policy.....	410
Reporting	411
System requirements	411
Install Microsoft SQL Server	413
Install Microsoft SQL Server 2016 or earlier.....	413
Install Microsoft SQL Server 2017 or 2019	416
Install Parallels RAS Reporting.....	416
Running Parallels RAS reports.....	418
GDPR compliance	421
Performance Monitor	422
Overview.....	422
Install RAS Performance Monitor.....	423
Using Parallels RAS Performance Monitor	423
Configure RAS Performance Monitor Security	426
Updating RAS Performance Monitor	427
Common Management Tasks	429
Recovery - add a root administrator	429
Host name resolution	430
Computer management tools.....	431
Site information	433
Site settings	433
Settings audit.....	435
Upgrading RAS agents	437

Licensing	438
Configure HTTP proxy settings.....	440
System event notifications.....	440
Configuring notification handlers.....	440
Configuring notification scripts.....	443
Configuring SMTP server connection for event notifications	446
RAS session variables	446
Maintenance and backup.....	448
Exporting and importing Farm settings from the command line	448
Problem reporting and troubleshooting	450
Logging.....	451
Suggest a feature.....	453
Parallels RAS Management Portal.....	454
Overview	454
Prerequisites	455
Installation.....	455
Log in to RAS Management Portal	456
Configure RAS Web Administration Service.....	456
RAS Management Portal user interface	457
Parallels RAS APIs.....	461
RAS PowerShell API	461
RAS REST API	463
Installation	463
Permissions.....	464
Getting started	464
Logging in and sending requests.....	465
More information	467
RAS HTML5 Gateway API and Parallels Client URL scheme	467
Appendix.....	469
Microsoft license requirements in Parallels RAS.....	469
Port reference	473
Parallels Client	474
Web browsers	474
HALB	475

Contents

RAS Secure Client Gateway	475
RAS Publishing Agent	476
RAS Console.....	477
SSRS	478
RAS Reporting	478
RAS Web Administration Service (REST/Management Portal)	478
RAS PowerShell	479
RAS VDI Agent	479
RAS Enrollment Server	480
RAS RD Session Host Agent.....	480
RAS Guest Agent	480
RAS Remote PC Agent	481
Tenant Broker	481
Active Directory and Domain Services ports	481
Azure Virtual Desktop	482
RAS performance counters	483
Index	485

CHAPTER 1

Introduction

Welcome to Parallels® Remote Application Server (Parallels RAS), an integrated solution to virtualize your applications, desktops and data. Parallels RAS publishes applications and delivers remote and virtual desktops to any device on your network, anywhere.

In This Chapter

Parallels RAS 18 release history.....	15
About Parallels RAS	15
About this guide	16
Terms and abbreviations used in this guide	17

Parallels RAS 18 release history

The following table lists the Parallels RAS 18 release history. Parallels RAS documentation is updated for every release. This guide refers to the latest Parallels RAS 18 release from the table below. If you are using a newer Parallels RAS release or version, please download the current version of the guide from <https://www.parallels.com/products/ras/resources/>.

Parallels RAS Version	Release	Date
18.0	Initial release	12/14/2020
18.0	Update 1	03/03/2021
18.1	Initial release	07/14/2021
18.2	Initial release	11/03/2021
18.3	Initial release	12/21/2021
18.3	Update 1	02/16/2022

About Parallels RAS

Parallels RAS provides vendor independent virtual desktop and application delivery from a single platform. Accessible from anywhere with platform-specific clients and web enabled solutions, like the Parallels RAS HTML5 Gateway, Parallels RAS allows you to publish remote desktops, applications and documents, improving desktop manageability, security and performance.

Parallels RAS extends Windows Remote Desktop Services by using a customized shell and virtual channel extensions over the Microsoft RDP protocol. Parallels RAS supports all major hypervisors from Microsoft, VMware, and other vendors including Hyperconverged solutions such as Nutanix and Scale Computing and Cloud platforms and services such as Microsoft Azure and Azure Virtual Desktop (formerly known as Windows Virtual Desktop), enabling the publishing of virtual desktops and applications to Parallels Client.

The product includes powerful universal printing and scanning functionality, as well as resource-based load balancing and management features.

With Parallels Device Manager Module for Parallels RAS you can also centrally manage user connections and PCs converted into thin clients using the free Parallels Client.

How does it work?

When a user requests an application or a desktop, Parallels RAS finds a least loaded RD Session Host or a guest VM on one of the least loaded VDI providers and establishes an RDP connection with it. Using Microsoft RDP protocol, the requested application or desktop is presented to the user. Note that in addition to RD Sessions Hosts and VDI, Parallels RAS can also be used to configure, manage and publish Azure Virtual Desktop resources.

Users can connect to Parallels RAS using Parallels Client (available at no charge), which can run on Windows, Linux, macOS, Android, Chrome, iOS and iPadOS. Users can also connect via an HTML5 browser or Chromebook.

As newer versions of Windows keep on being developed as time goes by, you need to defend the migration cost to your business. Parallels RAS can help. Desktop replacement allows you to extend the lifespan of your hardware and delay migration to the latest OSs to a time that suits you best. The Parallels RAS solution allows you to be very flexible: you can lock machine configurations on the user side, placing your corporate data in an extremely secure position; or you can opt to allow users to run some local and remote applications. Parallels Client Desktop Replacement is able to reduce the operability of the local machine by disabling the most common local configuration options, while guaranteeing the same level of service and security afforded by thin clients, directly from your existing PCs.

About this guide

This guide is intended for system administrators responsible for installing, configuring, and administering Parallels RAS. This guide assumes that the reader is familiar with Microsoft Remote Desktop Services and has an intermediate networking knowledge.

Terms and abbreviations used in this guide

Term/Abbreviation	Description
RAS Console	Parallels RAS Console. The RAS console is the primary interface you use to configure, manage, and run Parallels RAS. As an administrator, you use the RAS console to manage Farms, Sites, RD Session Hosts, published resources, client connections, etc.
Category	In the RAS console, categories are displayed in the left pane of the main window. Each category consists of a number of settings related to a specific task or operation. The categories include Start, Farm, Load Balancing, Publishing, Universal Printing, Universal Scanning, Connection, Device Manager, and others.
Farm	A Parallels RAS Farm is a logical grouping of objects for the purpose of centralized management. A Farm configuration is stored in a single database which contains information about all objects comprising the Farm. A Farm consists of at least one Site but may have as many sites as necessary (see Site below).
Site	A Site consists of at least one RAS Publishing Agent, RAS Secure Client Gateway (or multiple gateways), and RAS agents installed on RD Session Hosts, VDI providers, and Windows PCs. Note that a given RD Session Host, VDI provider, or PC can be a member of only one Site at any given time.
Licensing Site	The Site that manages Parallels RAS licenses in a Parallels RAS Farm. By default, the server on which you install Parallels RAS becomes the Licensing Site. If you create additional sites later, you can designate any one of them as the Licensing Site. There can be only one Licensing Site in a given Farm. All other sites are called secondary sites. Note: Parallels RAS updates or upgrades must be applied to the Licensing Site first.
RAS Secure Client Gateway	RAS Secure Client Gateway tunnels all traffic needed by applications on a single port and provides secure connections.
HTML5 Client	HTML5 client allows users to view and launch remote applications and desktops in a web browser. The HTML5 client functionality is a part of RAS Secure Client Gateway.
Publishing	The act of making items installed on a Remote Desktop Server, VDI provider or Remote PC available to the users via Parallels RAS.
RAS Publishing Agent	RAS Publishing Agent provides load balancing of published applications and desktops.
RAS RD Session Host Agent	RAS RD Session Host Agent collects information from Microsoft RDS hosts required by the Publishing Agent and transmits to it when required.
Remote PC Agent	Remote PC Agent collects information from Remote PC hosts required by the Publishing Agent and transmits to it when required.
RAS Guest Agent	RAS Guest Agent collects information from the VDI desktop required by RAS Publishing Agent and transmits to it when required.

<p>RAS VDI Agent / RAS Provider Agent</p>	<p>RAS VDI Agent collects information from the Parallels RAS Infrastructure and is responsible for controlling VDI through its native API. RAS VDI Agent is built into the RAS Publishing Agent and is available by default. It can be used to control multiple VDI providers in a Parallels RAS Farm.</p> <p>RAS Provider Agent is the same as RAS VDI Agent, but the term is used in the context of Azure Virtual Desktop (described at the end of this table).</p>
<p>RAS VDI Agent dedicated</p>	<p>RAS VDI Agent dedicated is similar to the RAS VDI Agent described above with one important difference — it is a separate component that must be installed from the Parallels RAS installer and can only control a single VDI provider.</p>
<p>RDSH or RD Session Host</p>	<p>RDSH makes applications and a full desktop accessible to a remote client that supports Remote Desktop Protocol (RDP). RDSH replaced Terminal Server beginning with Windows 2008 R2.</p>
<p>HALB</p>	<p>High Availability Load Balancing (HALB) is an appliance that provides load balancing for RAS Secure Client Gateways. Parallels HALB virtual appliance is available for the following hypervisors: Hyper-V, VMware. Multiple HALB Virtual Servers representing different HALB devices can be deployed in a single Site.</p> <p>Multiple HALB deployments can run simultaneously, one acting as the primary and others as secondaries. The more HALB deployments a Site has, the lower the probability that end users will experience downtime. Primary and secondary HALB deployments share a common or virtual IP address (VIP). Should the primary HALB deployment fail, a secondary is promoted to primary and takes its place.</p>
<p>Tenant Broker</p>	<p>Tenant Broker is a special RAS installation that hosts shared RAS Secure Client Gateways. It is an essential part of the RAS multi-tenant architecture.</p>
<p>Tenant</p>	<p>Tenants are RAS farms that join Tenant Broker (see above) and use shared RAS Secure Client Gateways and HALB thus eliminating the need to have their own Gateways and HALB deployed.</p>
<p>RAS Enrollment Server</p>	<p>RAS Enrollment Server is an essential component of the SAML SSO Authentication functionality. It communicates with Microsoft Certificate Authority (CA) to request, enroll, and manage digital certificates on behalf of the user for SSO authentication in the Parallels RAS environment.</p>
<p>RAS PowerShell</p>	<p>Parallels RAS PowerShell allows you to perform Parallels RAS administrative tasks using PowerShell cmdlets. You can execute cmdlets in the Windows PowerShell console or you can write scripts to perform common Parallels RAS administrative tasks. A complete guide to Parallels RAS PowerShell is available on the Parallels website together with other Parallels RAS documentation.</p>
<p>RAS REST API</p>	<p>Parallels RAS comes with various APIs to help you develop custom applications that integrate with it. The RAS REST API is one of them.</p>
<p>RAS Management Portal</p>	<p>Parallels RAS Management Portal is an HTML5 browser-based application that lets you manage Parallels RAS.</p>
<p>RAS Web Administration Service</p>	<p>A Web service that provides the user interface for RAS Management Portal and implements RESTful Web services for the RAS REST API (see above).</p>
<p>Azure Virtual Desktop</p>	<p>Azure Virtual Desktop is a desktop and app virtualization service running on Microsoft Azure, providing access to RD Session Hosts and VDI. Parallels RAS 18 provides the ability to integrate, configure, maintain, support and access Azure Virtual Desktop workloads on top of the existing technical capabilities of Parallels RAS.</p>
<p>FSLogix</p>	<p>FSLogix Profile Container is a remote profile solution for non-persistent environments. Parallels RAS supports FSLogix on RD Session Hosts, VDI, and Azure Virtual Desktop.</p>

CHAPTER 2

Installing Parallels RAS

This chapter describes how to install and activate Parallels RAS.

In This Chapter

System requirements.....	19
Install Parallels RAS	23
Log in and activate Parallels RAS	24

System requirements

Before installing Parallels RAS, please verify that your hardware and software meet or exceed hardware and software requirements described below. Please note that although Parallels RAS can be used in Workgroup environment, Parallels recommends using Active Directory to manage users, groups, and machine accounts via group policies.

Hardware requirements

Parallels RAS is extensively tested on both physical and virtual platforms. The minimum hardware requirements approved to run Parallels RAS are outlined below.

- Physical Machines – Dual Core Processor and a minimum of 4GB RAM.
- Virtual Machines – Two Virtual Processors and a minimum of 4GB of RAM.

The server hardware requirements to install and configure Parallels RAS can vary according to end-user requirements.

Typically for an installation of 30 users or under, Parallels RAS can be installed on one high specification server and the resources published directly from it. For more than 30 users, multiple servers may be required.

The below should be considered during the planning stage of a Parallels RAS deployment:

- High specification servers should be used, consisting of multiple CPU cores, a high specification disk transfer rate and plenty of RAM.
- A hypervisor-based virtual machine can be used as long as the resources needed to serve end-users are calculated accordingly.

- It is recommended that RAS Secure Client Gateway does not exceed 1200 users per server for incoming connections using the Gateway SSL mode.
- HALB usage should not exceed 2000 user sessions per HALB appliance. See <https://kb.parallels.com/125229>.
- When planning VDI Hypervisor resource requirements, extra requirements such as RAM usage per virtual machine and disk space should be taken into account.

When configuring RD Session Hosts, VDI, or Azure Virtual Desktop, please keep in mind that different types of workloads require different session host configurations. For the best possible experience, scale your deployment depending on your users' needs. The following table gives you an idea of how different workload types affect session host configurations.

Workload	Example users	Example apps	Max users per vCPU	Minimum
Light	Basic data entry tasks	Database entry applications, command-line interfaces	6	2 vCPUs 8 GB RAM 16 GB storage
Medium	Consultants and market researchers	Database entry applications, command-line interfaces, Microsoft Word, static web pages	4	4 vCPUs 16 GB RAM 32 GB storage
Heavy	Software engineers, content creators	Database entry applications, command-line interfaces, Microsoft Word, static web pages, Microsoft Outlook, Microsoft PowerPoint, dynamic web pages	2	4 vCPUs 16 GB RAM 32 GB storage
Power	Graphic designers, 3D model makers, machine learning researches	Database entry applications, command-line interfaces, Microsoft Word, static web pages, Microsoft Outlook, Microsoft PowerPoint, dynamic web pages, Adobe Photoshop, Adobe Illustrator, CAD, CAM	1	6 vCPUs 56 GB RAM 340 GB storage

Note: Sizing guidelines are based on Microsoft recommendations on RDS or Azure Virtual Desktop multi-session hosts.

For port requirements, please see the **Port Reference** section.

Software requirements

RAS Publishing Agent and RAS Secure Client Gateway (64-bit versions only)

RAS Publishing Agent and RAS Secure Client Gateway are supported on the following operating systems:

- Windows Server 2008 R2 up to Windows Server 2022
- On Windows Server 2016, 2019, and 2022 both Server Core and Desktop Experience installations are supported

Note: RAS Publishing Agent and RAS Secure Client Gateway should not be installed on a domain controller or any other machine where a DHCP server is running. This in general applies to any of the RAS components.

RAS Web Administration Service

Same OS requirements as for RAS Publishing Agent (see above). Note that for larger environments (2000 or more concurrent connections), it is recommended to install the component on a dedicated server. For details, please see <https://kb.parallels.com/en/124988>.

Please also note that Windows Server 2012 R2 and 2008 R2 must have the following updates installed:

- Windows Server 2012 R2 — KB2999226
- Windows Server 2008 R2 — KB2999226 and KB2533623

Newer versions of Windows Server do not require any specific updates.

RAS RD Session Host Agent

RAS RD Session Host Agent is supported on the following operating systems:

- Windows Server 2008 up to Windows Server 2022
- Windows Server 2016 and newer must be installed using the "Desktop Experience" installation option.
- Windows Server 2012 R2 — Server Core installation option is not supported.

RAS VDI Agent / RAS Provider Agent

- Windows Server 2008 R2 up to Windows Server 2022

For the list of supported VDI providers, see **RAS VDI Agent Installation Options** (p. 128).

RAS Guest Agent

- Windows Server 2008 R2 up to Windows Server 2022
- Windows 7 up to Windows 11

Remote PC Agent

- Windows Server 2008 R2 up to Windows Server 2022
- Windows 7 up to Windows 11

Parallels RAS PowerShell

- Windows Server 2008 with Service Pack 2 up to Windows Server 2022
- Windows 7 up to Windows 11
- Windows Management Framework 3.0 and .NET Framework 4.5.2 must be installed

Parallels RAS Console

- Windows Server 2008 R2 up to Windows Server 2022
- Windows 7 up to Windows 11

RAS HTML5 Gateway

- Windows Server 2008 R2 up to Windows Server 2022

RAS Enrollment Server

- Windows Server 2008 R2 up to Windows Server 2022

Parallels Client

Parallels Client is approved for the following operating systems (both 32-bit and 64-bit systems are supported, where applicable):

- Windows 7, 8.x, 10, 11
- Windows Server 2008 R2 up to Windows Server 2022
- macOS 10.11 El Capitan up to macOS 12 Monterey. Parallels Client runs natively on both Intel and Apple M1 Mac computers.
- iOS 11.x, 12, 13.x, 14, 15; iPadOS 13.x, 14, 15.
- Android 5.0 up to 12
- Chrome OS

Parallels Client for Linux supports the following Linux distributions (x64 versions only):

- Ubuntu 16.04 LTS, 18.04 LTS, 20.04 LTS
- Debian 9 (Stretch), Debian 10 (Buster)
- Fedora 33, 34
- CentOS 7, 8
- Linux Mint 19, 20
- ThinOS/ Dell Wyse Thin Clients. Note that ThinOS 9 and above are not supported.

For a list of supported thin clients and supported hardware from Technology Partners such as Igel, HP, 10Zig, and more, please see the following KB article: <https://kb.parallels.com/124606>.

Microsoft license requirements

For information about Microsoft license requirements, such as Remote Desktop Services Client Access Licenses (RDS CALs) and Virtual Desktop Access (VDA) licenses, please see **Appendix: Microsoft license requirements in Parallels RAS** (p. 469).

Install Parallels RAS

To install Parallels RAS:

- 1 Make sure you have administrative privileges on the computer where you are installing Parallels RAS.
- 2 Double click the `RASInstaller.msi` file to launch the Parallels RAS installation wizard. If you see a message that begins with "This version of Parallels RAS is only for testing purposes.", it means that it's not an official build and should not be used in a production environment.
- 3 Follow the onscreen instructions.

Note: Please ensure that the presented terms in the license agreement are read and accepted to complete installation and/or upgrade. For programmatic deployment, it is understood that the terms in the license agreement have been read and accepted.

- 4 Proceed to the **Select Installation Type** page and select from the following:
 - **Parallels Remote Application Server.** The default installation that will install RAS Console, RAS Management Portal, RAS Publishing Agent, RAS Secure Client Gateway, RAS RD Session Host Agent, RAS PowerShell, and RAS Web Administration Service on the same machine. This is ideal for testing or small production environments.
 - **Parallels RAS Tenant Broker.** This option installs Tenant Broker. Please note that Tenant Broker must be installed on a server separate from the existing RAS farms. For more information about Tenant Broker, please see the **RAS Multi-Tenant Architecture** chapter (p. 299).
 - **Custom.** Select and install only the components that you require. You can select individual components after you click **Next**. Note that if a component cannot be installed on the current server, it will not be available for installation. See **Software Requirements**.
- 5 Click **Next**.
- 6 Review the notice on the **Important Notice** wizard page. If there's a port conflict on your computer, the information will be displayed here. You can resolve the conflict later.
- 7 Click **Next**.
- 8 On the **Firewall Settings** page, select **Automatically add firewall rules** to configure the firewall on this computer for Parallels RAS to work properly. See **Port Reference** for details.
- 9 Click **Next** and then click **Install**. Wait for the installation to finish and click **Finish**.

When you need to install a particular Parallels RAS component on a different server, run the installation wizard again, select **Custom** and choose the component(s) you wish to install.

Log in and activate Parallels RAS

After you've installed Parallels RAS, run the RAS Console and activate your new Parallels RAS Farm.

Start the Parallels RAS Console

By default, the Parallels RAS Console is launched automatically after you click **Finish** on the last page of the installation wizard. To launch the console manually, navigate to **Start > Apps > Parallels** and click on **Parallels Remote Application Server Console**.

When the Parallels RAS Console is launched for the first time, you are presented with the login dialog. In the dialog, specify the following:

- **Farm:** A Parallels RAS Farm to connect to. Enter the FQDN or IP address of the server where you have RAS Publishing Agent installed.
- If you've installed the Parallels Single Sign-On component when installing the RAS Console, you will see the **Authentication type** field from which you can select whether to log on using your credentials or SSO. If you reboot after the installation and select SSO, select **Single Sign-On** and then click **Connect**. Your Windows credentials will be used to log in to the RAS Farm. If you select **Credentials**, enter your credentials as described below.
- **Username:** A user account with administrative privileges on the server where Parallels RAS is installed (usually a domain or local administrator). The account name must be specified using the UPN format (e.g. administrator@domain.local). The specified user will be automatically configured as the Parallels RAS administrator with full access rights.
- **Password:** The specified user account password.
- If you select the **Remember credentials** option, this dialog will not be shown the next time you launch the Parallels RAS Console.

The **Edit Connections** button opens a dialog where you can manage your RAS connection. This dialog becomes useful if this is not the first time you are connecting to one or more of your RAS Farms. The left pane of the dialog displays RAS Farms to which previously connected (you can remove a Farm from the list by clicking the **[-]** icon if you no longer need it). The right pane displays at least the primary Publishing Agent for the selected Farm. If you've added a secondary Publishing Agents to a Farm, you can add it to this list by clicking the **[+]** icon and typing its hostname or IP address (click the "recycle" icon to verify the agent status). This way the RAS Console will try to connect to the primary Publishing Agent first and if it fails (e.g. the agent is offline or cannot be reached), it will try to connect to the secondary Publishing Agent. For more information about secondary Publishing Agents, please see **Parallels RAS Publishing Agents** chapter (p. 57).

When you are done entering the connection information, click the **Connect** button to connect to the Parallels RAS Farm.

Sign in to Parallels My Account

To activate Parallels RAS, you must register for a Parallels business account. After you logged in to Parallels RAS, you'll see the **Sign In to Parallels My Account** dialog. If you already have an account, type the email address and password you used to register the account and click **Sign In**.

Note: If you use an HTTP proxy server on your network, you will see a dialog asking you to configure the proxy server connection settings. Click the **Configure Proxy** button. In the dialog that opens, select one of the following: **Use system proxy settings** (the default proxy settings from the Internet Explorer will be used) or **Manual HTTP proxy configuration** (specify the settings manually). If your proxy configuration changes, you can re-configure it later by navigating to **Administration > Settings** and clicking the **Configure Proxy** button.

If you don't have a Parallels business account, you can register for one as follows:

- 1 In the **Sign In to Parallels My Account** dialog, click **Register**. The **Register Parallels My Account** dialog opens.
- 2 Enter your name and email address, choose and type a password, and enter your company info (all fields are required).
- 3 Follow the links to Parallels Privacy Policy and Terms of Use. After reading them (and if you agree) select the **I have read and agree to the Parallels Privacy Policy and Terms of Use** checkbox.
- 4 Click **Register** to register an account. This will create a personal account for yourself and a business account for your organization to which you will be assigned as administrator.

Activate Parallels RAS

After you sign in to Parallels My Account, the **Activate Product** dialog opens asking you to activate the Parallels RAS Farm.

If you already have a Parallels RAS license key, select the **Activate using license key** option and enter the key in the field provided. You can click the button next to the field to see the list of subscriptions and/or permanent license keys you have registered in Parallels My Account. If the list is empty, it means that you don't have any subscriptions or license keys and need to purchase one first.

Note: You can manage your Parallels RAS license using the **Licensing** category in the Parallels RAS console. The management tasks include viewing the license information, switching to a different Parallels My Account, and activating the Parallels RAS Farm using a different license key. For more information, please see the **Licensing** section (p. 438).

If you don't have a Parallels RAS license key, you have the following options:

- Purchase a subscription online by clicking the **Purchase a license** link.
- Activate Parallels RAS as a trial by selecting the **Activate trial version** option.

After entering a license key (or selecting to activate a trial version), click **Activate**. You should see a message that the Parallels RAS Farm was activated successfully. Click **OK** to close the message box.

The first dialog that you see informs you that you have no servers configured that can be used to host published resources. This means that to begin using Parallels RAS, you need at least one RD Session Host, VDI provider, or a Remote PC configured. We'll talk about configuring a Parallels RAS Farm in the next chapter. For now, click **OK** to close the message box. You will then see the **Applying Settings** dialog. Wait for the initial configuration of Parallels RAS to complete and click **OK**. You will now see the main Parallels RAS Console window where you can begin configuring the Parallels RAS Farm.

Read on to learn how to quickly add an RD Session Host, publish resources, and invite your users to Parallels RAS.

Getting Started with Parallels RAS

This chapter will help you get started with Parallels RAS. Read it to learn how to use the Parallels RAS Console and how to set up a simple RAS environment.

In This Chapter

The Parallels RAS Console	27
Set up a basic Parallels RAS Farm.....	29

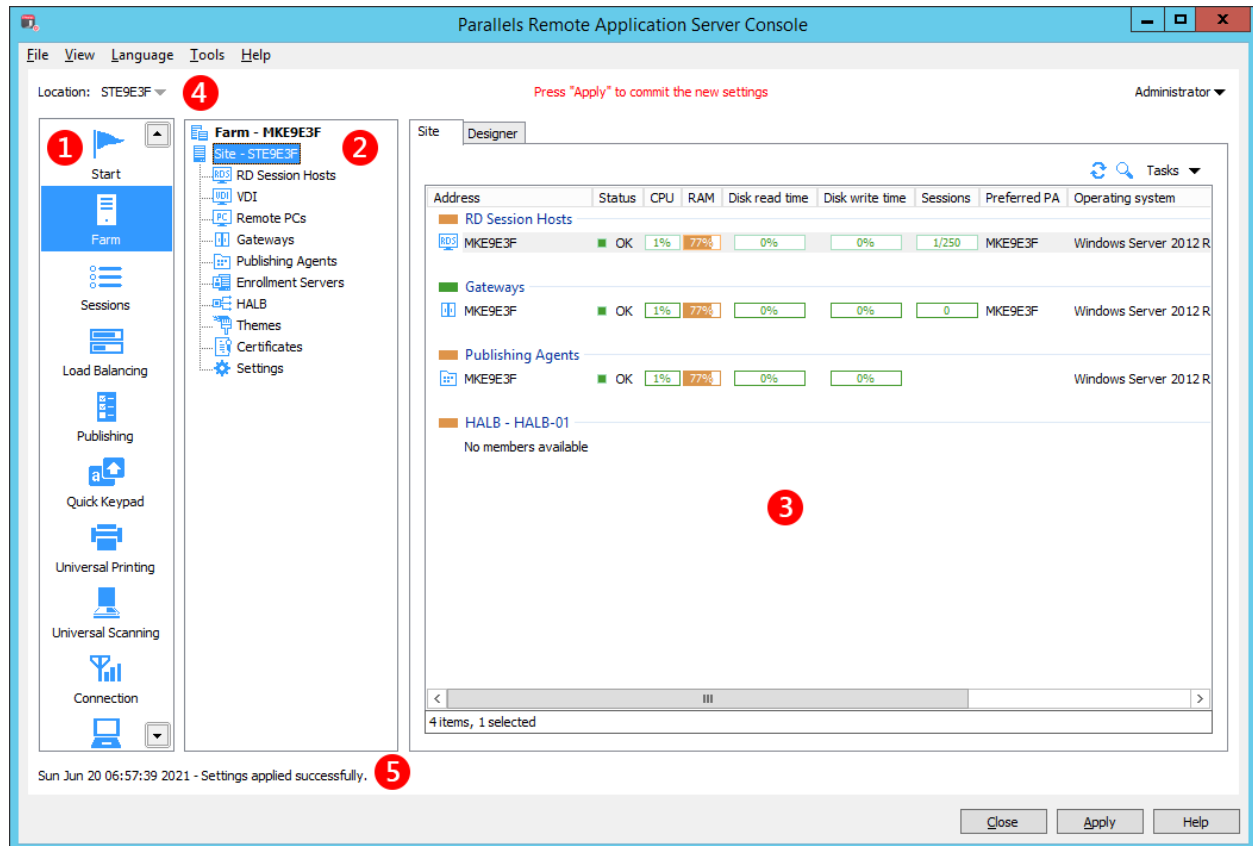
The Parallels RAS Console

The Parallels RAS Console is a Windows application used to configure and administer a Parallels RAS Farm.

To open the Parallels RAS Console, navigate to **Apps > Parallels** and click **Parallels Remote Application Server Console**. Note that you can open multiple instances of the Parallels RAS Console on the same computer if you want to manage more than one Farm or Site simultaneously without switching between them inside the console. This works with a locally installed Parallels RAS Console and when you run it as a remote application from Parallels Client.

Information: In addition to Parallels RAS Console, Parallels RAS 18 introduced Parallels RAS Management Portal, an HTML5 browser-based console that lets you manage Parallels RAS. Note that at the time of this writing, Parallels RAS Management Portal does not completely replace the desktop RAS Console as some management features are still in development. More features will be added in the upcoming releases. For more information, please refer to **Parallels RAS Management Portal Guide**, which is available on the Parallels website: <https://www.parallels.com/products/ras/resources/>.

The following screenshot and the description below it give you an overview of the Parallels RAS Console:



The Parallels RAS Console consists of the following sections:

- 1** This section lists categories. Selecting a category will populate the right pane with elements relevant to that category.
- 2** This section (the middle pane) is available only for the **Farm** and the **Publishing** categories. The navigation tree allows you to browse through objects related to that category.
- 3** This section displays the selected object or category properties, such as servers in a Farm or published application properties, etc.

- 4 The information bar at the top of the RAS Console displays the name of the Site you are currently logged in to on the left side (the **Location** field). If you have more than one Site, you can switch between them by clicking the drop-down menu (the Site name) and choosing a desired Site. If you used the RAS Console to connect to more than one Farm, the drop-down menu will also display the other Farm name(s), clicking on which will connect the console to that Farm.

Your administrator account name is displayed on the right side. Clicking on the name opens a drop-down menu from which you can initiate a chat with other administrators, show current sessions, and log off from the RAS Console.

The **Press 'Apply' to commit the new settings** message in the middle (in red) appears after you make any changes to any of the components or objects. It reminds you that you need to apply these changes to Parallels RAS for them to become effective. The following describes how it works.

When you make changes in the RAS Console, they are saved in the database as soon as you click **OK** in a dialog. If you close the console at this point, the changes will remain in the database and will not be lost. The changes, however, are not yet applied to running instances of the Parallels RAS processes, so they have no effect in the running RAS Farm. When you click the **Apply** button (at the bottom of the screen) the changes are applied to the runtime and become effective immediately.

When modifying anything in the RAS Console, follow these rules. When you make a small change, you can click **Apply** as soon as you are done with it. If you are working on something that requires many modifications in many places, you can wait until you are done with all changes and only then press **Apply** to apply all of them at the same time.

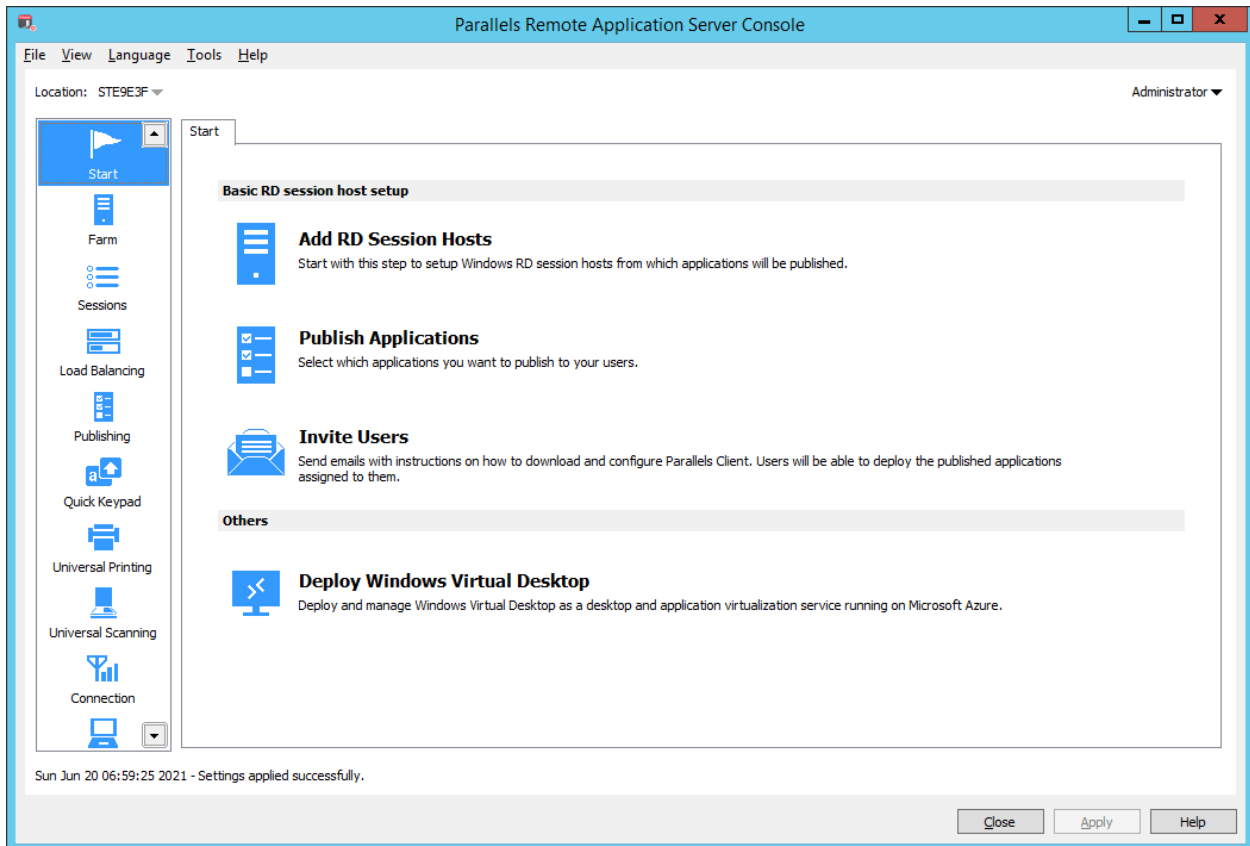
- 5 The information bar at the bottom of the screen is used to display the most recent console notification (if one is available).

Set up a basic Parallels RAS Farm

In this section, we'll set up a basic Parallels RAS Farm where all required components run on a single server.

To set up a Parallels RAS Farm:

- 1 Log in to the Parallels RAS Console.
- 2 In the console, select the **Start** category. This category gives you access to three wizards that you can use to easily perform essential tasks, such as adding RD Session Hosts, publishing applications, and inviting users to Parallels RAS.

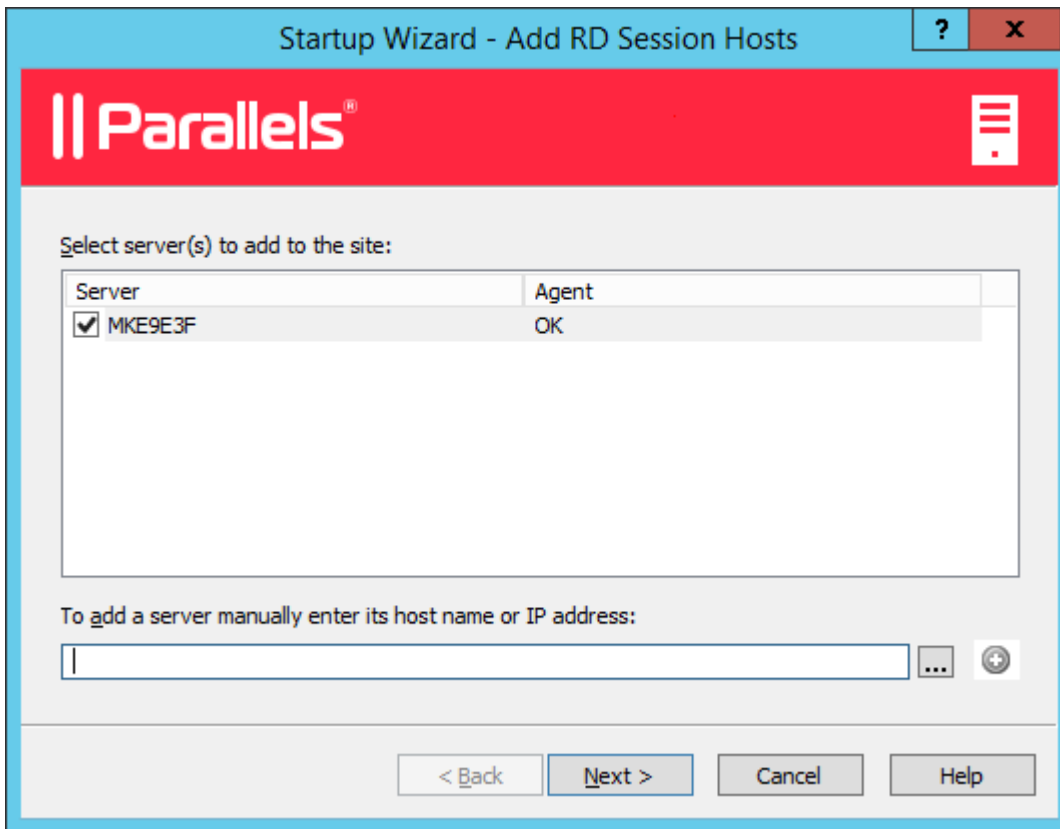


Add an RD Session Host

First, you need to add an RD Session Host to the Farm. In this tutorial, we'll add the local server on which Parallels RAS is installed.

To add an RD Session Host to the Farm:

- 1 Click **Add RD Session Hosts**. The **Add RD Session Hosts** wizard opens.

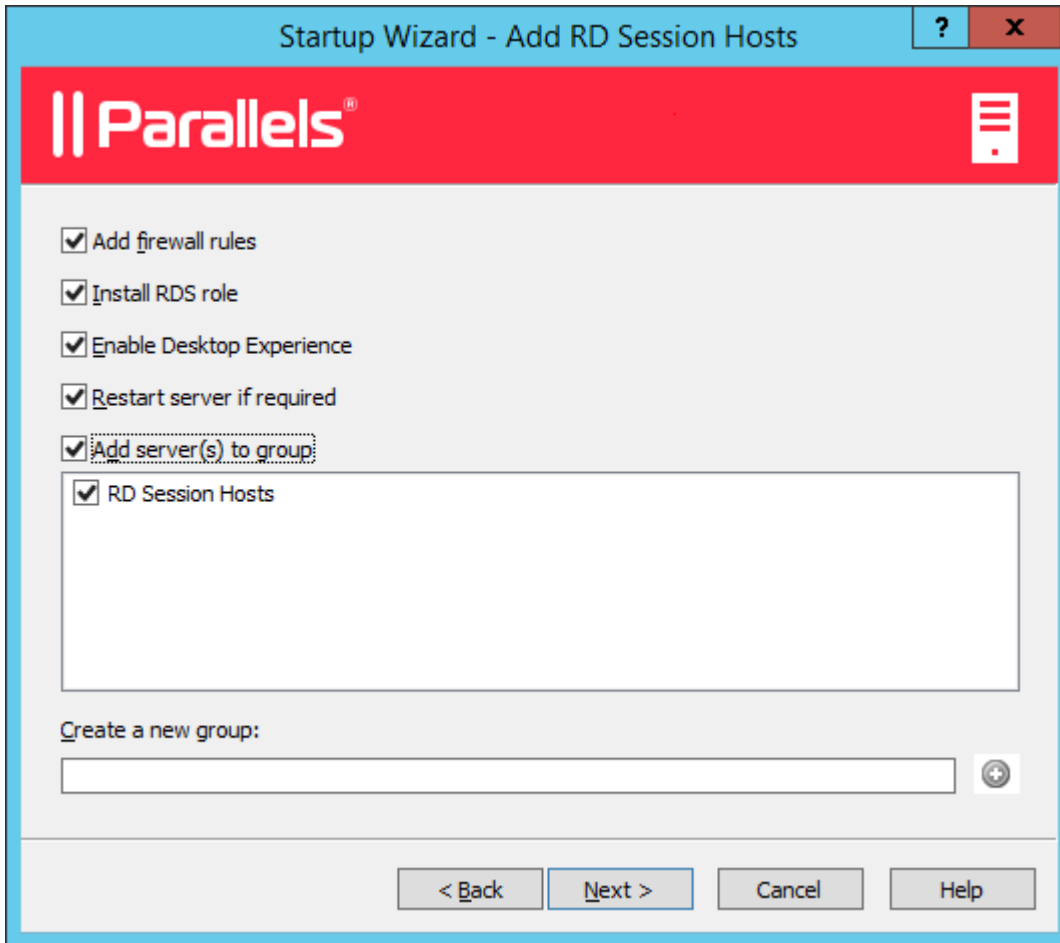


- 2 Select a server or type a server FQDN or IP address and then click the plus-sign icon to add the server to the list. If you are testing this by adding the localhost and see the FQDN warning, you can ignore it.

Note that if you enter the server FQDN, it will be used as the primary method of connecting to this server from other Parallels RAS components and clients. If you enter the IP address, it will be automatically resolved to FQDN, but only if the global option to resolve to FQDN is enabled. To see the current setting of this global option, click **Tools > Options** on the main menu. In the **Options** dialog, examine the **Always attempt to resolve to fully qualified domain name (FQDN) when adding hosts** option. When the option is selected, the IP address of every server/component in the RAS Farm is always resolved to FQDN. When the option is cleared, whatever is specified for a server (IP address or name) is used to communicate with a server. This makes a difference in deployments where an IP address cannot be used to access a server, such as when a server is hosted in the cloud. For more information, see **Host Name Resolution** (p. 430).

- 3 Click **Next**.

4 The page with general settings opens:



Specify the following settings:

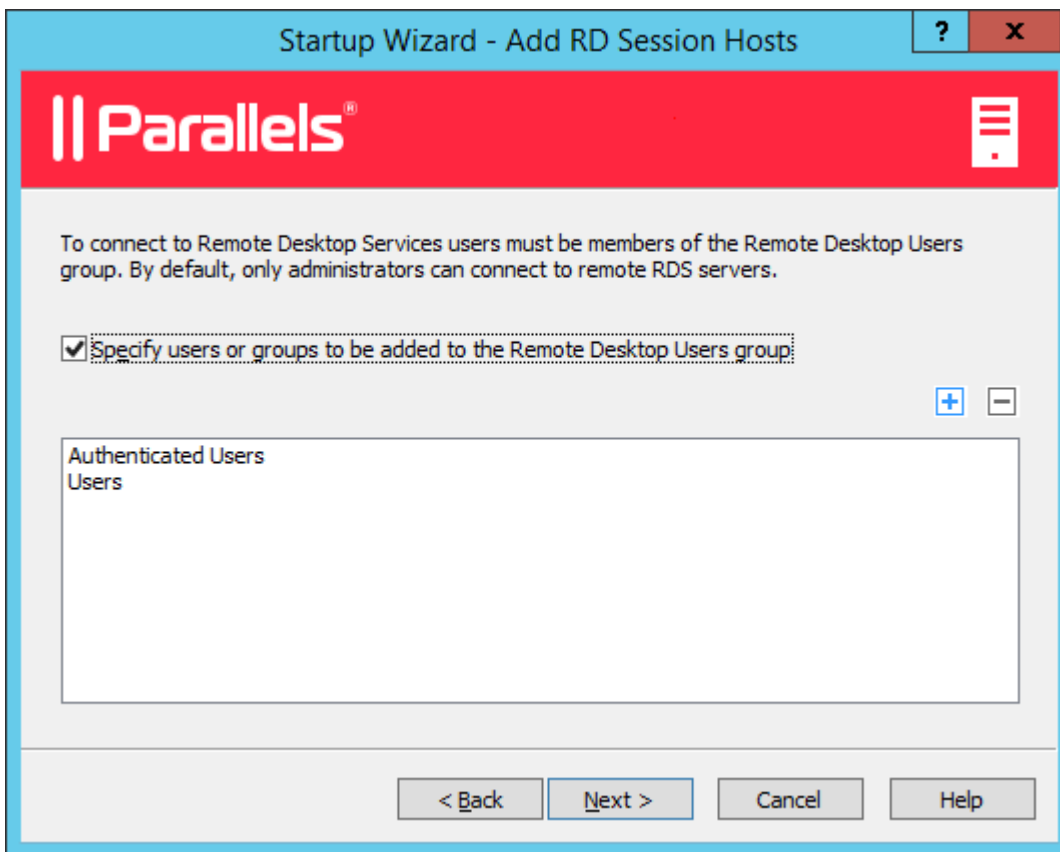
- **Add firewall rules.** Add firewall rules required by Parallels RAS in Windows running on the server. See **Port Reference** for details.
- **Install RDS role.** Install the RDS role on the server if it's not installed. You should always select this option.
- **Enable Desktop Experience.** Enable the Desktop Experience feature in Windows running on the server. This option is enabled only if the Install RDS role option (above) is selected. The option applies to Windows Server 2008 R1/R2 and Windows 2012 R1/R2 on which the Desktop Experience feature is not enabled by default.
- **Restart server if required.** Automatically restart the server if necessary. You can restart the server manually if you wish.
- **Add server(s) to group.** Add the server (or servers) to a group. Select the desired group in the list box located below this option. Groups are described in detail in the **Grouping RD Session Hosts** (p. 108) section. If you are just learning how to use this wizard, you can skip this option.

5 Click **Next**.

- 6 In order for end users to access published resources on the RD Session Host, they must be added to the Remote Desktop Users group in Windows running on the server. This can be done one of the following ways:
- Adding each user or group directly on the server using standard Windows administrative tools.
 - Adding users or groups through Active Directory.
 - Using the wizard page described below, which is provided for your convenience.

If you already added your users to the Remote Desktop Users group on the given server (or if for any reason you want to use one of the other methods listed above), you can simply click **Next** and skip this page.

To add users to the Remote Desktop Users group using the wizard, select the **Specify users or groups to be added to the Remote Desktop Users group** option and then click the **[+]** icon. In the **Select Users or Groups** dialog, specify a user or group and click **OK**. The selected user/group will be added to the list on the wizard page.



- 7 Click **Next**.

- 8 The **User profile** page allows you to select a technology to manage user profiles.

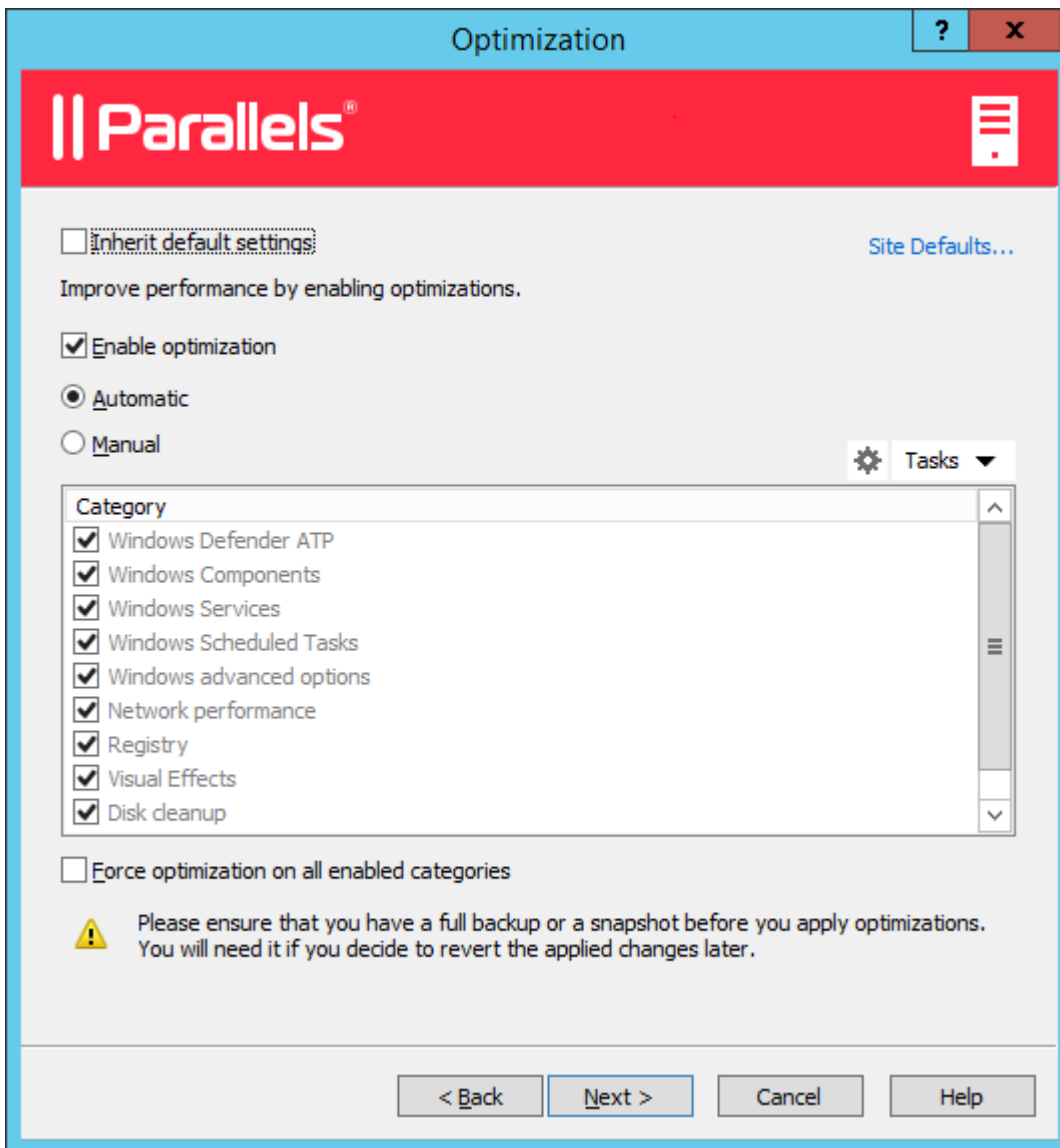
The screenshot shows the 'User profile' configuration window in Parallels RAS. The window has a blue title bar with a question mark and a close button. Below the title bar is a red header with the Parallels logo and a menu icon. The main content area is light gray and contains the following settings:

- Inherit default settings [Site Defaults...](#)
- Technology**: A dropdown menu with 'FSLogix' selected.
- Deployment method: Manual [Change...](#)
- Settings**
- Location type: SMB Location
- Location of profile disks: An empty text box with '+' and '-' buttons.
- Profile disk format: VHD
- Allocation type: Dynamic
- Default size: 30 GB (with up/down arrows) [Additional settings...](#)

At the bottom, there is an information icon and a note: 'Please ensure that FSLogix is not configured by GPO on the server(s). Storage permissions must be configured for use with Profile Containers.' Below the note are four buttons: '< Back', 'Next >', 'Cancel', and 'Help'.

You can select from **User profile disk** or **FSlogix**. User Profile Disks are virtual hard disks that store user application data on a dedicated file share. Microsoft FSLogix Profile Container is the preferred Profile Management solution as the successor of Roaming Profiles and User Profile Disks (UPDs). It is set to maintain user context in non-persistent environments, minimize sign-in times and provide native profile experience eliminating compatibility issues. You can keep the default settings for now. We will talk in detail about user profiles later in this guide (p. 96).

- 9 The **Optimization** page allows you to specify settings that will be used to optimize Windows on the RD Session Host for best performance in a Parallels RAS environment.



You can select Windows components, services, and other options that will be disabled, removed, or optimized to ensure a more efficient, streamlined, and improved delivery of virtual apps and desktops. You can keep the default settings or you can modify (or disable if not sure) optimization for now. Optimization is described in detail later in this guide (p. 102).

- 10 On the next page, review the settings and click **Next**.
- 11 The **Install RAS RD Session Host Agent** dialog opens. Follow the instructions and install the agent. When the installation is finished, click **Done** to close the dialog.
- 12 Back in the wizard, click **Finish** to exit.

If you would like to verify that the RD Session Host has been added to the Farm, click the **Farm** category (below the **Start** category in the left pane of the Parallels RAS Console window) and then click **RD Session Hosts** in the navigation tree (the middle pane). The server should be included in the **RD Session Hosts** list. The **Status** column may display a warning message. If it does, reboot the server. The **Status** column should now say, "OK", which means that your RD Session Host is functioning properly.

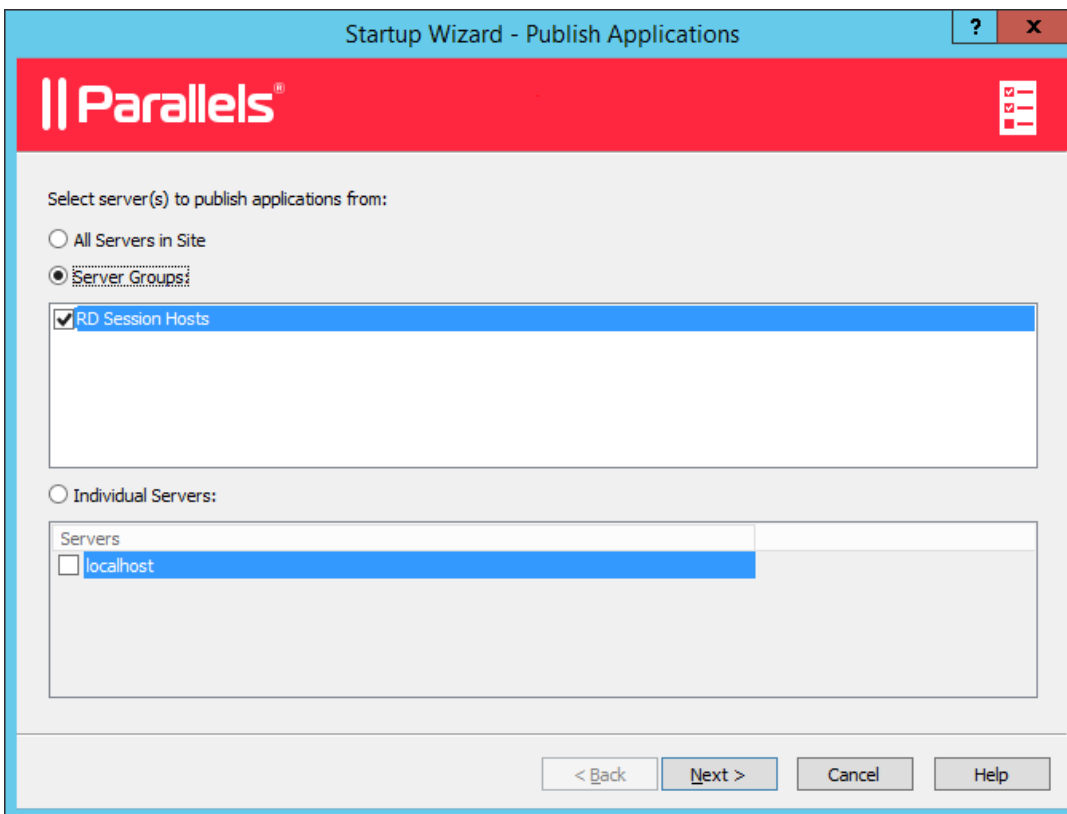
Read on to learn how to publish an application from an RD Session Host (p. 36)

Publish applications

After you added an RD Session Host, you can publish applications from it.

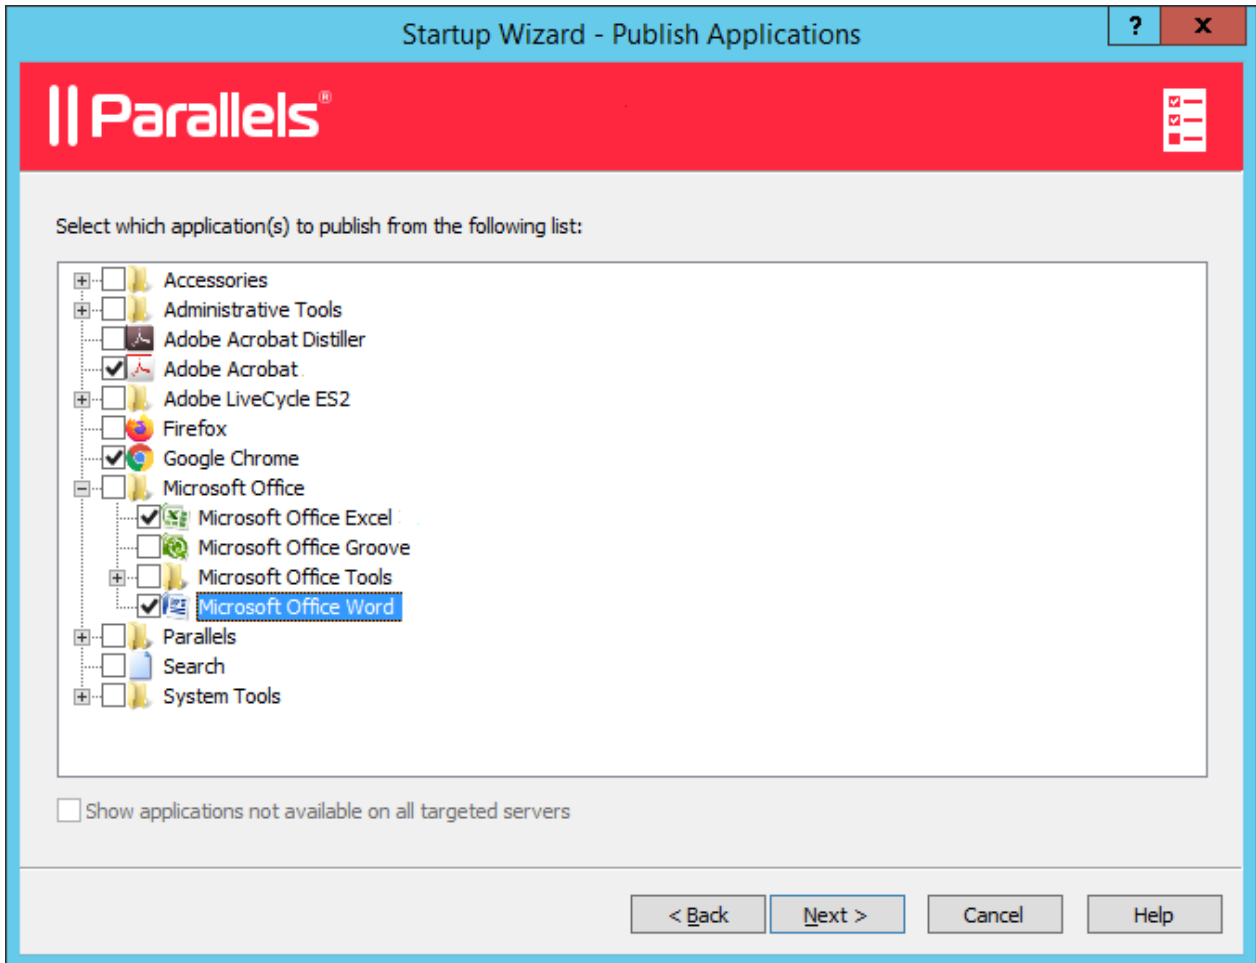
To publish an application:

- 1 In the Parallels RAS Console, select the **Start** category and click the **Publish Applications** item in the right pane.
- 2 The **Publish Applications** wizard opens. On the first page, select one or more servers from which the application should be published. You can select all servers, server groups, or individual servers.



- 3 Click **Next**.

- 4 On the next page, select one or more applications you want to publish.



If you've selected more than one server on the previous screen, the **Show applications not available on all target servers** option becomes enabled. If the option is cleared (default), the folder tree will contain applications that are available on each and every server that you selected. If the option is enabled, the tree will contain applications that may be available on some server(s), but not on the others.

- 5 Click **Next**. Review the summary information and click **Next** again.
 6 Click **Finish** when ready.

To verify that an application has been successfully published, select the **Publishing** category in the RAS Console. The application should be included in the **Published Resources** list (the middle pane).

Invite users

Your Parallels RAS Farm is now fully operational. You have an RD Session Host and published application(s). All you need to do now is invite your users to install the Parallels Client software on their devices and connect to the Parallels RAS Farm.

To invite users:

- 1 In the Parallels RAS Console, select the Start category and click the **Invite Users** item.
- 2 The **Invite Users** wizard opens. On the first page, specify the mailbox information that should be used to send invitation emails to users.

Startup Wizard - Invite Users

Parallels

Configure the mailbox from where the invitations will be sent from.

Mailbox configuration

Mailbox Server: mail.ras.com:500
Example: mail.yourcompany.com:500

Sender Address: admin@ras.com

TLS / SSL: Do not use

SMTP server requires authentication

Username: administrator

Password: ●●●●●●●●

Test email

Separate email addresses by a semicolon to send a Test Email to multiple addresses.

mikef@ras.com;andys@ras.com

Send Test Email

< Back Next > Cancel Help

Specify the following options:

- **Mailbox Server:** Enter the mailbox server name. For example, mail.company.com:500
- **Sender Address:** Enter the email address.
- **TLS / SSL:** Choose whether to use the TLS/SSL protocol.
- **SMTP server requires authentication:** Select this option if your SMTP server requires authentication. If it does, also type the username and password in the fields provided.

In the **Test Email** section, type one or more email addresses to which a test email should be sent (separate multiple address with a semicolon). Click the **Send Test Email** button to send the email.

- 3 Click **Next**.
- 4 On the next page of the wizard, specify target platforms and connection options:

Startup Wizard - Invite Users

Parallels

Specify target platform:

Name

- Windows
- macOS
- HTML5
- iOS
- Android
- ChromeApp
- Linux

Specify connection options:

Public address: (IP:Port/SSL Port) ...

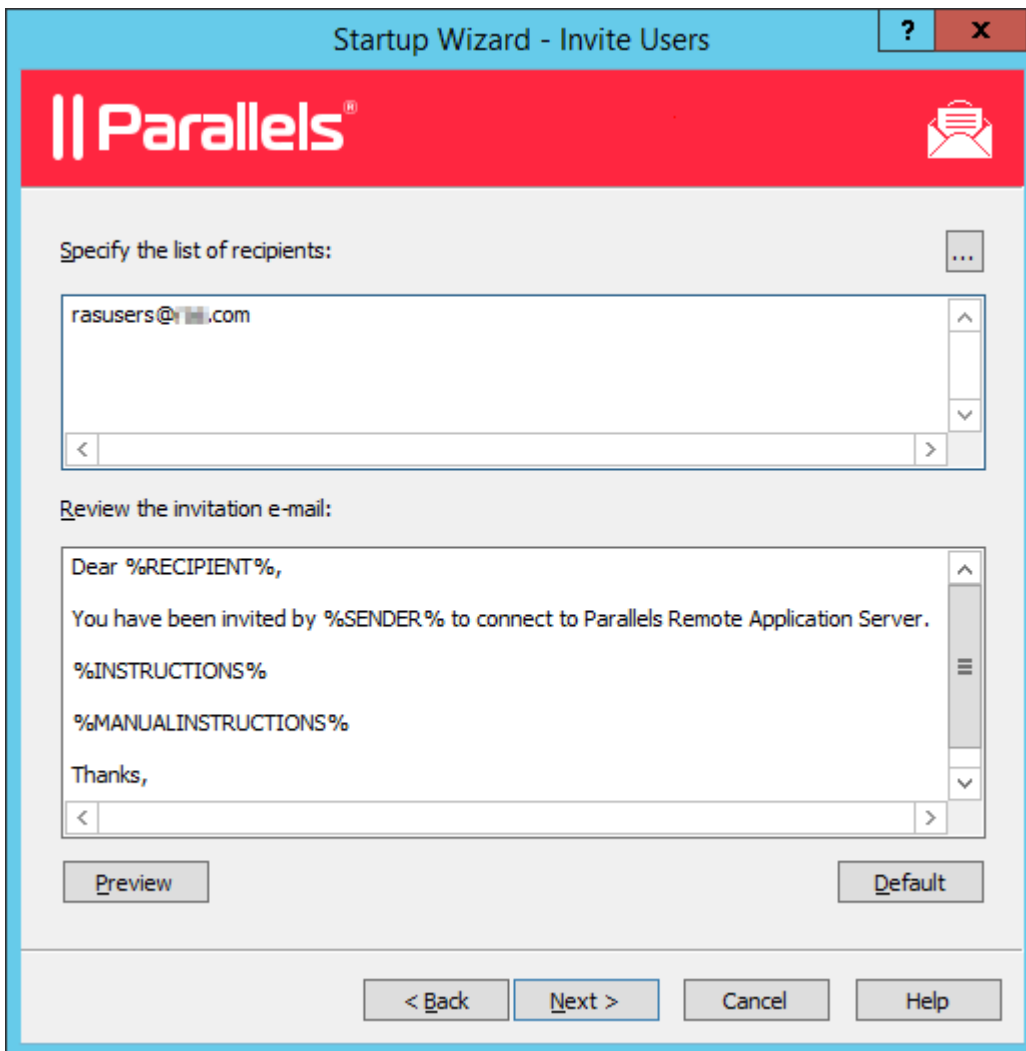
Connection Mode: ▼

Advanced

< Back Next > Cancel Help

- In the target devices list, select the types of devices to send an invitation to. Each target device of a particular type will receive an email with instructions on how to download, install, and configure the Parallels Client software on that device type.
- In the **Public address** field, specify a public FQDN or IP address. This setting is used by the Preferred routing functionality to redirect client connections. Please see **Configuring preferred routing** (p. 250).
- In the **Connection Mode** drop-down list, select the RAS Secure Client Gateway connection mode. Please note that SSL modes require the gateway to have SSL configured. More information can be found in the **Configuring RAS Secure Client Gateway** (p. 67) section.
- Optionally, click the **Advanced** button to open the **Advanced Settings** dialog. This dialog allows you to specify a third-party credential provider component. If you use such a component to authenticate your users, specify its GUID in this dialog. For more information, see **Configure Client Policy Options > Single Sign-On** (p. 401).

- 5 Click **Next**.
- 6 On the next page, specify the email recipients. Click the [...] button to select users or groups.



- 7 Review the invitation email template displayed in the **Review the invitation e-mail** box. You can modify the template text as needed. The template also uses variables, which are explained below.
 - %RECIPIENT% — Specifies the name of a recipient to whom the email message is addressed.
 - %SENDER% — The sender's email address that you specified in the first step of this wizard when you configured the outgoing email server settings.
 - %INSTRUCTIONS% — Includes a custom URL hyperlink for automatic configuration of Parallels Client. The URL uses the Parallels Client URL scheme. For more info, see **RAS HTML5 Gateway API and Parallels Client URL Scheme** (p. 467).
 - %MANUALINSTRUCTIONS% — Includes instructions for manual configuration of Parallels Client.

The variables are defined dynamically depending on the type(s) of the target devices and other settings. Normally, you should always include them in the message, so your users will receive all the necessary instructions and links. If you don't include any of the variables, you will see a warning message, but including all of them is not a requirement. To preview the message, click the **Preview** button. This will open the HTML version of the message in a separate window. This is the email message that your users will receive.

- 8 Click **Next**, review the summary and click **Next** again to send the invitation email to users.

When users receive the invitation email, they will follow the instructions that it contains to install and configure Parallels Client on their devices. Once that's done, the users will be able to connect to Parallels RAS and launch published resources.

Azure Virtual Desktop

The **Deploy Azure Virtual Desktop** section in the **Start** category is an optional feature, which allows you to deploy Azure Virtual Desktop in Parallels RAS. The feature is described in detail in the **Azure Virtual Desktop** chapter (p. 188).

Conclusion

In this tutorial, we have configured a simple Parallels RAS Farm with a single RD Session Host and one published application. We then configured a mailbox for outgoing emails and sent an invitation email to end users with instructions on how to install Parallels Client, connect to the Parallels RAS Farm, and run the published application. Essentially, we've created a fully functional Parallels RAS Farm serving remote applications to end users.

If you wish, you can repeat the tutorial and add more RD Session Hosts, publish more applications, or send an invitation email to users who use different types of devices. The instructions remain essentially the same.

The rest of this guide explains in detail how to configure and use various features of Parallels RAS.

CHAPTER 4

Farm and Sites

Parallels RAS Farm is a logical grouping of objects for the purpose of centralized management. A Farm configuration is stored in a single database which contains information about all objects comprising the Farm. A Site is the next level grouping in the Farm hierarchy which contains servers and other objects providing connection and remote application services.

In This Chapter

Connecting to a Parallels RAS Farm	42
About Sites	44
Sites in the RAS Console.....	45
Adding a Site to the Farm.....	47
Replicating Site settings	48
Managing Licensing Site.....	49
Managing administrator accounts.....	50

Connecting to a Parallels RAS Farm

If you have more than one Parallels RAS Farm in your organization, you can use the same Parallels RAS Console instance to manage any of them. By default, the Parallels RAS Console is installed on the same server where you install other Parallels RAS components, but you can install the console on any computer on your network.

Connecting to a Parallels RAS Farm for the first time

When you open the Parallels RAS Console for the first time, it displays the logon dialog on which you need to specify the following:

- **Farm:** A Parallels RAS Farm to connect to. Enter the FQDN or IP address of the server where you have RAS Publishing Agent installed.
- If you've installed the Parallels Single Sign-On component when installing the RAS Console, you will see the **Authentication type** field from which you can select whether to log on using your credentials or SSO. If you reboot after the installation and select SSO, select **Single Sign-On** and then click **Connect**. Your Windows credentials will be used to log in to the RAS Farm. If you select **Credentials**, enter your credentials as described below.

- **Username:** A user account with administrative privileges on the server where Parallels RAS is installed (usually a domain or local administrator). The account name must be specified using the UPN format (e.g. administrator@domain.com). The specified user will be automatically configured as the Parallels RAS administrator with full access rights.
- **Password:** The specified user account password.
- If you select the **Remember credentials** option, this dialog will not be shown the next time you launch the Parallels RAS Console.

After entering the connection properties, click **Connect** to connect to the Farm and open the RAS Console.

Note that the **Edit Connections** button will not display any information on first connect (it is used to edit Farm connections that already exist), so you can ignore it at this point. We will talk about using this button closer to the end of this section.

Connecting to a different Parallels RAS Farm

When you need to connect to a different Parallels RAS Farm, you first need to log off from the Parallels RAS Console in order to see the logon dialog again. To do so:

- 1 In the Parallels RAS Console, click on the arrow icon next to your user name in the upper right-hand corner and then choose **Log Off** in the context menu.
- 2 The console will close and the RAS logon dialog will open. The dialog will be populated with the current Farm connection properties.
- 3 To connect to a different Farm, type the FQDN or IP address of the server where the other Farm is located. Once again, this should be the server where you have the RAS Publishing Agent installed.
- 4 Specify a username and password and click **Connect**. The Parallels RAS Console will connect to the Farm using the connection properties that you specified.

Switching between Parallels RAS Farms

After you connect to more than one Farm from the same Parallels RAS Console instance, you can easily switch between them as follows:

- 1 In the Parallels RAS Console, click the **Location** drop-down menu in the upper left-hand corner (right below the main application menu, where the current Site name is displayed).
- 2 The lower portion of the drop-down list will contain names of the Farms to which you connected at least once in the past (the upper portion contains one or more Site names for the current Farm). Click a desired Farm name to connect to it.
- 3 When you click the Farm name, the console will close momentarily and will re-open connected to the Farm that you selected.

Note that you can also switch between Farms by logging off from the console and choosing a desired Farm from the **Farm** drop-down list in the RAS logon dialog. The method described above is more convenient, so this one is just another way to do it.

Editing Parallels RAS Farm connections

As was mentioned in the beginning of this section, the RAS logon dialog has the **Edit Connections** button. When you click it, the **Manage Parallels RAS Farm Connections** dialog opens.

On the left side of the dialog, the **Farm Connections** pane lists Parallels RAS Farms to which you connected at least once in the past. If a connection is no longer relevant, you can remove it by selecting it and clicking the "minus sign" icon at the top. Once a connection is removed, it will no longer appear in the RAS logon dialog and in the Parallels RAS Console (the **Location** drop-down list).

On the right side of the dialog, the **Publishing Agents** pane lists RAS Publishing Agents for the selected Farm connection. By default, the primary Publishing Agent is included in the list, but you can add more Publishing Agents if needed. When connecting to a Farm, the Parallels RAS Console will try the primary Publishing Agent first. If a connection cannot be established, it will try other Publishing Agents in the order they are listed in the **Publishing Agents** pane. To add a Publishing Agent to the list, click the "plus sign" icon and then specify the server FQDN or IP address.

About Sites

A Parallels RAS Farm consists of at least one Site, but may have as many sites as necessary.

Sites are often used to separate management and/or location functions. For example, by creating a Site, you can delegate permissions to a Site administrator without granting them full Farm permissions. Or you can have separate sites for different physical locations with the ability to copy the same settings to each Site while using RD Session Hosts, VDI providers, or PCs that are closer to end users or (depending on your needs) to back-end servers. For instance, it would make sense for a client/server application querying a database to be published from an RD Session Host which is located closer to the database server.

Each Site is completely isolated from other sites within the same Farm. The Farm simply groups sites logically and stores configuration properties of each Site (and the objects that comprise it) in a single database. sites don't communicate with each other and don't share any objects or data. The only exception to this rule is the RAS Licensing Site which periodically communicates with other sites to obtain statistics.

Individual object settings in a given Site can be replicated to all other sites. This does not mean that settings will be shared between sites. The settings that you choose will simply be applied to other sites. For more information, see the **Replicating Site Settings** section (p. 48).

When you install Parallels RAS, a Farm with a single Site is created automatically. This first Site becomes the RAS Licensing Site and the host for the main Parallels RAS configuration database. When you add more sites to the Farm, the data in this database is automatically synchronized with every Site that you add. When changes are applied to a particular Site, the main configuration database is automatically updated to reflect the changes.

Each Site must have at least the following components installed in order to publish remote applications and desktops for end users:

- Primary RAS Publishing Agent
- RAS Secure Client Gateway. Note that if a Site is joined as Tenant to RAS Tenant Broker, RAS Secure Client Gateway is not needed. For details, see **RAS Multi-Tenant Architecture** (p. 299).
- RD Session Host, VDI, or PC

When you install Parallels RAS using default installation options, the primary RAS Publishing Agent and the RAS Secure Client Gateway are automatically installed on the server on which you perform the installation. You can then add one or more RD Session Hosts to the Site to host published resources. You can also add more sites to the Farm if needed and configure individual components for each Site as you desire.

Sites in the RAS Console

To view existing sites in the Parallels RAS Console, select the **Farm** category in the left pane. Existing sites are listed in the right pane.

Note: The **Farm** node will only be visible to an administrator who has full permissions to manage the Farm. For more information about Farm/Site permissions, please refer to **Managing Administrator Accounts** (p. 50).

The **Farm** category displays the configuration of only one Site at a time. If you log in as the Farm administrator, the configuration of the RAS Licensing Site will be displayed. If you log in as an administrator who has access to a specific Site (but not the Farm), the configuration of that Site will be displayed.

Current Site

Click on the **Farm** item in the middle pane to view the list of available sites. The Site which configuration is currently loaded in the console is marked as "Current Site" in the **Type** column. The column also displays other Site attributes. For example, "Licensing Site / Local Site / Current Site".

Switching between sites

To switch to a particular Site, select **Farm** in the middle pane, then right-click the Site in the right pane and choose **Switch to this Site**. The Site configuration will be loaded into the RAS Console.

The other way of switching between sites is to click the **Location** drop-down menu in the upper left-hand side of the RAS Console. The menu lists sites for the current Farm and may also list other Farms if you used this RAS Console to connect to them. For more info, see **Connecting to a Parallels RAS Farm** (p. 42).

Renaming a Site

To rename a Site, right-click it and choose **Rename Site**.

Site configuration and health view

When you select the **Site** node in the middle pane, the **Site Info** tab in the right pane displays the list of Parallels RAS components that have been configured for the Site with interactive performance monitoring metrics for each component. Depending on the Site configuration, the list may include RD Sessions Hosts, VDI, Remote PCs, Gateways, Publishing Agents, Azure Virtual Desktop, HALB Virtual Servers and devices, Tenant Broker, Host pools, and Enrollment Server.

To collapse or expand a component group, click an "arrow up" or "arrow down" icon on the right side of the list. Note that if no servers of a particular type have been added to the Site, the group name will not be displayed in the list.

The following information is displayed for each component (the information is updated at an interval of approximately 2 minutes):

- **Address:** Server FQDN or IP address.
- **Status:** Indicates whether the agent software is installed on the server and is functioning properly.
- **CPU:** Current CPU utilization.
- **RAM:** Current RAM utilization.
- **Disk Read Time:** Disk read time.
- **Disk Write Time:** Disk write time.
- **Sessions:** The number of currently active user sessions.
- **Preferred PA:** The name of the RAS Publishing Agent designated as preferred for this server.
- **Operating System:** Operating system version installed on the server.
- **Agent Version:** The agent version installed on the server.

You can customize this view by clicking **Tasks > Monitoring Settings**. This opens a dialog where you can specify which colors should be used to display different performance counters and their values.

Performing tasks on a component

You can perform a number of tasks on a component displayed in the **Site Info** tab. These tasks are described below.

To configure a component, do one of the following:

- While the **Site** node is selected in the middle pane, right-click a component in the right pane and choose **Show in the editor**.
- Select a component category in the middle pane (e.g. RD Session Hosts, VDI providers, etc.).

To use server management tools, right-click a component (server), click **Tools** and choose a desired tool. For the complete description of tools, see **Computer Management Tools** (p. 431).

Using the Site Designer

Select the **Site** node in the middle and then click the **Designer** tab in the right pane. The tab displays a visual representation of the Site infrastructure. Use the icons at the top to add more components to the diagram as desired. Note that adding a component to the diagram will actually add it to the Site. Double-click a component to view and configure it in a corresponding editor.

Adding a Site to the Farm

To add a Site to the Farm:

- 1** In the RAS Console, select the **Farm** category in the left pane and then select the Farm in the middle pane.
- 2** In the **Tasks** drop-down menu (the right pane, above the Site list), click **Add** (or click the **+** icon).
- 3** In the **Add Site** dialog:
 - In the **Site** field, specify a Site name.
 - In the **Server** field, specify the IP address or FQDN of the server where the Primary Publishing Agent and Secure Client Gateway should be installed.
 - Select the **Add an SSL certificate and enable HTML5 Gateway** option to automatically create a self-signed certificate, enable SSL, and enable HTML5 support. For more info, please see **Configure HTML5 Client** (p. 74).
- 4** Click **Next**.
- 5** The **Site Properties** dialog opens. First, it verifies if RAS Publishing Agent is installed on the specified Site server. If it isn't, it will indicate this in the **Status** field.
- 6** Click the **Install** button to install the agent.

- 7 In the **Install RAS Publishing Agent** dialog, highlight the server name on which the RAS Publishing Agent is to be installed.
- 8 (Optional) Select the option **Override system credentials** to specify and use different credentials to connect to the server and install the agent.
- 9 Click **Install** to install the publishing agent and gateway. Click **Done** once it has been successfully installed.

Once a new Site is created, you can view and manage its configuration by right-clicking the Site in the RAS Console and choosing **Switch to this Site**.

Replicating Site settings

Site-specific settings configured for a given Site can be replicated to all other sites in a Farm. Refer to the table below for the information about which settings can be replicated to other sites.

Category	Section	Options
Farm	VDI > Templates	Auto removal timeout of guest VMs that fail preparation
Farm	VDI > Desktops	Auto removal timeout
Farm	Settings > Auditing	All settings
Farm	Settings > Global Logging	Logging settings
Farm	Settings > URL Redirection	All settings
Load Balancing	Load Balancing	All settings
Load Balancing	CPU Optimization	All settings
Publishing	Application	Site defaults are replicated. Other settings (name, description, icon, etc.) are global and are common to all sites
Publishing	Shortcuts	All settings
Publishing	Extensions	All settings
Publishing	Licensing	All settings
Publishing	Display	All settings
Publishing	Filtering (all types except Gateway)	All settings
Universal Printing	Universal printing	Printer renaming
Universal Printing	Printer drivers	All settings
Universal Printing	Fonts management	All settings
Universal Scanning	WIA	Scanner renaming
Universal Scanning	TWAIN	Scanner renaming
Universal Scanning	TWAIN > TWAIN applications	Scanning applications
Connection	Authentication	All settings

Connection	Settings	All settings
Connection	Multi-factor authentication	All settings
Connection	Allowed devices	All settings
Reporting	Reporting engine	Reporting engine type
Reporting	Engine specific settings	All settings

To replicate Site settings to all other sites, select **Farm** > <site>> **Settings** and then select the **Replicate settings** option (at the bottom of the **Auditing** tab). Please note that this option is disabled if you have just one Site in the Farm.

Overriding Site Replicated Settings

If an administrator who has permissions to enable or disable replication settings makes a change to a specific setting, such setting is replicated to all other sites. If an administrator has access to a particular Site only, upon modifying Site settings which have been replicated, the replicated settings are overridden and the option **Replicate Settings** is automatically cleared, therefore such settings will no longer be replicated to other sites.

Managing Licensing Site

The Licensing Site should always be online even if you have other sites in your Farm. If your Licensing Site goes offline, your other sites can still use the maximum number of individual licenses included in your subscription but only for a period of 72 hours. During this time, you need to do one of the following:

- Restore your Licensing Site.
- Promote a different Site to be the Licensing Site in the Farm (see below for instructions).

Please note that if the Licensing Site is offline from 48 to 72 hours and back online three times per month, you will be required to re-activate it using your Parallels RAS licensing key after the third time.

To promote a secondary Site to be the Licensing Site in the Farm:

- 1** In the RAS Console, navigate to **Farm** > **Farm**.
- 2** In the right pane select a Site and then click **Tasks** > **Set as licensing Site**.
- 3** You will be asked to activate the new Licensing Site using your Parallels RAS license. Follow the instructions and activate the Site.

Managing administrator accounts

You can have more than one administrator in Parallels RAS. At least one administrator (called the root administrator) must be present at all times. Other administrators can be given the following roles:

- **Root administrator.** Has full permissions to manage a Parallels RAS Farm.
- **Power administrator.** Has most permissions granted by default, but can be configured to have limited permissions to manage certain sites or categories.
- **Custom administrator.** Has no permission by default and can be granted specific permission to view or modify very specific areas or objects in the Parallels RAS Farm.

Read on to learn how to create and manage administrator accounts.

Adding an administrator account

To add an administrator account to the Parallels RAS Farm:

- 1 In the RAS Console, navigate to **Administration > Accounts**.
- 2 Click the **Tasks** drop-down menu and choose **Add** (or click the **[+]** icon).
- 3 The **Account Properties** dialog opens.
- 4 Click the **[...]** button next to the **Name** field. In the **Select User or Group** dialog, select a user or a group.
- 5 Specify an email address and mobile phone number. Both fields are optional and are disabled if the account specified in the **Name** field is a group.
- 6 In the **Permissions** drop-down list select a role to assign to the administrator:
 - **Root administrator.** Grants the administrator full permissions to manage the Farm.
 - **Power administrator.** Grants the administrator full permissions by default but allows you to limit them if needed. To grant or deny specific permissions, click the **Change Permissions** button. For additional info, see **Administrator Account Permissions** (p. 51).
 - **Custom administrator.** This role doesn't have any permissions by default and allows you grant very specific permissions for a particular category, area, or object in the RAS Console. See **Administrator Account Permissions** (p. 51) for details.
- 7 In the **Receive system notifications via** drop-down list, select **Email** to send all system notifications to the specified email address, or select **None** to disable email system notifications for this account.
- 8 Click **OK** to add the new administrator account to the Farm.

Modifying an administrator account

To modify an account, select it in the list and click **Tasks > Properties**. This opens the **Account Properties** dialog where you can modify the account information.

To enable or disable an account, select or clear the **Enable account** option at the top of the **Account Properties** dialog.

Administrator account permissions

To set permissions for a RAS administrator, do the following:

- 1 In the RAS Console, navigate to **Administration > Accounts**.
- 2 Select an administrator in the list and click **Tasks > Properties**.

When you click the **Change Permissions** button in the **Administrator Properties** dialog, the following happens depending on what is selected in the **Permissions** field:

- **Root administrator.** The **Change Permission** button is disabled because the root administrator always has full permissions.
- **Power administrator.** The **Account Permissions** dialog opens. In the left pane, select one or more sites for which to grant permissions to the administrator. In the right pane, select specific permissions. See the **Power administrator permissions** subsection below for details.
- **Custom administrator.** A different **Account Permissions** dialog opens where you can set custom permissions. Compared to the **Power administrator** role (see above), this option allows you to grant any permission (view, modify, add, etc.) for entire categories or specific areas or objects in the RAS Console. If a Custom administrator doesn't have permissions to even view a category or tab page, they will not even appear in the RAS Console. Using the **Custom administrator** role, you can limit permissions to one or more very specific tasks. For details, see **Custom administrator permissions** below.

Power administrator permissions

The following permissions can be set for a **Power administrator**:

- **Allow viewing of site information.** Whether the administrator can view the Site information.
- **Allow site changes.** Permissions to modify the following categories: **Site, Load Balancing, Universal Printing, Universal Scanning**. This option is disabled if the **Allow viewing of Site information** option is cleared.
- **Allow session management.** Permission to manage running sessions. This option is disabled if the **Allow viewing of site information** option is cleared.
- **Allow publishing changes.** Permission to modify the **Publishing** category.
- **Allow connection changes.** Permission to modify the **Connection** category.
- **Allow viewing of RAS reporting.** Permission to view reports generated by RAS Reporting.

- **Allow client management changes.** Permission to modify the **Device Manager** category.

In the **Global permission** area, set the following:

- **Allow viewing of policies.** Whether to allow the administrator to view the **Policies** category.
- **Allow policies changes.** Whether to allow the administrator to modify the **Policies** category.

Custom administrator permissions

To set custom administrator permissions, you must be either a root administrator or a power administrator with the "Allow site changes" permission granted.

When you first create an administrator of this type, they will have no permissions. To add permissions, select a Site in the left pane and then click the **Change permissions** button. The **Account Permissions** dialog opens. In the dialog, select a permission type in the left pane.

The permission types are:

- **RD Session Hosts.** The **RD Session hosts** tab in **Farm > RD Session hosts**.
- **RD Session hosts groups.** The **Groups** tab in **Farm > RD Session hosts**.
- **Remote PCs.** The **Farm > Remote PCs** view.
- **Gateways.** The **Farm > Gateways** view.
- **Publishing Agents.** The **Farm > Publishing Agents** view.
- **HALB.** The **Farm > HALB** view.
- **Themes.** The **Farm > Themes** view.
- **Publishing.** The entire **Publishing** category.
- **Connection.** The entire **Connection** category.
- **Device manager.** The entire **Device manager** category.
- **Certificates.** The **Farm > Site > Certificates** subcategory.

After you select a permission type, you can set the actual permissions in the right pane. Different permission types may have different sets of permissions. The following list describes all available permissions:

- **View.** View only.
- **Modify.** View and modify.
- **Add.** View, modify, and add new objects (e.g. servers).
- **Delete.** View, modify, and delete an object.

- **Control.** View and control an object. This permission enables the **Tasks > Control** menu (where available), which includes enable and disable logons, cancel pending reboot, install RDS role, reboot, and some other options. Also enables power operations (start, stop, etc., where available).
- **Manage sessions.** View and manage sessions.

The lower portion of the right pane lists individual objects (e.g. servers) if the selected permission type has them. Here, you can set individual permissions for a specific object (not the entire tab, for instance, which otherwise would include all available objects).

The **Global permissions** options at the top of the right pane enables all permissions for all objects for the selected permission type.

Clone permissions

As a root administrator (or a power administrator with sufficient privileges), you can apply (clone) permissions of an existing administrator account to another existing account. This way, you can configure permissions for one account and then quickly apply the same configuration to all other accounts that require them.

To clone permissions, select a source administrator account and click **Tasks > Clone permissions**. In the dialog that opens, select a destination account (or multiple accounts) and click **OK**.

Delegate permissions

There could be a situation when a power administrator needs to grant some permissions to a custom administrator. This cannot be done by modifying permissions because power administrators cannot manage administrator accounts directly. Instead, they can delegate some of their own permissions in a given Site to a custom administrator of their choice.

For example, if a power administrator wants the custom administrator to be able to manage a particular RD Session Host, he/she selects that host in the RAS Console and click **Tasks > Delegate permissions**. This opens a dialog where the administrator can select a custom administrator and specify which permissions (view, modify, etc.) that administrator should have. The **Tasks > Delegate permissions** menu option is available for many objects, such as RD Session Hosts, VDI providers, guest VMs (desktops), and some others. If the menu is not available for an object, it means that this functionality is not available for objects of this type.

Managing administrator accounts

To view existing administrator accounts, select the **Administration** category in the RAS Console. The **Accounts** tab lists existing accounts and their properties, including:

- **Group or user name.** Account name, which can be a user or group name.
- **Type.** Account type. Can be one of the following: **User**, **Group**, **Group User**. The **User** and **Group** are self-explanatory. The **Group User** is a user who receives Parallels RAS administrative permissions via a group membership. When you initially add a group to the list of Parallels RAS administrators, its members are not displayed on the **Accounts** tab. As soon as a member of the group logs in to Parallels RAS, the account name is added to the list of administrators as a **Group User** and remains there. Note that you cannot change Parallels RAS permissions for such an account individually outside the group permissions.
- **Permissions.** A security role assigned to an administrator.
- **Email.** Email address.
- **Mobile.** Mobile phone number.
- **Group.** Group name. This column has a value for Group Users only (see the **Type** column description above).
- **Last Modification By.** The name of the user who modified this account in Parallels RAS the last time.
- **Changed On.** The last account modification date.
- **Created By.** The name of the user who created this account in Parallels RAS.
- **Created On.** The date when this account was added to Parallels RAS.
- **ID.** Internal Parallels RAS ID.

Modifying an account

To modify an account:

- 1 Right-click an account and choose **Properties** in the context menu.
- 2 Use the **Administrator Properties** dialog to modify the necessary information. For more info, see **Adding an Administrator Account** (p. 50).

Handling locked objects

When an administrator is working with an object (e.g. a tab in the RD Session Host properties dialog), the object is locked for all other administrators. Therefore, upon trying to access a locked object, an administrator will be alerted with an error that the object is locked and will be denied access to it.

A root administrator (but not power or custom administrator) can release a locked object as follows:

- 1 On the **Administration > Accounts** tab, click the **Tasks** drop-down menu and choose **Show Sessions**.

- 2 In the **Sessions** dialog, select the administrator who is locking an object and then click the **Send Message** icon (at the top).
- 3 If the administrator doesn't reply and doesn't release the object, you have an option to click **Log Off**, which will log them off and will unlock the category.

Configure RAS Console idle sessions

If you have a number of administrators using the RAS Console to manage the same Farm, you can configure when an idle RAS Console session should be disconnected. By default, when an administrator opens the console and connects to a Farm but then forgets to log off and goes away, the session will stay active indefinitely possibly locking some of the categories for other administrators. You can change that by specifying the time period after which an idle session will be disconnected (thus unlocking the categories).

To configure idle sessions:

- 1 In the RAS Console, navigate to **Administration > Settings**.
- 2 Locate the **Miscellaneous** section (at the bottom) and choose a desired time period in the **Reset idle RAS Console session after** drop-down box.

When a session stays idle for close to the specified time period, the administrator (session owner) will be notified a few minutes in advance that the session is about to be disconnected. If the administrator chooses to stay connected, the time period is reset. If the administrator does nothing, the session will be disconnected when the time expires.

Using instant messaging

Parallels RAS administrators logged on to the same Farm can communicate with each other using a built-in instant messenger.

To use the instant messenger:

- 1 In the RAS Console, select the **Administration** category.
- 2 Expand the drop-down menu next to your name (top-right corner of the console screen) and click **Chat**.
- 3 The **Parallels Remote Application Server Chat** window opens.

To send a message:

- 1 Type the message text in the lower input panel.
- 2 In the **Logged on administrators** list box, select a specific administrator or **All** to send the message to an individual or all logged on administrators.
- 3 Click **Send**.

Your message history is displayed in the **Messages** panel. To clear the history, click **Clear All**.

You can also view the chat history listing all messages between all administrators (not just your own messages). To do so, select the **Administration** node in the console and then select the **Chat History** tab.

Joining Customer Experience Program

Parallels Customer Experience Program helps us to improve the quality and reliability of Parallels RAS. If you accept to join the program, we will collect information about the way you use Parallels RAS. We will not collect any personal data, like your name, address, phone number, or keyboard input.

To join the program:

- 1** In the RAS Console, select the **Administration** category.
- 2** In the right pane, click the **Settings** tab.
- 3** Select the **Participate in the Customer Experience Program** option.

After you join the program, CEP will automatically start to collect information about how you use Parallels RAS. Data collected from you and other participants is combined and thoroughly analyzed to help us improve Parallels RAS.

RAS Publishing Agent

RAS Publishing Agent provides load balancing of published applications and desktops. A RAS Publishing Agent is automatically installed on a server on which you install Parallels RAS and is designated as the primary Publishing Agent. Each Site must have a primary RAS Publishing Agent but can also have secondary Publishing Agents added to it. The purpose of a secondary Publishing Agent is to ensure that users do not experience any interruption of the service due to possible failure of the primary RAS Publishing Agent. This chapter describes how to add RAS Publishing Agents to a Site and how to configure them.

In This Chapter

Configuring RAS Publishing Agents	57
Secondary Publishing Agents	59
Managing Secondary Publishing Agents.....	61
Using computer management tools.....	63

Configuring RAS Publishing Agents

To view RAS Publishing Agents installed in a Site, navigate to **Farm > <Site> > Publishing Agents** in the RAS Console. The installed Publishing Agents are listed on the **Publishing Agents** tab in the right pane.

A Site must have at least the primary Publishing Agent installed, which is marked so in the **Priority** column. You can also add secondary agents to a Site for redundancy (described in the section that follows this one).

To modify the configuration of a Publishing Agent, select it and then click **Tasks > Properties** (or right-click > **Properties**). The **Properties** dialog opens where you can modify the following:

- **Enable Server in site:** Enables or disables the Publishing Agent. The option is enabled for secondary Publishing Agents only. It is disabled for the primary Publishing Agent.
- **Server:** Specifies the FQDN or IP address of the server that hosts the Publishing Agent. To automatically resolve IP address to FQDN, enable the global **Name Resolution** option. For details, see **Host Name Resolution** (p. 430).
- **IP:** Specifies the server IP address. Click the **Resolve** button to obtain the IP address automatically using the FQDN specified in the **Server** field. This IP address is used so that multiple Publishing Agents share information in real time.

- **Alternate IPs:** Specifies one or more alternate IP addresses separated by a semicolon. These addresses will be used if RAS Secure Client Gateways fail to connect to the RAS Publishing Agent using its FQDN or the address specified in the **IP** field. This can happen, for example, if Gateways are connecting from a network which is not joined to Active Directory.
- **Description:** A user-defined description.
- **Standby:** If selected, puts a secondary Publishing Agent into a standby mode. This means that no agent will connect to this Publishing Agent until another Publishing Agent goes offline. This option is enabled automatically for any new secondary Publishing Agent in excess of the three agents that already exist. It is not recommended to have more than three active Publishing Agents because it may degrade system performance. Using this option you can have more than three agents, but have them in standby mode until they are needed. For more information, see **Secondary Publishing Agents** (p. 59).

When done making the changes, click **OK** and then click **Apply** in the main RAS Console window.

The **Tasks** drop-down menu on the **Publishing Agents** tab has the following items:

- **Add.** Adds a RAS Publishing Agent to the Site. See the section that follows this one for the information on how to add secondary Publishing Agents.
- **Upgrade all Agents.** Upgrades agents to the current version. The item is disabled if all agents are up to date.
- **Tools.** Gives you access to a set of standard server management tools.
- **Troubleshooting.** The **Check agent** menu item verifies that the Publishing Agent is functioning properly. It opens a dialog where you can see the verification results and optionally install (or uninstall) the Publishing Agent. The **Logging** menu item allows you to configure logging and retrieve or clear log files. For more information, see **Logging** (p. 451).
- **Promote to primary.** Promotes a secondary Publishing Agent to primary. The current primary becomes a secondary Publishing Agent.
- **Refresh.** Refreshes the **Publishing Agents** list.
- **Delete.** Deletes a secondary Publishing Agent from the Site. To delete the primary Publishing Agent, you first need to promote a secondary Publishing Agent to primary.
- **Settings audit.** Opens the **Settings Audit** dialog where you can view the changes that were done to the Publishing Agent. For more information, see **Settings Audit** (p. 435).
- **Move up** and **Move down.** Changes the priority of a secondary Publishing Agent (moves it up or down in the priority list).
- **Properties.** Opens the Publishing Agent **Properties** dialog (see above).

RAS Publishing Agents overview

In addition to the Publishing Agent editor described above, you can also see the summary about the available RAS Publishing Agents.

To do so:

- 1 In the RAS Console, navigate to the **Farm** > <Site> .
- 2 The available RAS Publishing Agents are displayed in the **Publishing Agents** group on the **Site Info** tab.
- 3 To go to the Publishing Agents editor, right-click a RAS Publishing Agent and choose **Show in the editor**.

For additional info, see **Sites in the RAS Console** (p. 45).

Secondary Publishing Agents

A secondary Publishing Agent is added to a Site for redundancy. This way if the primary Publishing Agent fails, the secondary Publishing Agent is still available to handle the requests. Publishing Agents work in active/active manner to ensure high availability. In case of a Publishing Agent failure, the next agent is always ready to handle the load. In general, the N+1 redundancy approach should be used per Site. Note that for auto-promotion you shouldn't have more than three Publishing Agents (auto-promotion is described later in this section).

When you have one more secondary Publishing Agents installed, the runtime data is replicated on each agent, so if any service fails, the downtime is reduced to a minimum. In addition, any active Publishing Agent is used for authentication purposes with both the AD and any 2nd level authentication provider used.

The primary Publishing Agent performs the same tasks as secondary Publishing Agents but has additional responsibilities. It manages certain processes that must be managed by a single Publishing Agent. The following table lists processes managed by the primary Publishing Agent and secondary Publishing Agents:

Process	Primary Publishing Agent	Secondary Publishing Agents
Monitor PAs (counters)	Yes	Yes
Monitor RD Session Hosts (counters)	Yes	Yes
Monitor VDI providers (counters)	Yes	Yes
Monitor RDS Sessions (reconnection)	Yes	Yes
Monitor Deployed RDS applications	Yes	Yes
Monitor VDI session (reconnections)	Yes	Yes
Manage system settings	Yes	No
Send licensing information & heart beat	Yes	No

Process and send CEP information	Yes	No
Send information to reporting server	Yes	No
Manage RDS scheduler	Yes	No
Reporting engine information	Yes	Future versions
Shadowing	Yes	Future versions
Send email notifications	Yes	No

As a demonstration of how load distribution between multiple Publishing Agents works, consider the following example:

- Suppose we have two Publishing Agents: PA1 (primary) and PA2 (secondary).
- Suppose we also have 10 RD Session Hosts: RDS1, RDS2 ... RDS10

The resulting load will be distributed as follows:

- RDS1, RDS2 ... RDS4 will use PA1 as their preferred Publishing Agent.
- RDS5, RDS6 ... RDS10 will use PA2 as their preferred Publishing Agent.

Planning for secondary Publishing Agents

RAS Publishing Agents running on the same Site communicate with each other and share the load. The amount of data being transmitted from one agent to another is quite large, so a reliable high-speed communication channel must be ensured (e.g. a subnetwork can be configured for Publishing Agent communications).

When adding a secondary Publishing Agent to a Site, you specify an IP address for it. Make sure that the IP addresses of all agents belong to the same network segment. The port that Publishing Agents use to communicate with each other is TCP 20030.

There's no physical limit to how many Publishing Agents you can add to a Site. However, the best results are achieved with only two-three agents present. The three-agent scenario is highly recommended, especially when you have VDI providers and want to enable high availability for VDI (p. 142). Adding more than two secondary Publishing Agents to a Site may have a reverse effect and actually degrade the system performance. Note that this does not apply to secondary Publishing Agents in standby mode, which is explained in **Configuring RAS Publishing Agents** (p. 57).

Adding a secondary RAS Publishing Agent to a Site

To add a secondary Publishing Agent:

- 1 In the RAS console, navigate to **Farm** > <Site> > **Publishing Agents**.
- 2 Click the **Tasks** drop-down menu and choose **Add** to launch the **Add RAS Publishing Agent** wizard.
- 3 The **Server** field specifies the FQDN or IP address of the server that hosts the RAS Publishing Agent. To automatically resolve IP address to FQDN, enable the global **Name Resolution** option. For details, see **Host Name Resolution (p. 430)**.
- 4 The **IP** field specifies the server IP address. Click the **Resolve** button to obtain the IP address automatically using the FQDN specified in the **Server** field.
- 5 The **Alternative IPs** field specifies one or more alternative IP addresses, separated by a semicolon. These addresses will be used if RAS Secure Client Gateways fail to connect to the RAS Publishing Agent using its FQDN or the address specified in the **IP** field. This can happen, for example, if Gateways are connecting from a different network, which is not joined to Active Directory.
- 6 Select the **Install a gateway with a publishing agent** option if you also want to install a RAS Secure Client Gateway on the specified server. If you select this option, you may also select the **Add an SSL certificate and enable HTML5 Gateway** option (for more info, see **Configure HTML5 Client (p. 74)**).
- 7 Select the **Add Firewall Rules** option to automatically configure the firewall on the server. See **Port Reference** for details.
- 8 Click **Next**.
- 9 On the next page, click **Install** to install the RAS Publishing Agent on the server. The **Installing RAS Redundancy Service** dialog opens.
- 10 Select the server on which the RAS Publishing Agent is to be installed and click **Install**.
- 11 Click **Done**.
- 12 Click **OK** to add the server to the Farm.

Managing Secondary Publishing Agents

Enabling or disabling a secondary Publishing Agent

To enable or disable a secondary Publishing Agent in a Site, select it in the **Publishing Agents** list and then select or clear the check box at the beginning of the row.

Changing the secondary Publishing Agent priority

Each secondary Publishing Agent is given a priority. To change the priority, select a secondary Publishing Agent and use the "Up arrow" and "Down arrow" icons (or **Tasks > Move up, Move down**) to move it up or down the list. The higher the agent is in the list, the higher the priority.

Promoting a secondary Publishing Agent to primary

If the primary Publishing Agent cannot be recovered, you can promote a secondary Publishing Agent to primary as follows:

- 1 Open the RAS Console on the Publishing Agent server that you would like to promote (all required files are automatically installed when a server is added to a Site as a secondary Publishing Agent).
- 2 Select the **Farm** category and navigate to the **Publishing Agents** node.
- 3 Select the Publishing Agent and then click **Tasks > Promote to primary**.
- 4 Click **OK** once the process is finished.

Configuring auto-promotion

If the primary Publishing Agent goes offline, you will need to promote a secondary Publishing Agent to take its place. The auto-promotion feature can do it automatically after a specified time period.

By default, auto-promotion is turned off. To enable it, do the following:

- 1 In the RAS Console, navigate to **Farm > <Site> > Publishing Agents**.
- 2 Select the **Auto-promotion** tab in the right pane.
- 3 Select the **Enable auto-promotion** option and specify the time period after which the next secondary Publishing Agent should be promoted to primary. The time period can be set between 15 minutes and 72 hours (the default value is 30 min).
- 4 Select the **Enable fallback** option if you want the original Publishing Agent to become primary again should it go back online. For the Licensing Site, this eliminates license activation if fallback happens within 72 hours. The license activation countdown is always displayed in the RAS Console, so the administrator can check if the original primary Publishing Agent recovers within this time period or not. If the original agent goes back online after the 72-hour period (and if the Farm has been already reactivated), it will become a secondary Publishing Agent.

Note: To enable auto-promotion, you need at least three active Publishing Agents in a Site. If you have less than three, the auto-promotion is ignored.

Please also note that auto-promotion must be disabled if you have a single Site with Publishing Agents split across different locations with bad WAN links. If there's no link between Publishing Agent located remotely, the third Publishing Agent acts as a witness to prevent split-brain.

When auto-promotion takes place, the RAS administrator will receive notifications via email about the following events:

- A secondary Publishing Agent has been promoted to primary.
- Auto-promotion of a secondary Publishing Agent has failed.
- Auto-promotion failback completed.

Deleting a secondary Publishing Agent

To delete a secondary Publishing Agent, select it in the list and then click **Delete** in the **Tasks** drop-down menu.

Using computer management tools

You can perform standard computer management tasks on a server hosting the RAS Publishing Agent right from the RAS Console. These include Remote Desktop Connection, remote PowerShell, Computer Management, Service Management, Event Viewer, IPconfig, Reboot, and others. To access the **Tools** menu, select a server, click **Tasks > Tools** and choose a desired tool. For requirements and usage information, see **Computer Management Tools** (p. 431).

CHAPTER 6

RAS Secure Client Gateway

RAS Secure Client Gateway tunnels all Parallels RAS data on a single port. It also provides secure connections and is the user connection point to Parallels RAS.

At least one RAS Secure Client Gateway must be installed and configured in every Site. Note that if a Site is joined as Tenant to RAS Tenant Broker, RAS Secure Client Gateway is not needed. For details, see **RAS Multi-Tenant Architecture** (p. 299).

Multiple gateways can exist depending on your requirements. Read this chapter to learn how to add, configure, and manage RAS Secure Client Gateways.

In This Chapter

Overview	64
Adding a RAS Secure Client Gateway	66
Manually adding a RAS Secure Client Gateway	66
Checking the RAS Secure Client Gateway status	67
Configuring a RAS Secure Client Gateway.....	67
Gateway tunneling policies	80
Configure logging	81
Viewing Gateway summary and metrics	81
Using computer management tools.....	82

Overview

You need to install at least one RAS Secure Client Gateway for Parallels RAS to work. You can add additional Gateways to a RAS Site to support more users, load-balance connections, and provide redundancy.

Installing a RAS Secure Client Gateway on a dedicated server

If you are installing a RAS Secure Client Gateway on a dedicated server, you can also install the Parallels RAS console on the same server. The console will have limited functionality but will allow you to perform some important management operations on the Gateway, including:

- Setting the Gateway operation mode (normal or forwarding, see below for details).
- Assigning a RAS Publishing Agent that will manage the Gateway.

- Setting the Gateway communication port.
- Viewing the Gateway information, such as host OS version, Parallels RAS version, available IP addresses, and other.

The RAS Console in such an installation scenario (when connected to the local computer, not the RAS Farm) will only have two categories that you can select in the left pane: **Gateway** and **Information**. To manage the Gateway settings, select **Gateway** and then click **Change Ownership** in the right pane. To view the information select the **Information** category.

When the RAS console is connected to a Parallels RAS Farm (i.e. the server where RAS Publishing Agent is running), you can manage RAS Secure Client Gateways by navigating to **Farm > <Site> > Gateways**.

How a RAS Secure Client Gateway works

The following describes how a RAS Secure Client Gateway handles user connection requests:

- 1 A RAS Secure Client Gateway receives a user connection request.
- 2 It then forwards the request to the RAS Publishing Agent with which it's registered (the Preferred Publishing Agent setting by default).
- 3 The RAS Publishing Agent performs load balancing checks and the Active Directory security lookup to obtain security permissions.
- 4 If the user requesting a published resource has sufficient rights, the RAS Publishing Agent sends a response to the gateway which includes details about the RD Session Host the user can connect to.
- 5 Depending on the connection mode, the client either connects through the gateway or disconnects from it and then connects directly to the RD Session Host server.

RAS Secure Client Gateway operation modes

RAS Secure Client Gateway can operate in one of the following modes:

- **Normal Mode.** A RAS Secure Client Gateway in normal mode receives user connection requests and checks with the RAS Publishing Agent if the user making the request is allowed access. Gateways operating in this mode can support a larger number of requests and can be used to improve redundancy.
- **Forwarding Mode.** A RAS Secure Client Gateway in forwarding mode forwards user connection requests to a preconfigured gateway. Gateways in forwarding mode are useful if cascading firewalls are in use, to separate WAN connections from LAN connections and make it possible to disconnect WAN segments in the event of issues without disrupting the LAN.

Note: To configure the forwarding mode, a Parallels RAS Farm must have more than one RAS Secure Client Gateway.

Planning for high availability

When adding RAS Secure Client Gateways to a Site, the N+1 redundancy should be configured to ensure uninterrupted service to your users. This is a general rule that also applies to other Parallels RAS components, such as Publishing Agents or RD Sessions Hosts.

Adding a RAS Secure Client Gateway

To add a RAS Secure Client Gateway to a Site, follow these steps:

- 1 In the RAS Console, navigate to **Farm** > <Site> > **Gateways**.
- 2 With the **Gateways** tab selected in the right pane, click **Tasks** > **Add** to start the **Add RAS Secure Client Gateway** wizard.
- 3 Enter the server FQDN or IP address (or click the [...] button to select a server from the list). To automatically resolve IP address to FQDN, enable the global **Name Resolution** option. For details, see **Host Name Resolution (p. 430)**.
- 4 Select the gateway mode from the **Mode** drop down menu.
- 5 If you selected the **Forwarding** mode in the step above, select the destination gateway in the **Forward To** drop-down list. You can also select a specific IP address in the **On IP** drop-down list if the Gateway server has more than one.
- 6 Select the **Add an SSL certificate and enable HTML5 Gateway** option to automatically create a self-signed certificate, enable SSL, and enable HTML5 support. For more info, please see **Configure HTML5 Client (p. 74)**.
- 7 Select the **Add Firewall Rules** to automatically configure the firewall on the server hosting the gateway. See **Port Reference** for details.
- 8 Click **Next**.
- 9 On the next page, click **Install** to start the RAS Secure Client Gateway installation.
- 10 Click **Done** when the installation is finished.

Manually adding a RAS Secure Client Gateway

To manually install a RAS Secure Client Gateway and add it to the Farm, follow these steps:

- 1 Log into the server where you'll be installing the RAS Secure Client Gateway using an administrator account.
- 2 Copy the Parallels RAS installation file (`RASInstaller.msi`) to the server and double click it to launch the installation wizard.
- 3 Follow the onscreen instruction and proceed to the installation type page. Select **Custom** and click **Next**.

- 4 Click on **RAS Secure Client Gateway** in the feature tree and select **Entire Feature will be installed on local hard drive**.
- 5 Ensure that all other components in the selection tree are cleared and click **Next**.
- 6 Click **Install** to start the installation.
- 7 When the installation is completed, click **Finish** to close the wizard.
- 8 Open the RAS Console and specify the RAS Publishing Agent that will manage the gateway.

Checking the RAS Secure Client Gateway status

To check the status of a RAS Secure Client Gateway, right-click it in the list and then click **Check Status** in the context menu. The **RAS Secure Client Gateway Information** dialog opens.

The dialog displays the gateway information, including:

- **Server:** The name of the server on which the gateway is installed.
- **Gateway:** The gateway verification status (e.g. Verified).
- **Version:** The gateway software version number. The version number must match the Parallels RAS version number.
- **OS Type:** Operating system type and version.
- **Status:** Display the current RAS Secure Client Gateway status. If the status indicates a problem (e.g. the gateway did not reply or the gateway software version is wrong), click the **Install** button to push install the gateway software on the server. Wait for the installation to complete and check the status again.

Configuring a RAS Secure Client Gateway

To configure a RAS Secure Client Gateway:

- 1 In the RAS console, navigate to **Farm > <Site> > Gateways**.
- 2 In the right pane, right-click a gateway and click **Properties**.
- 3 The **RAS Secure Client Gateway Properties** dialog opens.

Read on to learn how to configure the RAS Secure Client Gateway properties.

Enable or disable a Gateway

A RAS Secure Client Gateway is enabled by default. To enable or disable a gateway, open the **RAS Secure Client Gateway Properties** dialog and select or clear the **Enable RAS Secure Client Gateway in site** option on the **General** tab.

Set public address

The **Public address** field on the **General** tab specifies a public FQDN or IP address of the Gateway. This setting is used by the Preferred routing functionality for redirecting a client connection. Please see **Configuring preferred routing** (p. 250).

Set IP addresses for client connections

IP addresses for incoming client connections for a gateway are specified on the **General** tab of the **RAS Secure Client Gateway Properties** dialog. RAS Secure Client Gateway recognizes both IPv4 and IPv6. By default, IPv4 is used.

You can specify the following IP options:

- **Use IP version:** Select the IP version(s) to use.
- **IP(s):** Specify one or more IP addresses separated by a semicolon, or click **Resolve** to resolve the IP address automatically. These are the available addresses on the gateway server. To specify IP addresses that should be used for client connections, use the **Bind to IP** section (see below).
- **Bind to IP:** Use this section to specify on which IP address (or addresses) the gateway will listen for client connections. You can select a specific address or **<All available addresses>**, in which case all of the IP addresses specified in the **IP(s)** field will be used.
- **Remove system buffers for:** These fields (one for each IP version) can be used when the connection between the gateway and the Parallels Client has a high latency (such as the Internet). This option will optimize traffic for better experience on the Parallels Client side. You can select a specific address, all available addresses, or none. What this option will do is delay the internal socket to match the performance of the external socket. If the internal network is fast and the external is slow, RDP detects the fast internal socket and sends a lot of data. The problem is that this data cannot be sent fast enough from the gateway to the Client, thus ending up with a bad user experience. Enabling this option will optimize the data exchange.

Site defaults (Gateways)

RAS Secure Client Gateway **Properties** dialog consists of tabs, each containing their own specific set of options. All tabs, except **Properties**, have one common option **Inherit default settings**. When you select this option, all fields on a tab are grayed out and the settings are inherited from Site defaults. To view (and modify if necessary) Site default properties for gateways, click the **Site Defaults** link, which is available on all tabs mentioned above. The link opens the **Site default properties** dialog. You can also open this dialog by clicking **Tasks > Site defaults** while on the **Farm > Site > Gateways** tab.

The subsequent sections describe individual tabs and available options in the gateway **Properties** dialog.

Gateway mode and forwarding settings

A RAS Secure Client Gateway can operate in normal and forwarding modes (p. 64). To set the desired mode and configure related settings click the **Mode** tab in the **RAS Secure Client Gateway Properties** dialog.

Using Site defaults

To use Site default settings, click the **Inherit default settings** option. To specify your own settings, clear the option. For more info, see **Site Defaults (Gateways)** (p. 68).

Setting the normal mode

To set the normal mode, in the **Gateway mode** drop-down list, select **Normal**.

The **Forward requests to HTTP Server** option allows you to forward requests that do not belong to RAS Secure Client Gateways (gateways handle HTML5 traffic, Wyse, and URL scheme). To specify multiple servers, separate them with a semicolon. An HTTP server can be specified using an IPv6 address if necessary. Please note that the HTTP server must support the same IP version as the browser making the request.

The **Preferred Publishing Agent** drop-down list allows you to specify a RAS Publishing Agent that the gateway should connect to. This is helpful when Site components are installed in multiple physical locations communicating through WAN. You can decrease network traffic by specifying a more appropriate Publishing Agent. For the gateway to select a Publishing Agent automatically, select the **Automatic** option.

Setting the forwarding mode

To configure the forwarding mode, in the **Gateway mode** drop-down list, select **Forwarding**.

Specify (or select) one or more forwarding gateways in the **Forwarding RAS Secure Client Gateway(s)** field.

Note: The forwarding mode allows you to forward data to a gateway listening on IPv6. It is recommended that forwarding gateways are configured to use the same IP version.

Gateway network options

The **Network** tab is used to configure RAS Secure Client Gateway network options.

Using Site defaults

To use Site default settings, click the **Inherit default settings** option. To specify your own settings, clear the option. For more info, see **Site Defaults (Gateways)** (p. 68).

Configuring network

By default RAS Secure Client Gateway listens on TCP ports 80 and 443 to tunnel all Parallels RAS traffic. To change the port, specify a new port in the **RAS Secure Client Gateway Port** input field.

RDP port 3389 is used for clients that require basic load balanced desktop sessions. Connections on this port do not support published resources. To change the RDP port on a gateway select the **RDP Port** option and specify a new port. When setting your own port, please make sure that the port number does not conflict with the standard "RD Session Host Port" setting.

Note: If RDP port is changed, the users need to append the port number to their connection string in the remote desktop client (e.g. [ip address]:[port]).

Broadcast RAS Secure Client Gateway Address. This option can be used to switch on the broadcasting of the gateway address, so Parallels Clients can automatically find their primary gateway. The option is enabled by default.

Enable RDP UDP Data Tunneling. To enable UDP tunneling on Windows devices, select this option (default). To disable UDP tunneling, clear the option.

Device Manager Port. Select this option to enable management of Windows devices from the **Device Manager** category. The option is enabled by default.

Enable RDP DOS Attack Filter. When selected, this option denies chains of uncompleted sessions from the same IP address. For example, if a Parallels Client initiates multiple successive sessions with each session waiting for the user to provide credentials, Parallels RAS will deny further attempts. The option is enabled by default.

SSL/TLS encryption

The traffic between Parallels RAS users and a RAS Secure Client Gateway can be encrypted. The **SSL/TLS** tab allows you to configure data encryption options.

Using Site defaults

To use Site default settings, click the **Inherit default settings** option. To specify your own settings, clear the option. For more info, see **Site defaults (Gateways)** (p. 68).

Enforcing HSTS

The **Configure** button in the HSTS section allows you to enforce HTTP Strict Transport Security (HSTS), which is a mechanism that makes a web browser to communicate with the web server using only secure HTTPS connections. When HSTS is enforced for a RAS Secure Client Gateway, all web requests to it will be forced to use HTTPS. This specifically affects the RAS HTML5 Gateway (p. 74), which can normally accept both HTTP and HTTPS requests.

When you click the **Configure** button, the **HSTS Settings** dialog opens where you can specify the following:

- **Enforce HTTP strict transport security (HSTS):** Enables or disables HSTS for the gateway.
- **Max-age:** Specifies the max-age for HSTS, which is the time (in our case in months) that the web browser should remember that it can only communicate with the gateway using HTTPS. The default (and recommended) value is 12 months. Acceptable values are 4 to 120 months.
- **Include subdomains:** Specifies whether to include subdomains (if you have them).
- **Preload:** Enables or disables HSTS preloading. This is a mechanism whereby a list of hosts that wish to enforce the use of SSL/TLS on their Site is hardcoded into a web browser. The list is compiled by Google and is used by Chrome, Firefox, Safari, Internet Explorer 11, and Edge browsers. When HSTS preload is used, a web browser will not even try to send a request using HTTP, but will use HTTPS every time. Please also read the important note below.

Note: To use HSTS preload, you have to submit your domain name for inclusion in Chrome's HSTS preload list. Your domain will be hardcoded into all web browser that use the list. **Important:** Inclusion in the preload list cannot easily be undone. You should only request inclusion if you are sure that you can support HTTPS for your entire Site and all its subdomains in the long term (usually 1-2 years).

Please also note the following requirements:

- Your website must have a valid SSL certificate. See **SSL server configuration** (p. 73).
- All subdomains (if any) must be covered in your SSL Certificate. Consider ordering a Wildcard Certificate.

Configuring SSL

By default, a self-signed certificate is assigned to a RAS Secure Client Gateway when the gateway is installed. Each RAS Secure Client Gateway must have a certificate assigned and the certificate should be added to Trusted Root Authorities on the client side to avoid security warnings.

SSL certificates are created on the Site level using the **Farm > Site > Certificates** subcategory in the RAS Console. Once a certificate is created, it can be assigned to a RAS Secure Client Gateway. For the information about creating and managing certificates, refer to the **SSL Certificate Management** (p. 258) chapter.

To configure SSL for a gateway:

- 1 Select the **Enable SSL on Port** option and specify a port number (default is 443).
- 2 In the **Accepted SSL Versions** drop-down list, select the SSL version accepted by the RAS Secure Client Gateway.
- 3 In the **Cipher Strength** field, select a desired cipher strength.
- 4 In the **Cipher** field, specify the cipher. A stronger cipher allows for stronger encryption, which increases the effort needed to break it.
- 5 The **Use ciphers according to server preference** option is ON by default. You can use client preferences by disabling this option.
- 6 In the **Certificates** drop-down list, select a desired certificate. For the information on how to create a new certificate and make it appear in this list, see the **SSL Certificate Management** (p. 258) chapter.

The **<All matching usage>** option will use any certificate configured to be used by gateways. When you create a certificate, you specify the "Usage" property where you can select "Gateway", "HALB", or both. If this property has the "Gateway" option selected, it can be used with a gateway. Please note that if you select this option, but not a single certificate matching it exists, you will see a warning and will have to create a certificate first.

Encrypting Parallels Client connection

By default, the only type of connection that is encrypted is a connection between a gateway and backend servers. To encrypt a connection between Parallels Client and the gateway, you also need to configure connection properties on the client side. To do so, in Parallels Client, open connection properties and set the connection mode to **Gateway SSL**.

To simplify the Parallels Client configuration, it is recommended to use a certificate issued by a well-known third-party Trusted Certificate Authority. Note the Windows certificate store is used by some web browsers (Chrome, Edge etc.) when connecting to the RAS HTML5 Gateway.

Parallels Clients configuration

In case the certificate is self-signed, or the certificate issued by Enterprise CA, Parallels Clients should be configured as follows:

- 1 Export the certificate in Base-64 encoded X.509 (.CER) format.
- 2 Open the exported certificate with a text editor, such as notepad or WordPad, and copy the contents to the clipboard.

To add the certificate with the list of trusted authorities on the client side and enable Parallels Client to connect over SSL with a certificate issued from an organization's Certificate Authority:

- 1 On the client side in the directory "C:\Program Files\Parallels\Remote Application Server Client\" there should be a file called `trusted.pem`. This file contains certificates of common trusted authorities.
- 2 Paste the content of the exported certificate (attached to the list of the other certificates).

Securing RDP-UDP connections

A Parallels Client normally communicates with a RAS Secure Client Gateway over a TCP connection. Recent Windows clients may also utilize a UDP connection to improve WAN performance. To provide the SSL protection for UDP connections, DTLS must be used.

To use DTLS on a RAS Secure Client Gateway:

- 1 On the **SSL/TLS** tab, make sure that the **Enable SSL on Port** option is selected.
- 2 On the **Network** tab (p. 69), make sure that the **Enable RDP UDP Data Tunneling** option is selected.

The Parallels Clients must be configured to use the **Gateway SSL Mode**. This option can be set in the **Connections Settings > Connection Mode** drop-down list on the client side.

Once the above options are correctly set, both TCP and UDP connections will be tunneled over SSL.

SSL server configuration

When configuring RAS Secure Client Gateway to use SSL encryption, you should pay attention to how the SSL server is configured to avoid possible traps and security issues. Specifically, the following SSL components should be rated to determine how good the configuration is:

- The certificate, which should be valid and trusted.
- The protocol, key exchange, and cipher should be supported.

The assessment may not be easy to perform without specific knowledge about SSL. That's why we suggest that you use the SSL Server Test available from Qualys SSL Labs. This is a free online service that performs an analysis of the configuration of an SSL web server on the public Internet. To perform the test on a RAS Secure Client Gateway, you may need to temporarily move it to the public Internet.

The test is available at the following URL: <https://www.ssllabs.com/ssltest/>.

You can read a paper from Qualys SSL Labs describing the methodology used in the assessment at the following URL: <https://github.com/ssllabs/research/wiki/SSL-Server-Rating-Guide>.

Configure HTML5 Gateway

The **HTML5** tab is used to configure the HTML5 Gateway.

Parallels HTML5 Gateway is a functionality built into RAS Secure Client Gateway that allows users to connect to Parallels RAS and open published resources from a web browser using the Parallels HTML5 Client. The client works similarly to a platform-specific Parallels Client, but does not require any additional software to be installed on users' computers or devices. All that users need is an HTML5-enabled web browser.

This section describes how to configure HTML5 Gateway in the Parallels RAS Console. For the information about how to use it, please refer to the **Parallels HTML5 Client** chapter (p. 345).

Note: To use HTML5 Client, SSL must be enabled on a RAS Secure Client Gateway. When enabling the client, please verify that SSL is enabled on the **SLL/TLS** tab or on your network load balancer. Please also note that the **HTML5** tab is only available if the gateway mode is set to "Normal". For more information, see **Gateway mode and forwarding settings** (p. 69).

To configure HTML5 Gateway, click the **HTML5** tab in the RAS Secure Client Gateway properties dialog and then set the options described in the subsequent sections.

For the information on how to configure the HTML5 Client URL and how to access the client from a web browser, please **Web request load balancing** (p. 78).

Using Site defaults

To use Site default settings on the **HTML5** tab, click the **Inherit default settings** option. To specify your own settings, clear the option. For more info, see **Site defaults (Gateways)** (p. 68).

Enable or disable HTML5 Client

To enable or disable HTML5 Client, select or clear the **Enable HTML5 Client** option. This disables the HTML5 Gateway, so users will no be able to run the HTML5 Client.

Client settings

The **Client** section allows you to specify application launch methods and other HTML5 Client settings.

- **Launch sessions using:** When a user tries to open a resource from the HTML5 Client web page, the resource can open right in the web browser or it can be launched in a platform-specific Parallels Client installed on the user's computer (e.g. Parallels Client for Windows). This option specifies which client will be used. Compared to HTML5 Client, platform-specific Parallels Client includes a richer set of features and provides end users with a better overall user experience. Select one of the following:

- a Launch apps in browser only (HTML5 only)** — Users can run remote applications and desktops using Parallels HTML5 Client only. Use this option if you don't want your users to install a platform-specific Parallels Client.
 - b Launch apps with Parallels Client** — Users can run remote applications and desktops in Parallels Client only. When a user connects to Parallels RAS using Parallels HTML5 Client, they will be asked to install the platform-specific Parallels Client before they can launch remote applications and desktops. A message will be displayed to the user with a link for downloading the Parallels Client installer. After the user installs Parallels Client, they can still select a remote application or desktop in Parallels HTML5 Client but it will open in Parallels Client instead.
 - c Launch apps with Parallels Client & fallback to HTML5** — Both Parallels Client and a browser (HTML5) can be used to launch remote applications and desktops. Parallels Client will be the primary method; Parallels HTML5 Client will be used as a backup method if a published resource cannot be launched in Parallels Client for any reason. A user will be informed if a resource couldn't be opened in Parallels Client and will be given a choice to open it in the browser instead.
- **Allow users to select a launch method:** If selected, users will be able to choose whether to open remote applications in a browser or in Parallels Client. You can enable this option only if the **Launch session using** option (above) is set to **Launch apps in Parallels Client and fallback to HTML5** (i.e. both methods are allowed).
 - **Allow opening applications in a new tab:** If selected, users will be able to open remote applications in a new tab in a web browser.
 - **Use Pre Windows 2000 login format:** Enables legacy (pre-Windows 2000) login format.
 - **Allow embedding of Parallels HTML5 Client into other web pages:** If selected, the Parallels HTML5 Client web page can be embedded in other web pages. Please note that this may be a potential security risk due to a practice known as clickjacking.
 - **Allow file transfer command:** Enables file transfer in a remote session. To enable file transfer, select this option and click the **Configure** button. In the dialog that opens, select **Client to server only** (transfer files from client to server only), **Server to client only** (transfer files from server to client only), **Bidirectional** (transfer files in both directions). For more information, see **Configuring Remote File Transfer** (p. 408).
 - **Allow clipboard command:** Enables clipboard operations (copy/paste) in a remote session. To enable the clipboard, select this option and click the **Configure** button. In the dialog that opens, select **Client to server only** (copy/paste from client to server only), **Server to client only** (copy and paste from server to client only), **Bidirectional** (copy and paste in both directions). For more information about using the clipboard, see **Using the Remote Clipboard** (p. 364).
 - **Allow cross-origin resource sharing:** Enables cross-origin resource sharing (CORS). To enable CORS, select this option and click the **Configure** button. In the dialog that opens, specify one or more domains for which access to resources should be allowed. If you don't specify any domains, the option will be automatically disabled. In the **Browser cache time** field, specify for how long the end-user's browser will cache a resource.

Network load balancers access

The **Network Load Balancers access** section is intended for deployment scenarios where third-party front-end load balancers such as Amazon Web Services (AWS) Elastic Load Balancers (ELBs) are used. It allows you to configure an alternate hostname and port number to be used by the Network Load Balancer (NLB). This is needed to separate hostnames and ports on which TCP and HTTPS communications are carried out because AWS load balancers don't support both specific protocols over the same port.

The following options are available:

- **Use alternate hostname:** Select this option and specify an alternate hostname. When the alternate hostname is enabled, all platform-specific Parallels Clients will use this hostname to connect to the RAS Farm or Site.
- **Use alternate port:** Select this option and specify an alternate port number. The port must not be used by any other component in the RAS Farm or Site. To reset the port number to the default value, click **Default**. When the alternate port is enabled, all platform-specific Parallels Clients will use this port to connect to the RAS Farm or Site. Note that RDP sessions in HTML5 Client will still be connecting to the standard SSL port (443).

Note: Please note that using an alternate host or port is not suitable in a multi-tenant environment as Tenant Broker RAS Secure Client Gateways are shared between Tenants, which would require different configurations.

In addition, the AWS Application Load Balancer (ALB), which handles HTTP/s traffic required by the Parallels HTML5 Client, only supports specific cookies that are usually automatically generated. When a load balancer first receives a request from a client, it routes the request to a target and generates a cookie named `AWSALB`, which encodes information about the selected target. The load balancer then encrypts the cookie and includes it in the response to the client. When sticky sessions are enabled, the load balancer uses the cookie received from the client to route the traffic to the same target, assuming the target is registered successfully and is considered healthy. By default, Parallels RAS uses its own ASP.NET cookie named `_sessionId`, however in this case you must customize the cookie specifying the mentioned AWS cookie for sticky sessions. This can be configured using the **Web cookie** field on the **Web Requests** tab. Please note that this functionality is available in Parallels RAS 17.1 or newer.

Wyse ThinOS support

To publish applications from the Parallels RAS to thin clients using the Wyse ThinOS, select the **Enable Wyse ThinOS support** option on the **Wyse** tab.

Note: The Wyse tab is only available if the gateway mode is set to normal. See **Gateway mode and forwarding settings** for more info (p. 69).

By enabling this option, the RAS Secure Client Gateway will act as a Wyse broker. You need to make sure that DHCP option 188 on your DHCP server is set to the IP address of this gateway for thin clients that will be booting via this gateway. Once the DHCP server is configured, click the **Test** button to verify the DHCP server settings.

The **Do not warn if server certificate is not verified** option can be selected (enabled) if a Wyse device shows an SSL warning when connecting to a RAS Secure Client Gateway because the hostname does not match the certificate. When the option is selected, the Gateway will send Wyse clients the following parameters in the wnos.ini file: SecurityPolicy=low TLSCheckCN=no, which will disable SSL checks. Note that the option is not required if a certificate has the following:

- The CNAME set to the FQDN of the RAS Secure Client Gateway.
- The SAN set to the RAS Secure Client Gateway IP address.

Note that if you use a custom wnos.ini in "C:\Program Files (x86)\Parallels\ApplicationServer\AppData\wnos" folder on Gateway, the Gateway will not send the SSL check parameters.

If you configure DHCP option 188 to set the broker address to a given Gateway, you can verify this by clicking the **Test** button.

Gateway security

You can allow or deny user access to a gateway based on a MAC address. This can be accomplished using the **Security** tab in the **RAS Secure Client Gateway Properties** dialog.

Using Site defaults

To use Site default settings, click the **Inherit default settings** option. To specify your own settings, clear the option. For more info, see **Site defaults (Gateways)** (p. 68).

Configuring security

To configure a list of allowed or denied MAC addresses, click the **Security** tab and select one of the following options:

- **Allow all except.** All devices on the network will be allowed to connect to the gateway except those included in this list. Click **Tasks > Add** to select a device or to specify a MAC address.
- **Allow only.** Only the devices with the MAC addresses included in the list are allowed to connect to the gateway. Click **Tasks > Add** to select a device or to specify a MAC address.

Please note that the Gateway MAC address filtering is based on ARP, so client and server must be on the same network for the filtering to work. It does not work across network boundaries.

Web request load balancing

Note: The **Web** tab is only available if the gateway mode is set to normal. See more in **Gateway mode and forwarding settings** (p. 69).

The **Web** tab allows you to tweak settings necessary for load balancing in certain scenarios. Here you can specify a redirection URL for web requests and a session cookie name to maintain persistence between a client and a server.

Redirection URL

An original web request can reach the gateway one of the following two ways:

- The request is sent directly to the gateway over the local network using its IP address or FQDN. For example, `https://192.168.10.10`.
- The request is sent to a HALB device that load-balances this and other gateways in the Farm. The HALB device often faces the Internet (i.e. located in DMZ) and so its DNS name can be used in the original request URL. For example, `https://ras.msp.com`. The HALB device is then distributes the request to a gateway.

When the gateway receives the web request, it takes the URL specified on the **Web** tab and sends it back to the web browser for redirection.

Technically, you can enter any URL here, and the original web request will be redirected to that URL. The primary purpose of this field, however, is to give end users an easy way to access the HTML5 Client from their web browsers. Here's how it works:

- 1** A user enters the Load Balancer DNS name in a web browser. For example, `https://ras.msp.com`.
- 2** The Load Balancer receives the request and distributes it to the least-busy RAS Secure Client Gateway for processing.
- 3** The gateway receives the original URL and replaces it with the URL specified in the **Default URL** field. See the **Default URL format** subsection below.
- 4** The replaced URL is then sent back to the web browser, which uses it to open the HTML5 Client login page.

Default URL format

The default URL format is the following:

```
https://%hostname%/RASHTML5Gateway
```

- The `%hostname%` variable is automatically replaced with the name of the server that received the original request, which in our example is the Load Balancer DNS name. If you wish, you can replace the variable with a specific host name or IP address (e.g. this or some other gateway). For example, `https://192.168.5.5/RASHTML5Gateway`. If you do this, the web requests will always be forwarded to the specified host and will open the HTML5 Client on it. Hard-coding a host may not be very practical, but you can do this nevertheless.
- `RASHTML5Gateway` is a constant and is the path to the HTML5 Client login page.

In our example, the resulting URL that the web browser will use to access the HTML5 Client is the following:

```
https://ras.msp.com/RASHTML5Gateway
```

The fact is, a user could simply use the above URL from the start, but thanks to the redirection feature, users only need to enter the server DNS name (or FQDN/IP-address on the local network) instead of the entire URL.

Opening a specific HTML5 Client Theme

HTML5 Client Themes is a feature that allows you to custom design the HTML5 Client look and feel for different groups of users. Themes are described in detail in the **Parallels HTML5 Client** chapter (p. 74).

The default web request URL opens the default Theme. To make it open a specific Theme, add the Theme name at end of the URL as follows:

```
https://%hostname%/RASHTML5Gateway/?theme=<theme-name>
```

where `<theme-name>` is the name of a Theme without brackets or quotes.

For users to open a specific Theme, the URL that they enter in a web browser must contain the Theme name, but in this case the format is as simple as the following:

```
https://<server-name>/<theme-name>
```

Using our Load Balancer DNS name example from above, the URL may look like the following:

```
https://ras.msp.com/Theme-E1
```

For additional information, please see **HTML5 Client Theme Settings > URLs** (p. 347).

Web cookie

The Web cookie field is used to specify a session cookie name. RAS HTML5 session persistence is normally set by the user IP address (source addressing). If you can't use source addressing in your environment (e.g. your security policy doesn't allow it), you can use the session cookie to maintain persistence between a client and a server. To do so, you need to set up a load balancer that can use a session cookie for persistence. The default cookie name is ASP.NET_SessionId. Note that if you are using Amazon Web Services (AWS) or other third-party load balancers, you may need to specify their own cookie name. See **Network load balancers access** (p. 76) for more information.

Gateway tunneling policies

Tunneling policies can be used to load balance connections by assigning a group of RD Session Hosts to a specific RAS Secure Client Gateway or RAS Secure Client Gateway IP address.

To configure tunneling policies, navigate to **Farm** > <Site> > **Gateways** and then click the **Tunneling Policies** tab in the right pane.

The <**Default**> policy is a preconfigured rule and is always the last one to catch all non-configured gateway IP addresses and load balance the sessions between all servers in the Farm. You can configure the <**Default**> policy by right-clicking it and then clicking **Properties** in the context menu.

Adding a new Tunneling Policy

To add a new policy:

- 1 Click **Tasks** > **Add**.
- 2 Select a gateway IP address.
- 3 Specify to which RD Session Host(s) the users connecting to that specific gateway should be forwarded. If you select **None** (no forwarding), read the **Restricting RDP access** section below.

Managing a Tunneling Policy

To modify an existing Tunneling Policy, right-click it and then click **Properties** in the context menu.

Restricting RDP access

You can use tunneling policies to restrict RDP accesses through the RAS Secure Client Gateway port. To do so, on the **Tunneling Policies** tab, select the **None** option at the bottom of the tab (this is the default setting in a new Parallels RAS installation). By doing so, you are restricting native MSTSC from accessing the gateway through its port (the default port is 80). As a result, when someone tries to use MSTSC at IP-address:80, the access will be denied. Same will happen for an RDP connection from a Parallels Client.

There are a couple of reasons why you would want to restrict RDP access. The first one is when you want your users to connect to the RAS Farm using the Parallels RAS connection only, but not RDP. The second reason is *to prevent a DDoS attack*.

A common indication of a DDoS attack taking place is when your users cannot login to a RAS Farm for no apparent reason. If that happens, you can look at the Controller.log file (located on the RAS Publishing Agent server, path C:\ProgramData\Parallels\RASLogs) and see that it is full of messages similar to the following:

- [I 06/0000003E] Mon May 22 10:37:00 2018 - Native RDP LB Connection from Public IP x.x.x.x, Private IP xxx.xxx.xx.xx, on gateway xxx.xxx.xx.xx, Using Default Rule
- [I 06/00000372] Mon May 22 10:37:00 2018 - CLIENT_IDLESERVER_REPLY UserName hello@DOMAIN, ClientName , AppName , PeerIP xxx.xxx.xx.xx, GatewayIP xxx.xx.x.xx, Server , Direct , desktop 0
- [I 05/0000000E] Mon May 22 10:37:00 2018 - Maximum amount of sessions reached.
- [I 06/00000034] Mon May 22 10:37:00 2018 - Resource LB User 'hello' No Servers Available!
- [W 06/00000002] Mon May 22 10:37:00 2018 - Request for "" by User hello, Client , Address xxx.xxx.xx.xx, was not served error code 14.

These messages tell us that a DDoS attack is in progress on the RDP port. By restricting RDP access through gateway tunneling polices, you can prevent this from happening.

Configure logging

A RAS Secure Client Gateway is monitored and logs are created containing relevant information. To configure logging and retrieve or clear existing log files, right-click a gateway, choose **Troubleshooting > Logging** in the context menu, and then click **Configure**, **Retrieve**, or **Clear** depending on what you want to do. For the information on how to perform these tasks, see the **Logging** (p. 451) section.

Viewing Gateway summary and metrics

You can view the summary information for all available RAS Secure Client Gateways in one place as follows:

- 1 In the RAS Console, select the **Farm** category and then select the **Site** node in the middle pane.
- 2 The available RAS Secure Client Gateways are displayed in the **Gateways** group in the right pane.
- 3 To go to the main Gateway view/editor, right-click a server and choose **Show in the Editor**.

You can also view the detailed information about a RAS Secure Client Gateway by navigating to **Information > Site** in the Parallels RAS Console. The information on this page includes general information, such as OS version, RAS version, Gateway mode, as well as the information about various types of connections, sessions, cached sockets, and threads.

Using computer management tools

You can perform standard computer management tasks on server hosting the RAS Secure Client Gateway right from the RAS Console. These include Remote Desktop Connection, PowerShell, Computer Management, Service Management, Event Viewer, IPconfig, Reboot, and others. To access the **Tools** menu, select a server, click **Tasks** (or right-click) > **Tools** and choose a desired tool. For requirements and usage information, see **Computer management tools** (p. 431).

CHAPTER 7

RD Session Hosts

RD Session Hosts are used to host published resources (applications, desktops, documents, etc.) in a Parallels RAS Farm. Read this chapter to learn how to add, configure, and administer RD Session Hosts.

In This Chapter

RD Session Host types.....	83
Add an RD Session Host.....	84
Planning for high availability	87
Viewing RD Session Hosts	87
Configuring an RD Session Host	89
Grouping and cloning RD Session Hosts	108
Using scheduler.....	112
Managing RD Session Host sessions	116
Managing logons.....	116
Using computer management tools.....	117
Publishing from an RD Session Host	118
Publishing containerized applications	123
Viewing published resources	125

RD Session Host types

Beginning with Parallels RAS v16.5, you can create and add to a RAS Farm the following types of RD Session Hosts:

- Individual servers. These can be physical boxes or virtual machines treated as physical servers.
- Virtual machines (VMs) created from a template, which is a part of RAS Virtual Desktop Infrastructure (VDI). The main advantage of using VMs is the ability to create as many of them as you require from a single template. RD Session Hosts based on a template are described in the **Grouping and cloning RD Session Hosts** section (p. 108).

Considering that template is a part of RAS VDI, some aspects of creating, provisioning, and managing RD Session Hosts based on a template differ from the regular RD Session Hosts (individual servers). For example, template-based hosts are added to a Farm automatically from a group, not manually by the administrator. There are some other differences which are described in various sections of this chapter. When reading these sections, please pay attention to whether or not a particular functionality applies to RD Session Hosts based on a template.

Add an RD Session Host

RD Session Host requirements

An RD Session Host must have the Remote Desktop Services (RDS) role installed. You can install RDS right from the RAS Console, as described later in this section.

To push install the RAS RD Session Host Agent on a server, the following requirements must be met:

- The firewall must be configured on the server to allow push installation. Standard SMB ports (139 and 445) need to be open. See also **Port reference** for the list of ports used by Parallels RAS.
- SMB access. The administrative share (\\server\c\$) must be accessible. Simple file sharing must be enabled.
- Your Parallels RAS administrator account must have permissions to perform a remote installation on the server. If it doesn't, you'll be asked to enter credentials of an account that does.
- The RD Session Host should be joined to an AD domain. If it's not, the push installation may not work and you will have to install the Agent on the server manually. See **Installing the Agent manually** section (p. 86).

Note: The rest of this section applies to regular RD Session Hosts only. If you are looking for the information on how to add an RD Session Host based on a template, see **Grouping and cloning RD Session Hosts** (p. 108).

Add an RD Session Host

To add an RD Session Host to a Site:

- 1** In the RAS Console, navigate to **Farm > Site > RD Session Hosts**.
- 2** Click **Tasks > Add**. This opens the **Add RD Session Hosts** wizard. Note that you can also open the wizard from the **Start** category as describe in **Set up a basic Parallels RAS Farm** (p. 29).

- 3 On the first page, select a server or type a server FQDN or IP address in the edit box and then click the plus-sign icon to add the server to the list. Note that if you enter the server name (hostname or FQDN), it will be used as the primary method of connecting to this server from other RAS components and clients. If you enter the IP address, it will be automatically resolved to FQDN, but only if the global option to resolve to FQDN is enabled. To see the current setting of this global option, click **Tools > Options** on the main menu. In the **Options** dialog, examine the **Always attempt to resolve to fully qualified domain name (FQDN) when adding hosts** option. When the option is selected, the IP address of every server/component in the RAS Farm is always resolved to FQDN. When the option is cleared, whatever is specified for a server (IP address or name) is used to communicate with a server. This makes a difference in deployments where an IP address cannot be used to access a server, such as when a server is hosted in the cloud. For more information, see **Host name resolution** (p. 430).
- 4 Click **Next**.
- 5 On the next page, specify the following options:
 - **Add firewall rules.** Add firewall rules required by Parallels RAS in Windows running on the server. See **Port reference** for details.
 - **Install RDS role.** Install the RDS role on the server if it's not installed. You should always select this option.
 - **Enable Desktop Experience.** Enable the Desktop Experience feature in Windows running on the server. This option is enabled only if the Install RDS role option (above) is selected. The option applies to Windows Server 2008 R1/R2 and Windows 2012 R1/R2 on which the Desktop Experience feature is not enabled by default.
 - **Restart server if required.** Automatically restart the server if necessary. You can restart the server manually if you wish.
 - **Add server(s) to group.** Add the server (or servers) to a group. Select the desired group in the list box located below this option. Groups are described in detail in the **Grouping RD Session Hosts** (p. 108) section. If you are just learning how to use this wizard, you can skip this option.
- 6 Click **Next**.
- 7 The next page allows you to add users and groups to the Remote Desktop Users group in Windows running on the server. This is necessary for your Parallels RAS users to be able to access published resources hosted by an RD Session Host. To specify users and/or groups, select the option provided and then click the **[+]** icon. In the **Select Users or Groups** dialog, specify a user or a group and click **OK**. The selected user/group will be added to the list on the wizard page.

Note: If you skip this step and your users are not members of the Remote Desktop Users group on an RD Session Host, they will not be able to access published resources. If you already used (or want to use later) standard Windows tools to add users to the Remote Desktop Users group, you can skip this page.

- 8 Click **Next**.

- 9** The **User profile** page allows you to select a technology to manage user profiles. You can select from **User profile disk** or **FSLogix**. User profile disks are virtual hard disks that store user application data on a dedicated file share. Microsoft FSLogix Profile Container is the preferred Profile Management solution as the successor of Roaming Profiles and User Profile Disks (UPDs). It is set to maintain user context in non-persistent environments, minimize sign-in times and provide native profile experience eliminating compatibility issues. For complete instructions, please see **User profile** (p. 96).
- 10** The **Optimization** page allows you to specify settings that will be used to optimize Windows on the RD Session Host for best performance in a Parallels RAS environment. You can select Windows components, services, and other options that will be disabled, removed, or optimized to ensure a more efficient, streamlined, and improved delivery of virtual apps and desktops. For the complete description, please see **Optimization** (p. 102).
- 11** On the next page, review the settings and click **Next**.
- 12** The **Install RAS RD Session Host Agent** dialog opens. Follow the instructions and install the agent. When the installation is finished, click **Done** to close the dialog.
- 13** Back in the wizard, click **Finish** to close it.

If you would like to verify that the RD Session Host has been added to the Farm, click the **Farm** category (below the **Start** category in the left pane of the Parallels RAS Console window) and then click **RD Session Hosts** in the navigation tree (the middle pane). The server should be included in the **RD Session Hosts** list. The **Status** column may display a warning message. If it does, reboot the server. The **Status** column should now say, "OK", which means that your RD Session Host is functioning properly.

Read on to learn how to publish an application from an RD Session Host (p. 36)

Installing the agent manually

You may need to install the RAS RD Session Host Agent manually if the automatic push installation cannot be performed. For instance, an SMB share may not be available or the firewall rules may interfere with the push installation, etc.

Installing RAS RD Session Host Agent manually

- 1** Log in to the server where the RAS RD Session Host Agent is to be installed using an administrator account and close all other applications.
- 2** Copy the Parallels RAS installation file (`RASInstaller.msi`) to the server and double-click it to launch the installation.
- 3** Follow the onscreen instructions and proceed to the installation type page. Select **Custom** and click **Next**.
- 4** Click on **RAS RD Session Host Agent** and select **Entire Feature will be installed on local hard drive** from the drop-down menu.

- 5 Ensure that all other components are deselected and click **Next**.
- 6 Click **Install** to start the installation.
- 7 Click **Finish** once the installation is finished.

The RAS RD Session Host Agent doesn't require any configuration. Once the agent is installed, highlight the server name in the RAS Console and click **Troubleshooting > Check Agent** in the **Tasks** drop-down menu to update the server status.

Uninstalling RAS RD Session Host Agent

To uninstall RAS RD Session Host Agent from a server:

- 1 Navigate to **Start > Control Panel > Programs > Uninstall a Program**.
- 2 Find **Parallels Remote Application Server** in the list of installed programs.
- 3 If you don't have any other Parallels RAS components on the server that you want to keep, right-click **Parallels Remote Application Server** and then click **Uninstall**. Follow the instructions to uninstall the program. You may skip the steps below.
- 4 If you have other RAS components that you want to keep on the server, right-click **Parallels Remote Application Server** and then click **Change**.
- 5 Click **Next** on the Welcome page.
- 6 On the **Change, repair, or remove** page, select **Change**.
- 7 On the next page, select **Custom**.
- 8 Select **RAS RD Session Host Agent**, then click the drop-down menu in front of it, and click **Entire feature will be unavailable**.
- 9 Click **Next** and complete the wizard.

Planning for high availability

When adding RD Session Hosts to a Site, the N+1 redundancy approach should be used to ensure uninterrupted service to your users. This is a general rule that also applies to other Parallels RAS components, such as Publishing Agents, RAS Secure Client Gateways, or possibly VDI providers.

Viewing RD Session Hosts

To view the list of RD Session Hosts for the current Site:

- 1 In the RAS Console, navigate to **Farm > <Site-name> > RD Session Hosts**.
- 2 The available RD Session Hosts are displayed on the **RD Session Hosts** tab in the right pane.

You can filter the **RD Session Hosts** list as follows:

- 1 Click the magnifying glass icon, which is located on a toolbar above the list.
- 2 An extra row is displayed at the top of the list where you can type a string in one or more columns that will be used to filter the list.
- 3 For example, if you want to search for a server by its name, enter the text in the **Server** column. You can type the entire server name or the first few characters until a match is found. The list will be filtered as you type and only the matching server(s) will be displayed.
- 4 If you type a filter string in more than one column, they will be combined using the logical AND operator.
- 5 To remove the filter and display the complete list, click the magnifying glass icon again.
- 6 If you click the magnifying glass icon one more time, you'll see that the filter that you specified earlier is still there. To remove it completely, simply delete the filter string(s) from the column(s).

Viewing RD Session Host summary

In addition to the RD Session Hosts editor described above, you can also see the summary about the available RD Session Hosts. To do so:

- 1 In the RAS Console, select the **Farm** category and then select the **Site** node in the middle pane.
- 2 The available servers are displayed in the **RD Session Hosts** group in the right pane.
- 3 To go to the RD Session Host editor (described above), right-click a server and choose **Show in the Editor**.

For additional info, see **Sites in the RAS Console** (p. 45).

Available menu options

You can perform a number of tasks on the an RD Session Host using menus. To do so, click the **Tasks** drop-down menu and choose a desired option, or right-click a host and choose an option from the context menu.

Please note that not all menu options are available for RD Session Hosts based on a template. If an option is not available for this host type, it will be either disabled or hidden. These include:

- **Remove from group.** Hosts based on a template can only be removed from a group using the **Group Properties** dialog.
- **Assign to group.** Group assignment is performed automatically for template-based hosts.
- **Delete.** Deleting a host (which is a VM) can only be done on the template level (the **Guest VM List** dialog).

- **Properties.** RD Session Hosts of this type don't have individual properties. Some essential properties are inherited from **Default Server Properties** (see View and Modify RD Session Host Properties > Agent Settings (p. 90)).
- **Control** (logon commands). Drain mode is managed automatically by the group to which a template-based host belongs.

Configuring an RD Session Host

This section describes how to configure and manage an existing RD Session Host.

Read on to learn how to:

- Check RAS RD Session Host Agent status (p. 89)
- Change an RD Session Host Site assignment (p. 90)
- View and modify RD Session Host properties (p. 90)
- Configure logging (p. 107)

Check an RD Session Host Agent status

An RD Session Host must have RAS RD Session Host Agent installed in order to publish remote applications and desktop from it. In addition to this, Remote Desktop Services (formerly Terminal Services) must also be installed.

Normally when you add an RD Session Host to a Site, the RD Session Host Agent and Remote Desktop Services are installed by default. However, if you skipped the installation (or uninstalled the agent or RDS from the server), you can check their status and take appropriate actions if needed.

To check the status of RD Session Host Agent and RDS, do the following:

- 1** First, check the **Status** column in the **RD Session Hosts** list. The column should display "OK". If so, the Agent is installed and functioning properly. If not, read on.
- 2** In addition to the description, the **Status** column uses a color code to indicate the agent status as follows:
 - Red — not verified
 - Orange — needs update
 - Green — verified
- 3** Right-click a server and click **Troubleshooting > Check agent** in the context menu. The **Agent Information** dialog opens.
- 4** If the agent is not installed on the server, click the **Install** button and follow the instructions on the screen.

After the agent installation is complete, you may need to reboot the RD Session Host. You can do it right from the Parallels RAS Console by selecting the server and clicking **Tasks > Control > Reboot**.

Change RD Session Host Site assignment

You can assign an RD Session Host to a different Site in your Farm if needed. Please note that this functionality is only available if you have more than one Site in your Farm.

To change the Site assignment:

- 1 Right-click an RD Session Host and then click **Change Site** in the context menu. The **Change Site** dialog opens.
- 2 Select a Site in the list and click **OK**. The server will be moved to the **RD Session Hosts** list of the target Site (**Farm > <new-site-name> > RD Session Hosts**).

View and modify RD Session Host properties

Note: The information in this section does not apply to RD Session Hosts based on a template. Hosts of that type don't have individual properties and are managed on the template level. For more information, see **Grouping and cloning RD Session Hosts** (p. 108) and **Templates** (p. 147).

To configure an RD Session Host:

- 1 In the RAS Console, navigate to **Farm > <site> > RD Session Hosts**.
- 2 Select a server and click **Tasks > Properties**.
- 3 The server properties dialog opens where you can configure the RD Session Host properties.

The dialog is described in the subsections that follow this one:

- Using default settings (p. 91)
- General (p. 91)
- Agent settings (p. 91)
- User profile (p. 94)
- Optimization (p. 94)
- Desktop access (p. 95)
- RDP printer (p. 95)

Using default settings

The server properties dialog consists of tabs, each containing their own specific set of properties. All tabs, except **General**, have one common link **Site Defaults**, which allows you to view and modify Site default settings. If you want the properties on a particular tab to inherit default settings, select the **Inherit default settings** option. When you do, the default settings will be inherited from one of the following:

- **Group defaults** if the server is assigned to an RD Session Host group. Groups are described in **Grouping and cloning RD Session Host Servers** (p. 108).
- **Site defaults** if the server is not assigned to an RD Session Host group. Note that a group may also inherit Site defaults, but this can be overridden in the group properties dialog where you can specify custom settings for a group.

To view or modify Site default settings, click the **Site Defaults** link. This will open either the **Group default properties** or **Site default properties dialog** depending on whether the server is assigned to a group (see above). Note that each individual tab can inherit default settings independently from other tabs.

To specify custom settings for an RD Session Host, clear the **Inherit default settings** option and use the controls on a given tab to set the desired options.

General

Select or clear the **Enable Server in site** option to enable or disable the server. A disabled server cannot serve published applications and virtual desktops to clients.

Other elements on this tab are:

- **Server:** Specifies the server FQDN or IP address.
- **Description:** An optional server description.
- **Change Direct Address:** Select this option if you need to change the direct address that Parallels Client uses to establish a direct connection with the RD Session Host.

Agent settings

Each RD Session Host in a RAS Farm has an RAS RD Session Host Agent installed through which it communicates with other Parallels RAS components. Use the **Agent Settings** tab to configure the agent.

To use default settings, select the **Inherit default settings** option. See **Using default settings** (p. 91).

To configure the agent, set the options as described below.

Application session lingering

- **Disconnect active session after:** Specifies the amount of time each session remains connected in the background after the user has closed a remote application. This option is used to avoid unnecessary reconnections with the server.
- **Logoff disconnected session after:** This setting allows you to control how long it takes for a session to be logged off after it is marked as "disconnected".

Other settings

- **Port:** Specifies a different remote desktop connection port number if a non-default port is configured on the server.
- **Max sessions:** Specifies the maximum number of sessions.
- **Preferred Publishing Agent:** Select a Publishing Agent to which the RD Session Host should connect. This is helpful when Site components are installed in multiple physical locations communicating through WAN. You can decrease network traffic by specifying a more appropriate Publishing Agent.

Allow Client URL/Mail redirection

When a user tries to open a URL or an HTML Mailto link in a remote application, the link can be redirected to the client computer and open in a local default application (a web browser or email client) instead of an application on the remote host. To enable this functionality, select the option and click the **Configure** button. In the dialog that opens, select one of the following:

- **Replace Registered Application:** This option uses an alternative method of redirecting a link. It replaces the default web browser and mail client with "dummy" apps on the remote server side. By doing so, it can intercept an attempt to open a link and redirect it to the client computer.
- **Support Windows Shell URL namespace objects:** The Shell URL namespace objects support means that Parallels RAS can intercept actions in published applications that use the Shell namespace API to open links, which is a standard behavior in most applications. The ability to disable support for Shell URL namespace objects is for compatibility with older versions of Parallels RAS. You may disable this option if you want the behavior of an older version of Parallels RAS (RAS v16.2 or earlier).

Please note that you can configure a list of URLs that should never be redirected, even if the redirection is enabled. This can be done on the **Farm > Site > Settings > URL Redirection** tab. See more in **Site settings** (p. 433).

Enable Drag and drop

Allows you to set how the drag and drop functionality works in Parallels Clients. To enable drag and drop, select the option, click the **Configure** button and then select from the following:

- **Server to client only:** Drag and drop to a local application, but not in the opposite direction.
- **Client to server only:** Drag and drop to a remote application only.
- **Bidirectional:** Note that this option has changed since Parallels RAS 17.1. In the past, it was a checkbox that would enable or disable drag and drop which worked in the "Client to server only" mode. When upgrading from an older version of Parallels RAS, and if the checkbox was enabled, the "Client to server only" option is selected by default. If the option was disabled, the "Disabled" option will be set. You can change it to any of the new available options if you wish.

Note: At the time of this writing, the drag and drop functionality is only supported on Parallels Client for Windows and Parallels Client for Mac.

Allow 2XRemoteExec to send command to the client

Select this option to allow a process running on the server to instruct the client to deploy an application on the client side. Read more about 2XRemoteExec in the **Using RemoteExec** subsection at the end of this topic.

Use RemoteApp if available

Enable this option to allow use of remote apps for shell-related issues when an app is not displayed correctly. This feature is supported on the Parallels Client for Windows only.

Enable application monitoring

Enable or disable monitoring of applications on the server. Disabling application monitoring stops the WMI monitoring to reduce CPU usage on the server and network usage while transferring the information to RAS Publishing Agent. If the option is enabled, the collected information will appear in a corresponding RAS report. If the option is disabled, the information from this server will be absent from a report.

Allow file transfer command (HTML5 and Chrome clients)

Enables file transfer in a remote session. To enable file transfer, select this option and click the **Configure** button. For more information, see **Configuring remote file transfer** (p. 408).

Enable drive redirection cache

Improves user experience by making file browsing and navigation on redirected drives much faster. For details, see **Drive redirection cache** (p. 107).

Using 2XRemoteExec

2XRemoteExec is a feature that facilitates the servers ability to send commands to the client. This is done using the command line utility `2XRemoteExec.exe`. Command line options include:

Command Line Parameter	Parameter Description
-s	Used to run the 2XRemoteExec command in 'silent' mode. Without this parameter, the command will display pop up messages from the application. If you include the parameter, the messages will not be displayed.
-t	Is used to specify the timeout until the application is started. Timeout must be a value between 5000ms and 30000ms. Note that the value inserted is in 'ms'. If the timeout expires the command returns with an error. Please note that the application might still be started on the client.
-?	Shows a help list of the parameters that 2XRemoteExec uses.
"Path for Remote Application"	The Application that will be started on the client as prompted from the server.

2XRemoteExec examples:

The following command displays a message box describing the parameters that can be used.

```
2XRemoteExec -?
```

This command runs Notepad on the client.

```
2XRemoteExec C:\Windows\System32\notepad.exe
```

In this example, the command opens the `C:\readme.txt` file in the Notepad on the client. No message is shown and 2XRemoteExec would wait for 6 seconds or until the application is started.

```
2XRemoteExec C:\Windows\System32\notepad.exe "C:\readme.txt"
```

User profile

Use this tab to configure user profile settings.

To use default settings, select the **Inherit default settings** option. See **Using default settings** (p. 91).

For complete instructions about configuring user profiles, see **User profile** (p. 96).

Optimization

The **Optimization** tab allows you to specify settings that will be used to optimize the RD Session Host for best performance in a Parallels RAS environment. You can select Windows components, services, and other options that will be disabled, removed, or optimized to ensure a more efficient, streamlined, and improved delivery of virtual apps and desktops.

To use default settings, select the **Inherit default settings** option. See **Using default settings** (p. 91).

For the complete description, please see **Optimization** (p. 102).

Desktop access

The **Desktop Access** tab allows you to restrict remote desktop access to certain users.

To use default settings, select the **Inherit default settings** option. See **Using default settings** (p. 91).

By default, all users who have access to remote applications on an RD Session Host can also connect to the server via a standard RDP connection. If you want to restrict remote desktop access to certain users, do the following:

- 1 On the **Desktop Access** tab, select the **Restrict direct desktop access to the following users** option. If you have the **Inherit default settings** option selected, click the **Edit Defaults** link to see (and modify if needed) the default configuration. The rest of the steps apply to both the **Server Properties** and **Default Server Properties** dialogs.
- 2 Click the **Add** button.
- 3 Select the desired users. To include multiple users, separate them by a semicolon.
- 4 Click **OK**.
- 5 The selected users will appear in the list on the **Desktop Access** tab.

Users in this list will still be able to access remote applications using Parallels Client, but will be denied direct remote desktop access to this server.

Note: **Computer Configuration > Administrative Templates > Windows Components > Remote Desktop Services > Remote Desktop Session Host > Connection > Allow users to connect remotely using remote desktop services** must be set to **Not configured**, otherwise it takes precedence.

Please note that members of the Administrator group will still be able to connect to the remote desktop even if they are included in this list.

RDP printer

The **RDP Printer** tab allows you to configure the renaming format of redirected printers. The format may vary depending on which version and language of the server you are using.

To use default settings, select the **Inherit default settings** option. See **Using default settings** (p. 91).

The **RDP Printer Name Format** drop-down list allows you to select a printer name format specifically for the configured server.

Select the **Remove session number from printer name** and the **Remove client name from printer name** options to exclude the corresponding information from the printer name.

User profile

User profile is a collection of settings and application data associated with a specific user. In a non-persistent remote environment, such as Parallels RAS, user profiles must be maintained to provide consistent user experience. This is achieved by storing user profile data in a network location to minimize sign in times and optimize file I/O between host, client, and the profile storage.

Parallels RAS supports the following technologies to manage user profiles:

- **User profile disk:** [RD Session Hosts only] These are virtual hard disks that store user application data on a dedicated file share. This disk is mounted to the user session as soon as the user signs in to a session host. The disk is unmounted when the user logs out.
- **FSLogix Profile Container:** A remote profile solution for non-persistent environments. FSLogix Profile Container redirects the entire user profile to a remote location and maintains user context in non-persistent environments, minimizing sign-in times and providing native profile experience eliminating compatibility issues. FSLogix Profile Container is the preferred profile management solution as the successor of Roaming Profiles and User Profile Disks.

User profiles can be configured for the following:

- RD Session Hosts
- VDI
- Azure Virtual Desktop

User profile settings are configured for the above on the Site level (Site defaults) and can also be configured for individual components if the RAS administrator decides to use custom settings for a given component.

To configure user profile on the Site level, navigate to **Farm > Site**, click the **Tasks > Site defaults** menu and choose one of the following:

- **RD Session Host**
- **VDI**
- **AVD multi-session hosts**
- **AVD single-session hosts**

In a Site defaults dialog that opens, select the **User profile** tab. The user interface for configuring optimization is the same for all of the above. Please note that the **User profile disk** option is not available for VDI and Azure Virtual Desktop due to obsolescence.

The subsequent sections describe in detail how to configure the user profile functionality.

User Profile Disks

To configure User Profile Disks, specify the following settings:

- 1 When in the host "Properties" dialog, clear the **Inherit default settings** if you want to specify different settings for this host.
- 2 In the **Technology** section, select **User profile disk**.
- 3 In the drop-down list, select one of the following:
 - **Do not change**: Keep the current server settings (default).
 - **Enabled**: Enable the User Profile Disks functionality.
 - **Disabled**: Disable the functionality.
- 4 **Disk location**: If you selected **Enabled** in the previous step, specify a network location where the User Profile Disks should be created. Use the Microsoft Windows UNC format to specify a location (e.g. \\RAS\users\disks). Please note that the server must have full control permissions on the disk share.
- 5 **Maximum size**: Enter the maximum allowed disk size (in gigabytes).
- 6 **User profile disks data settings**: Click this button to open the **User Profile Disks Data Settings** dialog. In the dialog, you can specify which user folders should be stored on the user profile disk. Select one of the following:
 - **Store all user settings and data on the user profile disk**: All folders, except those specified in the exclusion list, will be stored on the user profile disk. To add or remove folders to/from the exclusion list, click the **[+]** or **[-]** buttons.
 - **Store only the following folders on the user profile disk**: Only folders specified in the inclusion lists will be stored on the user profile disk. There are two inclusion lists. The first one contains standard user profile folders (e.g. Desktop, Documents, Downloads, etc.) and allows you to select the folders that you want to include. The second list allows you to specify additional folders. Click the **[+]** or **[-]** buttons to add or remove folders.

Note that when you enable User Profile Disks, you need to restart the server for the changes to take effect.

FSLogix Profile Container

Note: If you have existing FSLogix Profile Containers and would like their configurations to be managed by Parallels RAS, please read additional instructions in **Configure managing existing profiles by Parallels RAS** (p. 100).

Supported FSLogix releases

Parallels RAS has been tested with FSLogix releases up to and including release 2105.

Configure installation method

Before you configure FSLogix Profile Container for a specific server or a template (described later in this guide), you need to configure the FSLogix installation method on the Site level as follows:

- 1 Navigate to **Farm > Site > Settings** and select the **Features** tab. Here you need to select a method that Parallels RAS will use to install FSLogix on individual hosts. You can select from one of the following:
 - **Install manually:** Select this option if you want to install FSLogix on every host yourself. If this option is selected, Parallels RAS will not attempt to install FSLogix on a host.
 - **Install online:** This option installs FSLogix on session hosts from the Internet. Select one of the supported FSLogix versions from the drop-down list or select **Custom URL** and specify a download URL. Click the **Detect latest** button to automatically obtain a URL of the latest FSLogix version.
 - **Install from a network share:** Select this option if you have the FSLogix installation files on a network share and specify its location.
 - **Push from RAS Publishing Agent:** This option allows you to upload the FSLogix installation archive to the RAS Publishing Agent server. When you enable FSLogix on a session host, it will be push installed on the host from the RAS Publishing Agent server.
- 2 When done, click **Apply** in the RAS Console to apply your changes to Parallels RAS.

Upgrade FSLogix

The dialog described above can also be used to upgrade FSLogix to a newer version. To upgrade, do one of the following:

- Select **Install online** and choose from one of the provided FSLogix builds or specify a custom URL. The **Detect latest** button obtains a URL for the latest stable FSLogix build.
- Download a new version from the Microsoft website, place it on a network share or upload it to the RAS Publishing Agent server and then select **Install from a network share** or **Push from RAS Publishing Agent**, whichever applies.

If FSLogix is already installed on one or more hosts and a new version of FSLogix becomes available when you do one of the above, FSLogix will be upgraded on hosts that have it installed. Note that if you specify a version that is earlier than the version installed on a host, then FSLogix will be downgraded.

Configure Site defaults and hosts for FSLogix

To configure Site defaults or individual hosts for FSLogix, do one of the following:

- For Site defaults, navigate to **Farm > Site** and click **Tasks > Site defaults > RD Session Hosts** (or **VDI** to configure defaults for VDI, or one of the **AVD** options to configure site defaults for Azure Virtual Desktop).

- To configure individual hosts, navigate to **Farm > Site > RD Session Hosts**. Right-click a host and choose **Properties**.
- When you add an RD Session Host to a Farm, the FSLogix settings are specified on the **User profile** page.

In the **Site defaults** or **Properties** dialog, select the **User profile** tab and specify the following options:

- 1** If you are in the host **Properties** dialog (or in a wizard where you add a new host or template), clear the **Inherit default settings** option if you want to specify different settings for this host.
- 2** In the **Technology** section, select **FSLogix**.
- 3** The **Deployment method** field shows the currently set deployment method as configured on the Site level (see the description above). You can click the **Change...** link and select a different method. Note that this will modify the Site setting, which will change it for all hosts in the Site.
- 4** In the **Settings** section, specify the following:
 - **Location type:** Select a location type for profile disks (SMB Location or Cloud Cache) and then specify one or more locations (see below).
 - **Location of profile disks:** Location(s) of profile disks. These are the locations of VHD(X) files (the VHDLocations setting in the registry as specified in the FSLogix documentation).
 - **Profile disk format:** Select from VHD or VHDX according to your requirements. VHDX is a newer format and has more features.
 - **Allocation type:** Select Dynamic or Full. This setting is used in conjunction with the **Default size** setting (see below) to manage the size of a profile. Dynamic causes the profile container to use the minimum space on disk, regardless of the allocated Default size. As a user profile is filled with more data, the amount of data on disk will grow up to the size specified in Default size, but will never exceed it.
 - **Default size:** Specifies the size of newly created VHD(X) in megabytes.
- 5** Click the **Additional settings** button to configure more settings:
 - **Users and Groups:** Specify include and exclude user and group lists. By default, Everyone is added to the FSLogix profile include list. If you want some user profiles remain local, you can add those users to the exclude list. Users and group can exist in both lists but exclude takes priority.
 - **Folders:** Specify include and exclude lists for folders. You can select from common folders or you can specify your own. Please note that folders must reside in user profile path.
 - **Advanced:** This tab allows you to modify advanced FSLogix registry settings. To modify a setting, select it and click **Tasks > Edit**. By default, the settings are disabled. To enable a setting, select the checkbox in front of its name. A description for each setting is provided in the RAS console. For further information regarding FSLogix Profile Containers configurations, visit <https://docs.microsoft.com/en-us/fslogix/profile-container-configuration-reference>.

Rebooting a host

When you enable FSLogix for a new host while running the wizard, no additional steps are necessary. On wizard completion, the host is rebooted and is added to the active load balancing. An existing host must be rebooted manually using the **Tasks > Tools > Reboot** menu option.

Configure managing existing profiles by Parallels RAS

This topic describes how to configure existing FSLogix Profile Containers to be managed by Parallels RAS. FSLogix Profile Container configuration defines how and where the profile is redirected. Normally, you configure profiles through registry settings and GPO. Parallels RAS gives you the ability to configure profiles from the Parallels RAS Console or RAS Management Portal without using external tools.

Before you begin

Before you configure FSLogix Profile Containers in Parallels RAS, make note of the following:

- You don't have to change the profiles themselves; existing profiles stay the same.
- You can keep using your existing FSLogix Profile Container locations, such as SMB network shares or Cloud Cache.

Preliminary steps

Perform the following preliminary steps:

- 1** Back up your existing profiles. It is highly unlikely that profile data can be lost or corrupted, but it is best practice to have a valid backup prior to any change in profile configuration.
- 2** Turn off the GPO configuration of FSLogix Profile Containers. This step is important because you cannot have both GPO and Parallels RAS management of FSLogix profiles enabled at the same time.
- 3** Before configuring FSLogix profiles for a server in a RAS Farm, make sure there are no user sessions running on the server. As a suggestion, you can make the transition in a maintenance window out of working hours.

Replicate GPO and FSLogix configuration

To configure existing FSLogix Profile Containers in Parallels RAS, you need to replicate your existing GPO to the FSLogix configuration in Parallels RAS. This can be done in the Parallels RAS Console or the Parallels RAS Management Portal.

To configure profiles in the RAS Console:

- 1 Follow the instruction from the **FSLogix Profile Containers** section (p. 97) and open the **User profiles** tab.
- 2 In the **Location of profile disks** list box, specify existing SMB or Cloud Cache locations where you keep your FSLogix profiles. Also, specify the profile disk format, allocation type, and default size.
- 3 Click the **Additional settings** button and configure the rest of FSLogix settings you may have on your servers, such as user exclusions, folder exclusions, and others.

To configure profiles in the RAS Management Portal:

- 1 Navigate to **Infrastructure > RD Session Hosts**.
- 2 Click a host in the list and then click **Properties**.
- 3 In the middle pane, click **User Profile**.
- 4 Specify the settings as described in steps above for the RAS Console.

Please note that at the time of this writing RAS Management Portal can only be used to configure RD Session Hosts to use FSLogix Profile Containers. For other host types, please use the desktop-based RAS Console.

Recommendations and testing

When performing steps in the previous section, do not configure multiple (or all) servers in a RAS farm right away. Begin with a single server (e.g. an RD Session Host) and then test it with a single user connection. After that, configure some other servers and test the same user logging in to multiple servers consecutively to confirm the profile is loaded and personalization is retained irrespective of a session host. If all is good, configure other servers, server groups, or Site defaults.

Your RAS users can now connect to Parallels RAS using pre-existing FSLogix Profile Containers, which are now managed centrally through Parallels RAS.

FSLogix antivirus exclusions

Make sure to configure the following antivirus exclusions for FSLogix Profile Container virtual hard drives. Make sure to check the following information with your security team.

Exclude files:

- %Programfiles%\FSLogix\Apps\frxdrv.sys
- %Programfiles%\FSLogix\Apps\frxdrvvt.sys
- %Programfiles%\FSLogix\Apps\frxccd.sys
- %TEMP%*.VHD
- %TEMP%*.VHDX

- %Windir%\TEMP*.VHD
- %Windir%\TEMP*.VHDX
- \\storageaccount.file.core.windows.net\share**.VHD
- \\storageaccount.file.core.windows.net\share**.VHDX

Exclude processes:

- %Programfiles%\FSLogix\Apps\frxccd.exe
- %Programfiles%\FSLogix\Apps\frxccds.exe
- %Programfiles%\FSLogix\Apps\frxsvc.exe

When configuring optimizations, you can specify files and processes to exclude in the Windows Defender ATP category. For more information, please see **Optimization** (p. 102).

Optimization

Beginning with version 18, Parallels RAS includes built-in automated optimization capabilities for RD Session Hosts, VDI, and Azure Virtual Desktop workloads. Different preconfigured optimizations for multi-session (such as RD Session Hosts) or single-session (such as VDI) hosts are available for administrators to choose from manually or automatically to ensure a more efficient, streamlined and improved delivery of virtual apps and desktops.

Preconfigured optimizations were designed to be easily updated to support future releases of Microsoft Windows. Moreover, custom scripts may also be used within the tool to make use of already available optimizations to be deployed on Parallels RAS workload machines.

Over 130 image optimizations are available out-of-the-box and divided into the following main categories:

- UWP application packages (removal; available for VDI only)
- Windows Defender ATP (turn ON or OFF, disable real-time scan, exclude files, folder, processes, and extensions)
- Windows components (removal)
- Windows services (disable)
- Windows scheduled tasks (disable)
- Windows advanced options (Cortana, system restore, telemetry, custom layout)
- Network performance (disable task offload, ipv6, etc.)
- Registry (service startup timeout, disk I/O timeout, custom, etc.)
- Visual effects (best appearance, best performance, custom)
- Disk cleanup (delete user profiles, image cleanup, etc.)
- Custom scripts (.ps1, .exe, .cmd, and other extensions/formats)

For the complete list of optimization categories and components, please see <https://kb.parallels.com/125222>.

Optimizations are applicable to RD Session Hosts, VDI desktops, Azure Virtual Desktop, and Remote PC pools (through VDI) based on:

- Windows Server 2012 R2 and later
- Windows 7 SP1
- Windows 10
- Windows 11

Configure optimization

Optimization can be configured for the following:

- RD Session Hosts
- VDI
- Azure Virtual Desktop

Optimization settings are configured for the above on the Site level (Site defaults) and can also be configured for individual components if the RAS administrator decides to use custom settings for a given component.

To configure optimizations on the Site level, navigate to **Farm > Site**, click the **Tasks > Site defaults** menu and choose one of the following:

- **RD Session Host**
- **VDI**
- **AVD multi-session hosts**
- **AVD single-session hosts**

In a Site defaults dialog that opens, select the **Optimization** tab. The user interface for configuring optimization is the same for all of the above.

Note: Before applying optimization, make sure you have a saved state of session hosts as you will not be able to revert changes after they are applied.

To configure optimization:

- 1** If you are in the host **Properties** dialog or in a wizard, clear the **Inherit default settings** options if you want to modify them for this host.
- 2** Select the **Enable optimization** option.

- 3 Choose optimization type from the following:
 - **Automatic:** Predefined and preconfigured optimization will be used automatically.
 - **Manual:** Gives you full control over which optimization options to use and allows you to configure each one. This option also gives you an option to use a custom optimization script that will be executed on the host.
- 4 If you selected **Manual** in the previous step, configure optimization categories and components according to your requirements. See **Configure optimization** below.
- 5 **Force optimization on all enabled categories:** This is a special option that should only be used in situations when some parts of optimization failed to apply to a host for some unforeseen reason (e.g. the host went offline unexpectedly). When you select this option, then click **OK** and then **Apply** in the RAS Console, the entire optimization configuration will be applied to the host. This way you can make sure that changes that you made to optimization components last time, and that were not applied to the host, will be applied again. The state of the **Force optimization on all enabled categories** option (selected or cleared) is not saved because this is a one-time action, so the next time you open the dialog, the option will be cleared again. Note that in a standard scenario, when you make changes and then apply them to a host, you don't need to select this option, because normally you want to apply just the changes that you made, not the entire optimization configuration.
- 6 The **Category** list contains optimization categories that can be configured. To include a category in optimization, select the corresponding checkbox. Some categories contain multiple components, which can be configured individually, some have settings that can be customized. To configure category settings or components, highlight the category and click the gear icon (or click **Tasks > Properties**, or simply double-click a category). Depending on the category selected, you can do the following:
 - Configure category settings (choose from available options, select or clear individual settings, specify values, add or remove entries).
 - Add or remove underlying components to include or exclude them from optimization (use the plus- and minus-sign icons). When adding a component (where available), you can select from a predefined list or you can specify a custom component.
 - In some cases (specifically registry entries) you can double-click an entry and specify multiple values for it.
 - If you remove a predefined component, you can always get it back in the list by clicking **Tasks > Reset to default**. You can also use this menu to reset category settings to default values if they were modified.
 - The last optimization category in the list is **Custom script**. You can use it to execute an optimization script that you may have available. Read the **Using custom script** subsection below for details.
- 7 When done, click **OK** to close the dialog.

Using custom script

The **Custom script** optimization category is used to execute an optimization script on a target host. Before configuring this category, make sure that the script exists on target hosts and that the path and file name are the same on each host.

To configure the Custom script optimization:

- 1 Enable the **Custom script** category in the list (select the checkbox), then highlight it and click **Tasks > Properties**.
- 2 In the dialog that opens, specify the command to execute, arguments (if required), the initial directory, and credentials that will be used to execute the script.
- 3 Click **OK**.

When you apply the optimization to a host, the script will be executed as part of applying other optimization parameters.

Applying optimization

After you enable optimization for a host and then click **Apply** in the RAS Console, the following will happen the next time the host communicates with Parallels RAS:

- 1 The host status changes to **Optimization pending** and the host enters the drain mode. At this stage, you can stop optimization by selecting a host in the list and clicking **Tasks > Stop optimization**.
- 2 Once all users are logged off, the host status changes to **Optimization in progress**.
- 3 After all optimization settings are applied, the host will reboot.
- 4 After the reboot, the host returns to operation and its status changes to **OK**.

Note: By design, the host will be rebooted after optimization completion even if it is failed.

Optimization results are logged on a host at the following location:

%ProgramData%\Parallels\RASLogs\ImageOptimizer.log. Open the file and search for entries similar to the following:

- [I 78/00000009/T10C4/P0FD4] 11-30-20 10:09:19 - Image Optimization completed with 98 successful and 0 unsuccessful optimizations.

Note: By design, Optimization has less priority than Reboot/Disable schedule. For example, it is expected if a host changes the status from "Optimization pending" to Disabled/Reboot when schedule starts.

Upgrade

When Parallels RAS is upgraded from an older version:

- The optimization feature is disabled.
- The inheritance is off.

To use optimization after the upgrade, the administrator needs to enable it manually either in Site defaults or in the group/host pool settings.

Inheritance

	Optimization	Inherits from
RDSH Site defaults	Yes	None
RDSH Group	No	None
RDSH standalone	Yes	RDSH Site defaults
RDSH template	Yes	RDSH Site defaults
RDSH from template	No	None
VDI Site defaults	Yes	None
VDI Desktop standalone	Yes	VDI Site defaults
VDI Desktop template	Yes	VDI Site defaults
VDI Desktop from template	No	None
Azure Virtual Desktop Site defaults	Yes	None
Azure Virtual Desktop host pool - hosts from a template	No	None
Azure Virtual Desktop host pool - standalone hosts	Yes	AVD multi-session hosts Site defaults or AVD single-session hosts Site defaults.
Azure Virtual Desktop template	Yes	AVD multi-session hosts Site defaults or AVD single-session hosts Site defaults.
Azure Virtual Desktop hosts from template	No	None

Additional information

Please note the following:

- Some optimizations may fail and generate warnings if they had been already applied.
- Some optimizations may fail and generate warnings depending on OS specifics. For example, removal of UWP apps may fail because apps are already absent.

Drive redirection cache

This topic explains the **Enable drive redirection cache** option, which is available in a dialog where you configure RAS RD Session Host, VDI, Azure Virtual Desktop, or Remote PC agents. When the option is enabled, browsing folders on redirected drives becomes much faster thanks to the caching mechanism explained below.

Native RDP is not efficient for file and folder enumeration when using drive redirection, which results in slow and sluggish user experience. The **Enable drive redirection cache** option forces the session host to run the kernel-based driver (RasRdpFs). This optimizes how the communication is carried out compared to standard RDP and also adds caching of the folder structure on the session host (RDSH, VDI, or Azure Virtual Desktop). The driver starts as soon as the setting is pushed to the session host via **Apply** in the RAS Console. When this happens, all new sessions will have this functionality enabled. The existing sessions need to be reconnected to use this optimization.

Notes

- A session host must run a 64-bit operating system.
- The cache is per session and is paged into the driver memory.
- On log off or disconnect, the cache is purged.
- If the number of cached folders in the session exceeds the threshold, and the user accesses a new non-cached folder, then the oldest accessed folder is replaced in the cache.
- When the option is switched off, all currently active user sessions will lose the cache (the driver is stopped and the cache is purged). This happens transparently to the user, but file and folder enumeration become slow.
- When the option is switched on, all currently active user sessions will not automatically have the cache enabled. To use this functionality, the existing sessions will need to be reconnected.

Limitations

- The option is applicable only to sessions initiated by the full version of Parallels Client for Windows version 18.
- Similar to native RDP changes made on the client side (in a remote session), requires manual refresh (F5) in a redirected folder on the server side.

Configure logging

An RD Session Host is monitored and logs are created containing relevant information. To configure logging and retrieve or clear existing log files, right-click a server, choose **Troubleshooting > Logging** in the context menu, and then click **Configure**, **Retrieve**, or **Clear** depending on what you want to do. For the information on how to perform these tasks, see the **Logging** (p. 451) section.

Grouping and cloning RD Session Hosts

When you publish resources in Parallels RAS, you need to specify one or more servers that host them. Groups allow you to combine multiple RD Session Hosts and then publish the resources from the group instead of specifying individual servers.

The main benefits of using RD Session Host groups are as follows:

- They simplify the management of published resources and are highly recommended in multi-server environments.
- They allow you to use RD Session Hosts created from a template by utilizing the VDI infrastructure. More on this later in this section.

Note that an RD Session Host can be a member of one group only. You cannot add the same server to multiple groups.

Creating a group

To create an RD Session Host group:

- 1** In the RAS console, navigate to **Farm > <Site> > RD Session Hosts**.
- 2** Click the **Groups** tab.
- 3** Click **Tasks > New Group** (or click the **[+]** icon). To modify an existing group, right-click it and then choose **Properties** in the context menu.
- 4** The **Group Properties** dialog opens where you can specify the group settings as described below.

On the **General** tab, select **Enable Group in site** to enable the group. Type a name and an optional description for the group.

You now need to add one or more RD Session Hosts to the group. You can do this by using the following options (both can be used at the same time):

- Specify a template on which the RD Session Hosts are based. To do so, select the **RD session hosts based on a template** option and then select a template from the drop-down list. Note that you need to create a template of type RD Session Host before you can select it here. Templates of type RD Session Host utilize the VDI functionality available in Parallels RAS. A template is based on a virtual machine running on a hypervisor or a cloud-based VDI provider. When you create a template, you select a preconfigured VM with the operating system and applications already installed. Individual hosts (VMs) are then created as clones of the template. The clones can be created in advance or on as-needed basis (configurable when you create a template). This functionality allows you to create and configure an RD Session Host running in a virtual machine and then create as many copies of it as you require. For the complete information about using VDI in Parallels RAS see the **VDI and Virtual Desktops** chapter (p. 126). Once you are familiar with adding and configuring a VDI provider, read the **Templates** section (p. 147) which explains how to create a template of type RD Session Host.
- Add servers manually one by one by clicking **Tasks > Add** and then selecting a server from the list. You can also add a server later by right-clicking it in the main list and choosing **Assign to group**.

Autoscale

When you select the **RD session hosts based on a template** option (see above), the **Autoscale** tab appears in the **Group Properties** dialog. The settings on this tab determine how RD Session Hosts are created from the specified template. The settings are described below.

Min number of servers to be added to the group from the template: Specifies the minimum number of servers that will be added to the group automatically when the template is assigned to the group. This number of servers will remain in the group irrespective of utilization.

Max number of servers to be added to the group from the template: This option allows you to set a limit on how many servers in total can be added to the group from the template. A template can be shared between groups. By setting a limit for each group, you can ensure that the combined number of servers in each group will not exceed the template limit. Consider the following examples:

- If the template is used by a single group, then this number can be up to the "Maximum guest VMs" setting of the template.
- If two or more groups share the same template, then the combined number from all groups must be less or equal to the "Maximum guest VMs" settings of the template.

When you save a group, a validation will be performed against other groups (if any) and you will see an error message if the numbers don't match. Note that when a server cannot be created on request due to an error, a "Template error" event is triggered and the administrator will receive an alert message.

Add servers from template when workload is above (%): Specifies a workload threshold in percent. When the actual workload is above this value, a new server (or servers) will be created and added to the group (if not already available). The group workload percentage is calculated using the following formula:

$$\text{Group Workload} = (\text{Current Sessions} / \text{Max Sessions}) * 100$$

In the formula above:

- Current Sessions is the total number of all sessions on all servers in the group. This includes static (standalone) servers and servers created from the template (guest VMs). Note that servers that are disabled, being drained, or have the agent status of 'Not Verified' are not included in the calculation.
- Max Sessions is a setting that you specify on the **Agent Settings** tab (either inherited from Site defaults or overridden for this group) and the maximum number of sessions allowed for the group.

Consider the following examples:

RAS Group 1 — mixed server types (static and guest VMs), different agent status:

- RDSH-1, Status: OK, Max Sessions 10, Current Sessions: 2, Type: Static
- RDSH-2, Status: Disabled, Max Sessions 20, Current Sessions: 0, Type: Static
- RDSH-3, Status: OK, Max sessions 10, Current Sessions: 4, Type: Guest VM
- RDSH-4, Status: Drain Mode, Max sessions 10, Current Sessions: 3, Type: Guest VM

For the group above, the workload is calculated as $(\text{Current Sessions} / \text{Max Sessions}) * 100$ or $((2 + 4) / 20) * 100 = 30\%$

Note that servers RDSH-2 and RDSH-4 are not included in the workload because the former has the agent disabled and the latter is in drain mode.

RAS Group 2 — mixed server types (static and guest VMs), different agent status:

- RDSH-1, Status: OK, Max Sessions 10, Current Sessions: 0, Type: Guest VM
- RDSH-2, Status: OK, Max Sessions 10, Current Sessions: 2, Type: Guest VM
- RDSH-3, Status: Not Verified, Max sessions 10, Current Sessions: 0, Type: Guest VM

$$\text{Group Workload} = (\text{Current Sessions} / \text{Max Sessions}) * 100 \text{ or } ((0 + 2) / 20) * 100 = 10\%$$

Please note that a group will always make sure that it has at least one server available, even if the workload is zero percent.

Number of servers to be added to the group per request: Specifies how many servers should be created when the workload goes above the threshold value. This setting works together with the **Add servers from template when workload is above (%)** setting described above. When a group sends a request to the template to create additional servers, the value specified here will determine the number of servers that will be created.

Drain and unassign servers from group when workload is below (%): Specifies a workload threshold in percent. When the actual workload is below this value and remains there for a period specified in the "remains below this level" field, excessive hosts will be switched to drain mode or unassigned from the group. The period of time can be selected from the drop-down list or you can type your own integer value using "weeks", "days", "hours", "minutes", or "seconds" as a unit measure. The server(s) with the least number of sessions will be switched to drain mode. As soon as all users are logged off from a server, it is unassigned from the group. At that point, the server becomes available to other groups on demand.

Tip: Servers are unassigned from the group only when all user sessions on that particular server are logged off. In case user sessions are still present, such as user sessions in idle, active or disconnected state, autoscaling does not log off user sessions and does not unassign the server from a group.

Note: Parallels recommends setting viable timeouts for idle time and disconnected sessions either in Windows Group Policies or in the **Site Default Properties** dialog to make the drain mode effective. GPOs can be used to forcibly log off a user session, however this should be used carefully as this may result in data loss.

Using group defaults

RD Sessions Hosts assigned to a group have various settings that they can inherit from the group defaults. This makes it simpler to configure a single set of settings for all servers instead of configuring each server individually. A Site also has its own default settings (Site defaults). Moreover, an RD Session Host group can inherit these Site defaults. This gives you the following choices when inheriting default settings by an RD Session Host:

- Configure Site defaults and make the group inherit these settings. The RD Session Hosts assigned to the group will therefore also inherit Site defaults. This is the default scenario for a new group. Site defaults can be configured by navigating to **Farm > <Site> > RD Session hosts** and clicking **Tasks > Site defaults**.
- Configure default settings for a given group. This way you can have multiple groups, each having its own group defaults (different from Site defaults). Therefore, the servers assigned to a group will inherit the group's defaults.

To configure default settings for a group, open the **Group Properties** dialog (**Tasks > Properties**), select a desired tab (except the **General** tab, which doesn't have any defaults) and select or clear the **Inherit default settings** option. If you clear the option, you can specify your own defaults. All servers that are (or will be) assigned to this group will inherit these settings. Note that inheritance works independently for each individual tab on the group properties dialog.

For the information on how default settings are configured for an RD Session Host, see **View and modify RD Session Host properties** (p. 90).

Removing a server from a group

To remove a regular RD Session Host from a group, do one of the following:

- On the **RD Session Hosts** tab, right-click a server and choose **Remove from group**.
- On the **Groups** tab, right-click a group and choose **Properties**. In the **Group Properties** dialog, select a server and click **Tasks > Delete**.

To remove an RD Session Host that was added to a group from a template:

- 1 Go to the **Groups** tab.
- 2 Select a group and click **Tasks > Properties**.
- 3 In the **Group Properties** dialog, select a server and click **Tasks > Delete**.

Note that this is the only place in the RAS Console where you can remove an RD Session Host of this type from a group. Please also note that when you delete such a host, it is drained first and only then unassigned from the group, which may take a considerable amount of time.

After you create a group and later publish resources from it, you can view the list of resources by right-clicking a group and choosing **Show published resources** (or click **Tasks > Show Published Resources**). For more information, see **Viewing published resources hosted by RD Session Hosts** (p. 125).

Using scheduler

The **Scheduler** tab in the **RD Session Hosts** view allows you to reboot or temporarily disable servers according to a schedule.

To create a new scheduler task or modify an existing one:

- 1 In the RAS Console, navigate to **Farm > <Site> > RD Session Hosts**.
- 2 In the right pane, select the **Scheduler** tab.
- 3 To create a new task, click **Tasks > Add** and select one of the following options:
 - **Disable Server**
 - **Disable Server Group**
 - **Reboot Server**
 - **Reboot Server Group**

The **RDSH Schedule Properties** dialog opens. The dialog consists of three tabs, which are described below.

Properties

On the **Properties** tab, specify the following:

- Select **Enable Schedule** to enable the scheduled task.
- Specify the task name and an optional description.
- In the **Available** list, select target servers or groups and click **Add** (repeat to add more servers or groups). To add all servers, click **Add all**. To remove a server or servers from the **Target** list, click **Remove** or **Remove all**.

Trigger

On the **Trigger** tab, specify when the scheduled task should trigger:

- In the **Date**, **Start**, and **Duration** fields, specify the start date, time, and duration.
- In the **Recur** field, specify the task recurrence. If you select **Never**, the task will still run as scheduled but only once. If you select **On specific day(s) of the week**, you need to select one or more days of the week.

Options

On the **Options** tab, you can do the following:

- Compose a message that will be sent to users before or after (in certain scenarios) the scheduled task is triggered. Composing a message is described later in this subsection.
- Specify additional options. Please note that the options are different depending on the task type, as described below.

If the task is **Disable Server** or **Disable Server Group**, the available options are:

- **On Disable:** Use this option to specify how active sessions should be handled when the task is triggered. Please note that disabling a server group with an assigned template will drain and remove RD Session Hosts from the group. See **Maintaining RD Session Hosts based on a Template** (p. 115).
- **Enforce schedule for currently inactive RD Session Hosts:** This option is only enabled when you have an active message in the list, which is displayed above these options. If the option is enabled, RD Session Hosts that are currently offline are also monitored, and if such a server comes back online during the scheduled task execution, the task is applied to it too.
- If you enable this option, the schedule will be applied to a currently inactive RD Session Host when it comes back online. If the option is disabled (default), the schedule will have no effect on such servers. Note that it is assumed that a server is inactive (offline) if it is disabled or cannot be reached over the network (registered on RAS Publishing Agent).

If a task is **Reboot Server** or **Reboot Server Group**, the available options are:

- **Enable Drain Mode** and **Force server reboot after**: The two options work together. If you enable the drain mode, the following will happen. When the task triggers, new connections to a server are refused but active connections will continue to run. The server will be rebooted when all active users close their sessions or when **Force server reboot after** time is reached, whichever comes first. For active users not to lose their work, create a message that will advise them to save their work and log off (see below for details). Please also see **RD Session Host drain mode examples** (p. 115).
- **Enforce schedule for currently inactive RD Session Hosts**: This option is enabled when the **Enable Drain Mode** option is selected. If the option is enabled, RD Session Hosts that are currently offline are also monitored and if such a server comes back online during the scheduled task execution, the task is applied to it too.

To create a text message to be sent to users, click the **Tasks > Add** and specify the following:

- Select the **Enable Message** option to enable the message. If the option is cleared, the message will still exist, but will not be sent to users. You can also enable or disable an existing message by selecting or clearing a checkbox in the list on the **Options** tab.
- Specify the message title and body. This is what users will see when the message is displayed on their screens.
- In the **Send message** drop-down list, select the time interval specifying when the message should be sent. By default, this is the time "before" the task is triggered. However, for **Reboot Server** and **Reboot Server Group** tasks, it can also be the time "after" the task is triggered, i.e. the server is put to drain mode. This may be specifically useful when you want to send multiple messages to users at different time intervals while the scheduled task is already in progress. See the explanation below.

Sending multiple messages to users

For **Disable Server** and **Disable Server Group** tasks, you can only send a message before the scheduled task is triggered. Hence, when creating a message, you can only select the "before" option when specifying when the message should be sent. You can create more than one message if needed and send them at different time intervals, so the users are notified more than once before the task executes.

For **Reboot server** and **Reboot server group** tasks, you can send a message before or after the scheduled task is triggered. The "after" option is available for these tasks because you have the ability to enable the drain mode, which will keep the active sessions running for some time. During this time, you can send multiple messages to active users reminding them that they should finish their work and close their sessions. To use the "after" option, the **Enable Drain Mode** option must be selected. Please also note that the "after" time interval and the **Force server reboot after** setting should be coordinated. For example, if the force reboot occurs before the "after" time elapses, active users will not have a chance to see the message.

RD Session Host drain mode examples

Example 1: Scheduling a server group for reboot without the drain mode

A server group contains 3 servers: A, B, C

- Date: 1/24/2020
- Start time: 10:45am
- Send message: 2 minutes before

Users with active sessions are notified 2 minutes before the server reboot task is triggered.

Example 2: Scheduling a server group for reboot with the drain mode enabled

A server group containing 3 servers: A, B, C

- Date: 1/24/2020
- Start Time: 10:45am
- Drain mode: enabled
- Force reboot after: 1 hour
- Send messages: 2 minutes before, 15 minutes after, 30 minutes after.

The session users are notified 2 minutes before the server reboot task is triggered and then twice more, 15 and 30 minutes after the task is triggered. Because the drain mode is enabled, the user sessions will continue to run, so they will see the messages and will be able to close their sessions before the server reboots. Note that since the force reboot time is set at 1 hour, the users will see the last message, which will be sent 30 minutes after the task is triggered.

When the task is triggered:

- 1 The drain mode is enabled on the servers.
- 2 Servers A and B have no active or disconnected sessions, so they are restarted immediately.
- 3 Server C still has open/disconnected sessions, so it continues to run until all users end their sessions. If in 1 hour the server still has active sessions, they are terminated and the server is restarted.

Maintaining RD Session Hosts based on a template

If you need to perform a scheduled maintenance of RD Session Hosts that were created from a template, please follow these steps:

- 1 Create a "Disable server group(s)" schedule that fits your maintenance window and apply to RD Session Host group(s) with the assigned template.

When the scheduler disables the group:

- All servers in a group have Agent Status as "Disabled (scheduler)", the Logon Status stays 'Enabled'.
 - New sessions are prohibited.
 - If the administrator specifies the option on disable as "Reset all sessions", sessions are logged off, but templated RD Session Hosts are not removed from group.
- 2** During maintenance window (or right before it) switch the template to the maintenance mode. Then apply the necessary changes.
 - 3** If you want to recreate all guest VMs on exiting the maintenance mode, you need to disable RD Session Host group(s) on the **Groups** tab. To do so, clear the checkbox in front of a group name (on the left side) and click **Apply**.

When the group(a) are disabled:

- Templated RD Session Hosts with 0 user sessions are removed from the group(s) (unassigned).
 - Since the templated RD Session Hosts do not belong to any group anymore, they can be recreated on exiting the template maintenance mode.
- 4** Release the template from maintenance and click **Yes** when asked whether to recreate all clones.
 - 5** Enable groups which were disabled earlier. At this point, the groups will begin receiving recreated guest VMs to comply with the "Keep Available Buffer" setting.
 - 6** From this point forward, groups are provisioned with updated templated RD Session Hosts on demand.

Managing RD Session Host sessions

Please see **Session Management (p. 227)**.

Managing logons

The logon management feature allows you to enable or disable logons from RD Session Hosts. The feature performs the same tasks as the `change logon` command-line utility.

Note: For RD Session Hosts based on a template, the drain mode (which disables logons) is handled automatically by the group to which a host belongs. For more information see **Using Scheduler (p. 112)**.

To manage logons:

- 1** In the Parallels RAS Console, navigate to **Farm** > <Site> > **RD Session Hosts**.

2 Select an RD Session Host, click **Tasks > Control** and choose one of the following:

- **Enable logons:** Enables logons from client sessions, but not from the console. This option performs the same action as the `change logon /enable` command.
- **Disable logons:** Disables subsequent logons from client sessions, but not from the console. Does not affect currently logged on users. This option performs the same action as `change logon /disable` command.
- **Drain:** Disables logons from new client sessions, but allows reconnections to existing sessions. Drain is kept even after reboot until the admin enables logons.

Note that while a server is in drain mode, administrators may still log on to the physical console or remotely log on using the `/admin` or `/console` command-line option for MSTSC. This allows administrators to remotely maintain the RDS server via **Tools > Remote Desktop**.

- **Drain until reboot:** Disables logons from new client sessions until the computer is restarted, but allows reconnections to existing sessions. Drain is kept until the server is restarted. Same action as the `change logon /drainuntilrestart` command.

To see the current logon control mode for an RD Session Host, click **Tasks > Control**. The checked-out option indicates the current logon control mode of the selected RD Session Host. To do this check from the command line, execute the `change logon /QUERY` command on the server.

Please also note the following:

- When applying a logon control mode on a server, ensure that the agent status is updated accordingly.
- You must set the logon control options for the servers one-by-one. If you need to do it for a group of servers, you can use the scheduler (see **Using scheduler** (p. 112)).
- `/Drain` disables logons from new client sessions, but allows reconnections to existing sessions. Drain differs from `Drainuntilrestart` as far as it requires administrator intervention whereas the latter enables logons automatically after restart.
- **Computer Configuration / Administrative Templates / Windows Components / Remote Desktop Services / Remote Desktop Session Host / Connection / Allow users to connect remotely using remote desktop services** must be set to **Not configured**, otherwise it takes precedence.

Using computer management tools

You can perform standard computer management tasks on an RD Session Host right from the RAS Console. These include Remote Desktop Connection, PowerShell, Computer Management, Service Management, Event Viewer, IPconfig, Reboot, and others. To access the **Tools** menu, select a server, click **Tasks** (or right-click) > **Tools** and choose a desired tool. For requirements and usage information, see **Computer management tools** (p. 431).

Publishing from an RD Session Host

This section describes how to publish resources hosted by an RD Session Host. The publishing functionality described here is accessed from the **Publishing** category in the RAS Console.

You can also publish resources using a publishing wizard in the **Start** category, as described in the **Set up a basic Parallels RAS Farm** section (p. 29). The **Start** category publishing wizard is a simplified version that gives you convenient options of selecting the resources that you want to publish. You may try both approaches and choose the one that better suits your needs.

Read on to learn how to publish resources from an RD Session Host.

Publishing a desktop

To publish a remote desktop from an RD Session Host:

- 1 In the RAS Console, select the **Publishing** category and click the **Add** icon below the **Published Resources** tree. This will launch the publishing wizard.

Note: If the wizard has all options disabled, it means that there are no resources (servers) in the Farm from which publishing can be configured.

- 2 In the first step of the wizard, select **Desktop** and click **Next**.
- 3 In the **Select Desktop Type** step, select **RD Session Host Desktop** and click **Next**.
- 4 Select one or more RD Session Hosts which desktops you want to publish. You can select all available servers, server group(s), or individual servers. Please note that if you have just one RD Session Host, this page will not be displayed.
- 5 Click **Next**.
- 6 In the next step:
 - Specify a name and description for the desktop, and optionally change the icon.
 - Select the **Connect to administrative session** option if you want users to connect to the administrative session.
 - Select **Exclude from session prelaunch** if needed. For details, see **Understanding session prelaunch** (p. 252).
 - Select the **Start automatically when user logs on** option if you want to open a desktop as soon as a user logs on.
 - Specify the desired screen resolution using the **Desktop Size** drop-down list. To set a custom width and height of the screen, select **Custom** in the **Size** drop-down list and specify the desired values in the fields provided.
 - In the **Multi-Monitor** drop-down list, select whether the multi-monitor support should be enabled, disabled, or whether the client settings should be used.

- 7 On the next page, specify the initial status of the resource. Choose from **Enabled** (end users can launch the resource), **Disabled** (the resource will not appear in Parallels Client), **In maintenance** (the resource will appear in Parallels Client but users will not be able to launch it). When a resource is in maintenance and a user tries to launch it, they will see a message. To customize the message, click the **Configure** button. For more info, see **Site defaults (Publishing)** (p. 245).
- 8 When done, click **Finish** to publish the desktop.

Publishing an application

To publish an application from an RD Session Host follow the below procedure:

- 1 In the RAS Console, select the **Publishing** category and then click the **Add** icon below the **Published Resources** tree (or right-click inside the **Published Resources** box and click **Add** in the context menu). This will launch the publishing wizard.

Note: If the wizard has all options disabled, it means that there are no resources (servers) in the Farm from which publishing can be configured.

- 2 On the **Select Item Type** wizard page, select **Application** and click **Next**.
- 3 On the **Select Server Type** page, select **RD Session Host** and click **Next**.
- 4 On the **Select Application Type** page, select one of the following available options:
 - **Single Application.** Choose this option to fully configure the application settings yourself such as the executable path etc.
 - **Installed Application.** Choose this option to publish an application that is already installed on the server, therefore all of the application settings are automatically configured.
 - **Predefined Application.** Choose this option to publish a commonly used Windows application such as Windows Explorer.
- 5 Click **Next**.
- 6 On the **Publish From** page, specify from which RD Session Hosts the application should be published. You have the following options:
 - **All Servers in Site.** If selected, the application will be published from all servers that are available on the Site.
 - **Server Groups.** Select this option and then select individual server groups to publish the application from.
 - **Individual Servers.** Select this option and select individual servers to publish the application from.

Please note that the **Publish From** wizard page will appear only if you have multiple RD Session Hosts. If you have just one server, this page will be skipped by the wizard. The page will also be skipped if the application type that you are installing is **Predefined Application**.

- 7 Click **Next**.

8 Depending on the application type that you selected on the **Select Application Type** page, the next wizard page will be one of the following:

- If you selected **Single Application**, the **Application** page will open where you have to specify the application settings manually (more about this option later in this section).
- If you selected **Installed Applications**, the **Installed Applications** page will open listing available applications (the applications are grouped by functionality). Select an application you wish to install and click **Next**. Follows the instructions to complete the wizard.
- If you selected **Predefined Application**, the **Select Predefined Applications** page will open listing available applications. Select an application you wish to publish and click **Finish**.

9 If you selected **Single Application** on the **Select Application Type** wizard page, the **Application** page will open. Specify the application settings as follows (see the screenshot below):

Note that if you populate the **Target** field first using the "browse" button ([...]), the application **Name**, **Description**, and icon will be chosen automatically. You can override this selection if you wish.

- **Name**. Choose and type a name for the application.
- **Description**. Type an optional description.
- **Run**. Select the application window state (normal window, minimized, maximized).
- **Exclude from session prelaunch**. For details, see **Understanding session prelaunch** (p. 252).
- **Start automatically when user logs on**. Select this option if you want to start an application as soon as a user logs on. This option works on desktop versions of Parallels Client only.
- **Change Icon**. Change the application icon (optional).
- **Server(s)**. Allows you to specify the rest of the server parameters individually for each server the application was published from. Select a server from the drop-down list box and specify the parameters. Repeat for other servers in the list.
- **Target**. Specify the application executable path and file name.
- **Start in**. If the **Target** field is valid, this field will be populated automatically. You can specify your own path if needed.
- **Parameters**. If the application accepts startup parameters, you can specify them in this field.

10 On the next page, specify the initial status of the resource. Choose from **Enabled** (end users can launch the resource), **Disabled** (the resource will not appear in Parallels Client), **In maintenance** (the resource will appear in Parallels Client but users will not be able to launch it). When a resource is in maintenance and a user tries to launch it, they will see a message. To customize the message, click the **Configure** button. For more info, see **Site defaults (Publishing)** (p. 245).

11 When done, click **Finish** to publish the application.

Publishing a web application

A web application is like any other application that you can publish using the standard application publishing functionality. However, to simplify publishing of straight URL links to web applications, a separate publishing item type is available that allows you to accomplish this task with minimal number of steps.

To publish a web application:

- 1 In the RAS Console, select the **Publishing** category and then click the **Add** icon below the **Published Resources** tree (or right-click inside the **Published Resources** box and click **Add** in the context menu). This will launch the publishing wizard.

Note: If the wizard has all options disabled, it means that there are no resources (servers) in the Farm from which publishing can be configured.

- 2 On the **Select Item Type** wizard page, select **Web Application** and click **Next**.
- 3 On the **Select Server Type** page, select **RD Session Host** and click **Next**.
- 4 On the **Publish From** page, select the server(s) to publish from. Note that if you have just one RD Session Host, the **Publish From** page will not appear.
- 5 On the **Web Application** wizard page that opens, specify the web application name, description, window state, and the URL. Select the **Force to use Internet Explorer** option if needed. To browse for a specific application icon, click **Change Icon**.
- 6 On the next page, specify the initial status of the resource. Choose from **Enabled** (end users can launch the resource), **Disabled** (the resource will not appear in Parallels Client), **In maintenance** (the resource will appear in Parallels Client but users will not be able to launch it). When a resource is in maintenance and a user tries to launch it, they will see a message. To customize the message, click the **Configure** button. For more info, see **Site defaults (Publishing)** (p. 245).
- 7 When done, click **Finish** to publish the application.

When published, the web application will appear in the **Publishing > Published Resources list**, just like any other application.

Publishing a network folder

You can publish a filesystem folder via UNC path to open in Windows explorer. To minimize the number of configuration steps, a special publishing item is available that allows you to publish a network folder from an RD Session Host.

To publish a network folder:

- 1 In the RAS Console, select the **Publishing** category and then click the **Add** icon below the **Published Resources** tree (or right-click inside the **Published Resources** box and click **Add** in the context menu). This will launch the publishing wizard.

Note: If the wizard has all options disabled, it means that there are no resources (servers) in the Farm from which publishing can be configured.

- 2 On the **Select Item Type** wizard page, select **Folder on the file system** and click **Next**.
- 3 On the **Select Server Type** page, select **RD Session Host** and click **Next**.
- 4 On the **Publish From** page, select the server(s) to publish from. Note that if you have just one RD Session Host, the **Publish From** page will not appear.
- 5 On the **UNC Folder** wizard page, specify the usual application properties.
- 6 In the **UNC path** field, enter the UNC path of the folder you wish to publish. Click the **[...]** button to browse for a folder (it may take some time for the **Browse for Folder** dialog to open).
- 7 On the next page, specify the initial status of the resource. Choose from **Enabled** (end users can launch the resource), **Disabled** (the resource will not appear in Parallels Client), **In maintenance** (the resource will appear in Parallels Client but users will not be able to launch it). When a resource is in maintenance and a user tries to launch it, they will see a message. To customize the message, click the **Configure** button. For more info, see **Site defaults (Publishing)** (p. 245).
- 8 Click **Finish** to publish the folder and close the wizard.

When published, the network folder will appear in the **Publishing > Published Resources list**, just like any other application. If you select it and then click the **Application** tab, the application settings will be as follows:

- The **Target** property will always be set to `PublishedExplorer.exe`. This binary is created automatically (via agents pushing) and is simply a copy of the standard `explorer.exe` executable.
- The **Parameters** property specifies the network folder that we want to publish. The folder path can be in any format that the `explorer.exe` can handle.

Please note that although you have all standard application property tabs enabled for this publishing item, at least the following items should be ignored, as they are completely irrelevant:

- **Publish From**
- **File Extensions**

Publishing a document

To publish a document from an RD Session Host, follow the below procedure:

- 1 In the RAS Console, select the **Publishing** category and then click the **Add** icon below the **Published Resources** tree (or right-click inside the **Published Resources** box and click **Add** in the context menu). This will launch the publishing wizard.

Note: If the wizard has all options disabled, it means that there are no resources (servers) in the Farm from which publishing can be configured.

- 2 On the **Select Item Type** wizard page, select **Document** and click **Next**.

- 3 Select **RD Session Host** and click **Next**.
- 4 Specify the content type of the document you want to publish. You can select the content type from the predefined list or specify a custom content type in the **Custom content types** input field.
- 5 Click **Next** when ready.
- 6 On the **Publish From** page, specify from which RD Session Hosts the application should be published. You have the following options:
 - **All Servers in Site**. If selected, the application will be published from all servers that are available on the Site.
 - **Server Groups**. Select this option and then select individual server groups to publish the application from.
 - **Individual Servers**. Select this option and select individual servers to publish the application from.

Please note that the **Publish From** wizard page will appear only if you have multiple RD Session Hosts. If you have just one server, this page will be skipped by the wizard.

- 7 On the **Application** page, enter a name, an optional description, a Window state, and an icon if needed.
- 8 Use the [...] button next to the **Target** input field to browse for the document. All other fields will be automatically populated. To edit any of the auto populated fields, highlight them and enter the required details.
- 9 (Optional) In the **Parameters** input field, specify the parameters to pass to the application when it starts.

Note: Use the **Server(s)** drop down list to specify different document settings for a specific server in case the document is configured differently on that particular server. The settings will be saved for each server you select individually.

- 10 On the next page, specify the initial status of the resource. Choose from **Enabled** (end users can launch the resource), **Disabled** (the resource will not appear in Parallels Client), **In maintenance** (the resource will appear in Parallels Client but users will not be able to launch it). When a resource is in maintenance and a user tries to launch it, they will see a message. To customize the message, click the **Configure** button. For more info, see **Site defaults (Publishing)** (p. 245).
- 11 Click **Finish** to publish the document.

Publishing containerized applications

Parallels RAS supports publishing of the following containerized applications:

- App-V applications (p. 124)

Publishing App-V applications

Microsoft Application Virtualization (or App-V) is an application streaming solution from Microsoft. Beginning with Parallels RAS v16.5, a support for App-V application publishing is available in the Parallels RAS console.

At the time of this writing, the App-V support implements scenarios where application provisioning is performed by means of App-V components:

- Applications are sequenced by the administrator according to Microsoft guidelines.
- Applications are stored on a network share created by the administrator (SMB, HTTPs).
- App-V Management and Publishing servers are used to publish applications for a specific AD groups that must be synced manually by the administrator with RAS publishing groups used for App-V application publishing.
- App-V client is installed and configured manually by the administrator.

The process of deploying and publishing an App-V application is as follows:

- 1 Package an applications using the App-V Sequencer.
- 2 Deploy the application to an RD Session Host using the App-V Management Console, Microsoft SCCM, etc.
- 3 Provision the application.
- 4 Verify that users can launch the application from the RD Session Host.
- 5 Publish the application from RAS Console (see below for instructions).
- 6 Launch the application from a Parallels Client.

Publish an App-V application from the Parallels RAS console

To publish an App-V application:

- 1 In the Parallels RAS Console, select the **Publishing** category.
- 2 Click the **[+] Add** icon at the bottom of the right pane. The publishing wizard opens.
- 3 On the **Select Item Type** page, select the **App-V application** option.
- 4 Click **Next**.
- 5 Select the server type from which to publish an application and click **Next**.

- 6 Select a server or a group to publish from and click **Next**.
- 7 On the **Installed Applications** page, select one or more App-V applications and click **Next**.
- 8 Review the summary and complete the wizard.

Once an App-V application is published, it can be launched from a Parallels Client.

Note: To avoid launch issues, use AutoLoad=2. More details in https://blogs.technet.microsoft.com/technetsto_blog/2013/11/12/autoload-setting-in-app-v-5-0/.

Viewing published resources

When you want to remove an RD Session Host or an RD Session Host group from a Site, you might want to see the list of published resources hosted by the server or servers in a group. This way you can see which resources will be affected. You can do so as follows:

- 1 In the Parallels RAS Console, select **Farm \ RD Session hosts**.
- 2 To see published resources for a specific RD Session Host, select the **RD Session hosts** tab. To see published resources for a group, select the **Groups** tab.
- 3 Right-click a server or a group and choose **Show published resources** (or click **Tasks > Show published resources**).
- 4 The **Published Resources** window opens displaying the list of published resources for the selected server or group. Resource information includes:
 - **Name.** Resource name.
 - **Status.** Enabled or disabled.
 - **Type.** "Application" is used for published applications, URLs, network folders, etc. "Desktop" is used for published desktops.
 - **Path.** For published applications, specifies a path to the execute file, URL, or UNC path.
 - **Parameters.** Published application parameters (if any).
 - **Published from.** Site, group(s), or individual server(s).
- 5 To refresh the list, press F5 or click the "recycle" icon (top-right).
- 6 To filter the list, press Ctrl-F or click the magnifying glass icon and then specify the filter criteria for desired column(s).

Virtual Desktop Infrastructure (VDI)

Parallels RAS VDI (Virtual Desktop Infrastructure) enables you to use server virtualization to reduce the number of physical servers required to host published resources. Parallels RAS VDI supports numerous virtualization technologies, including hypervisor and cloud-based platforms.

Parallels RAS VDI also includes the Template functionality, which gives you the ability to create a template from a preconfigured guest VM (virtual machine) and then automatically clone guest VMs and RD Session Host VMs from it.

In This Chapter

Supported VDI providers	126
RAS VDI Agent information.....	127
Add a VDI provider	129
Modifying VDI Provider configuration	138
Configure logging	142
Enabling high availability for VDI.....	142
Change VDI provider Site assignment.....	144
Site defaults (VDI)	144
Viewing guest VMs on a VDI provider	146
Templates	147
VDI pool management.....	170
Managing guest VMs.....	172
Persistent guest VMs.....	174
Using a VDI provider in multiple farms.....	175
Using computer management tools.....	176
Publishing from a guest VM.....	176
Viewing VDI provider summary	181
Managing VDI sessions	181
Remote PC pools.....	181

Supported VDI providers

Parallels RAS supports hypervisor-based VDI providers and cloud-based VDI providers.

Hypervisors

The following hypervisors are supported:

- Microsoft Hyper-V (Windows Server 2012 R2 up to Windows Server 2022)
- Microsoft Hyper-V Failover Cluster (Windows Server 2012 R2 up to Windows Server 2022)
- VMware vCenter 7.x, 6.7.0, 6.5.0, 6.0.0, 5.5.0, 5.1.0, 5.0.0, 4.1.0, 4.0.0
- VMware ESXi 7.x, 6.7.0, 6.5.0, 6.0.0, 5.5.0, 5.1.0, 5.0.0, 4.1.0, 4.0.0
- Scale Computing HC3 8.6.5, 8.8, 8.9*
- Nutanix AHV AOS 5.5, 5.10, 5.15, 5.20
- Remote PC — This is a special type that allows you to create pools of Remote PCs. See **Remote PC pools** (p. 181)

* Experimental support

Cloud VDI providers

The only cloud VDI provider supported at this time is Microsoft Azure.

RAS VDI Agent information

In order to function in a RAS Farm, a VDI provider (hypervisor or cloud-based) needs RAS VDI Agent to be installed in the Farm. RAS VDI Agent acts as an interface between other RAS components and a VDI provider. RAS VDI Agent conducts all communications with a VDI provider through the provider's native API.

Parallels RAS has two types of RAS VDI Agents that can be installed in a Farm:

- **Built-in:** This RAS VDI Agent is built into the RAS Publishing Agent and is installed automatically when you install Parallels RAS. The agent can handle multiple VDI providers and can also be configured for high availability.
- **Dedicated:** This RAS VDI Agent is installed manually. It can handle only a single VDI provider. If you want to use this agent type with more than one provider, you need to install a separate instance for each provider.

Both built-in and dedicated RAS VDI Agents are compatible with all types of VDI providers supported by Parallels RAS. Which agent you choose to install depends only on your requirements. When possible, it is always recommended to use the built-in VDI Agent for high availability and business continuity.

What to read next:

- If you are adding a VDI provider that will use the built-in RAS VDI Agent, you may skip to **Add a VDI provider** (p. 129).
- If you want to install a dedicated RAS VDI Agent on a server of your choice, read the **RAS VDI Agent installation options** section (p. 128), which follows this one.

RAS VDI Agent installation options

If you are installing a dedicated RAS VDI Agent, you first need to determine where it will be installed. Depending on the VDI provider type, the following options are available:

- The host on which the hypervisor is running. This option is available for Microsoft Hyper-V only.
- A supported version of Windows Server running on a physical box or in a virtual machine. For supported Windows Server versions, see **Software requirements > RAS VDI Agent**.

The following table lists RAS VDI Agent installation options for each supported VDI provider:

VDI provider	Built-in Agent (part of PA)	Agent on a VDI provider	Agent on a Windows Server (VM or HW)
Microsoft Hyper-V	Yes	Yes	Yes*
Microsoft Hyper-V Failover Cluster	Yes	No	Yes*
VMware VCenter	Yes	No	Yes*
VMware ESXi	Yes	No	Yes*
Scale Computing HC3	Yes	No	Yes*
Nutanix Acropolis	Yes	No	Yes*
Remote PC (see the Note below)	Yes	No	Yes*
Microsoft Azure	Yes	No	Yes*

* High Availability is not available with these VDI Agent installation options. For details, see **Enabling high availability for VDI** (p. 142).

Note: The **Remote PC** is a special type that can be used to create and manage pools of Remote PCs as part of hosted desktop infrastructure (HDI). When you add a VDI provider of this type, you can manage it like one of the real VDI providers with some limitations, such as you cannot create templates and use some other strictly VDI-specific functions. The main feature when using this type is the ability to create pools of HDI-based Remote PCs (e.g. HPE Moonshot System, Atrust Remote PC Array) and making PCs persistent by assigning an individual PC to a specific user. For more info, see **Remote PC pools** (p. 181).

In the table above, find the VDI provider type that you are using and see where the RAS VDI Agent can be installed.

Depending on the available choices, do one of the following:

- **Built-in Agent:** The agent is a part of RAS Publishing Agent, so it is already installed. When possible, it is always recommended to use the built-in VDI Agent for high availability and business continuity.
- **Agent on a VDI provider:** This option is only available if you are using Microsoft Hyper-V. You can simply install the agent on the host, as described in **Add a VDI provider** (p. 129).
- **Agent on a Windows Server (VM or HW):** To use this option, make sure you have a physical box or a virtual machine running a supported version of Windows Server. You will need to specify its FQDN or IP address when adding a VDI provider to the Farm.

Add a VDI provider

In this section:

- **Add a hypervisor VDI provider** (p. 129)
- **Add a cloud VDI provider** (p. 131)

Add a hypervisor VDI provider

This section describes how to add a hypervisor-based VDI provider (p. 126). For the information on how to add a cloud-based VDI provider, see **Add a cloud VDI provider** (p. 131).

To add a VDI provider:

- 1 In the RAS Console, navigate to **Farm > Site > VDI**.
- 2 On the **Providers** tab, click **Tasks > Add**.
- 3 The **Add VDI Provider** wizard opens.
- 4 On the **Select VDI provider type** page, select **Virtualization** to use a hypervisor deployed on your network. The **Cloud computing** option is described in the **Add a cloud VDI provider** section (p. 131).
- 5 Click **Next**.
- 6 In the **Type** field, select the hypervisor type. For the information on what the "Remote PC" type is, see **Remote PC pools** (p. 181).
- 7 In the **Address** field, specify the host's FQDN or IP address.
- 8 Specify a user name and password to log in to the server.
- 9 Type an optional description.

10 Click the **Advanced Settings** link to open the **Advanced VDI Provider Settings** dialog. The dialog allows you choose the following options:

- **Use dedicated VDI Agent:** Select this option if you will install (or have installed) the RAS VDI Agent yourself. Clear the option if you will use the built-in RAS VDI Agent (p. 127).
- **Agent address:** This option becomes enabled if you select the option above it. Specify the FQDN or IP address of the server where the RAS VDI Agent is (or will be) installed. This can be either a physical box or virtual machine.
- **Preferred Publishing Agent:** Select a RAS Publishing Agent to be the preferred agent for this VDI provider. Select **Automatic** to let the system select an agent (this option is enabled and selected by default if you have at least three Publishing Agents installed). The automatic selection is especially important to ensure that the VDI provider and guest VMs always have a VDI Agent they communicate with in case of other Publishing Agent(s) failure. For more info, see **Enabling high availability for VDI** (p. 142).

11 Click **Next**.

12 The wizard will now try to connect to the RAS VDI Agent. If you specified **Use dedicated VDI Agent** option in the previous (optional) step, but haven't installed the agent yet, click **Install** and follow the instructions to push install the agent on the specified host.

Please note that for the remote installation to work, the following requirements must be met:

- The firewall must be configured on the server to allow push installation. Standard SMB ports (139 and 445) need to be open. See also **Port reference** for the list of ports used by Parallels RAS.
- SMB access. The administrative share (\\server\c\$) must be accessible. Simple file sharing must be enabled.
- Your Parallels RAS administrator account must have permissions to perform a remote installation on the server. If it doesn't, you'll be asked to enter credentials of an account that does.
- The target server should be joined to an AD domain.

If push installation cannot be performed for any reason, you can install the agent manually using the installer. See **Installing RAS VDI Agent using the installer** (p. 137)

13 If you've selected **Microsoft Hyper-V Failover Cluster** as the VDI provider type, the page opens where you can disable MAC address management for guest VMs. Note that you should only do it if you are using Microsoft System Center Virtual Machine Manager (SCVMM) or other solution to manage MAC addresses. See the explanation below.

MAC address management is required when using Microsoft Hyper-V Failover Cluster as a VDI provider. This is to avoid duplicate MAC addresses, which may occur when a guest VM is migrated to a different node in the cluster and the MAC address is released and reused on the original node. If that happens, such a guest VM can no longer be managed in a Farm. Parallels RAS uses a pool of static MAC addresses at the VDI provider level to automatically generate and assign MAC addresses to guest VMs. This way, when a guest VM is migrated to a different node in the cluster, its MAC address will not be reused for a different VM and no duplicate MAC addresses will occur. The pool has 10,000 reserved MAC addresses in the range displayed in the **Starting MAC address** and **Ending MAC address** fields on the wizard page.

As was said above, if you are already managing MAC addresses using SCVMM or other solution, clear the **Enable MAC address management** option.

14 Click **Next**.

15 If you've selected **VMware vCenter** as the **VDI provider**, another page opens (the page will not open for any other host type). On this page, you can specify a vCenter resource pool. This allows you to enumerate VMs by selecting a cluster (root resource pool) or an individual resource pool within a cluster. To choose a resource pool, select the **Use specific resource pool** option and then click the [...] button next to the **Resource Pool** field. In the dialog that opens, select a desired resource pool. Note that if you leave the **Use specific resource pool** option cleared, all VMs from the entire vCenter cluster will be retrieved (max number is 35,000). Click **OK** when done.

16 Click **Finish** to close the wizard.

Add a cloud VDI provider

This section describes how to add a cloud-based VDI provider (p. 126). For the information on how to add a hypervisor provider, see **Add a hypervisor VDI provider** (p. 129).

Note: At the time of this writing, Parallels RAS supports Microsoft Azure as the only cloud VDI provider.

In this section:

- Introduction and prerequisites (p. 131)
- Create a Microsoft Azure AD application (p. 132)
- Add Microsoft Azure as a VDI provider (p. 135)
- Microsoft Azure and templates (p. 136)

Introduction and prerequisites

Introduction

Organizations using or interested in using Microsoft Azure can provision, scale, and manage VDI and RD Session Host workloads directly from the Parallels RAS console and deploy on to Microsoft Azure using Azure Resource Manager (ARM). Parallels RAS uses a service principal with required permissions on relevant Azure resources (subscription and resource groups) to authenticate, provision and manage the resources.

Prerequisites

To use Microsoft Azure as a VDI provider, you need the following:

- An existing Microsoft Azure account and subscription.
- The necessary Microsoft Azure providers must be enabled, including Microsoft.ResourceGraph, Microsoft.Resources, Microsoft.Compute, Microsoft.Network.
- An ARM virtual network and subnet in your preferred region with connectivity to AD services. Azure Active Directory with Active Directory Domain Services (AADDs), Domain Controller in Azure IAAS or hybrid with connectivity to on-premises domain can be used.
- Site-to-site VPN or ExpressRoute is required if hybrid RAS deployment is used.
- A configured VM to be used for VDI or RD Session Host as a template.

Adding Microsoft Azure as a VDI provider is a two-step process:

- 1 First, you need to create an application in Microsoft Azure to access the resources in your subscription. This step is described in the **Create a Microsoft Azure AD application** (p. 132) section.
- 2 Once the application is created and registered, you can add Microsoft Azure as a VDI provider in the Parallels RAS Console. This step is described in **Add Microsoft Azure as a VDI provider** (p. 135).

Read on to learn how to perform the steps above.

Create Microsoft Azure AD application

To complete the steps below, you must have a Microsoft Azure subscription and account. If you don't have a subscription, you need to purchase one first.

Create an Azure Active Directory application

An Azure Active Directory application is used with the role-based access control. You need to create an Azure AD application to access resources in your subscription from Parallels RAS.

To create an Azure AD application:

- 1 Log in to the Microsoft Azure portal.
- 2 Open the portal menu and select **Azure Active Directory**.
- 3 In the left pane, select **App registrations**.
- 4 Click **New registration** (at the top of the right pane).
- 5 The **Register an application** blade opens.
- 6 In the **Name** field, type a name you want to use for the application.

- 7 In the **Redirect URI (optional)** section, make sure that **Web** is selected in the drop-down list. Leave the URI field empty.
- 8 Click **Register** (at the bottom left).
- 9 The new Azure AD app is created and its blade is displayed in the portal.

Note the following app properties, which are displayed at the top of the right pane:

- **Display name**
- **Application (client) ID***
- **Directory (tenant) ID***
- **Object ID***

* Copy and save these properties. You will need to specify them later when adding Azure as a VDI provider in the RAS Console.

Create a client secret

A client secret is a string that the application uses to prove its identity when requesting a token. It essentially acts as an application password. You will need to specify this string in the RAS Console when adding Azure as a VDI provider.

To create a client secret:

- 1 If you are not on the application page anymore, navigate to it from the **Home** page by selecting **Azure Active Directory > App registration** and then clicking the app in the right pane.
- 2 In the left pane, click **Certificates & secrets**.
- 3 In the right pane, click **New client secret**.
- 4 Type a client name and select a desired expiration option.
- 5 Click **Add**. The new client secret appears in the **Client secrets** list.
- 6 **IMPORTANT:** Copy and save the client secret (the **Value** column). If you leave this page without copying the secret, it will be hidden and you will not be able to retrieve it later.

Give the application read and write access to resources

The Azure AD app that you created must have read and write access to Azure resources. The following instructions demonstrate how to give the application read and write access to a resource group. You can also give access to a specific resource or to your entire Azure subscription. For more information, please see the Microsoft Azure documentation.

To give the app write access to the resource group where new VMs will reside:

- 1 In the Azure portal menu, select **Resource groups**.
- 2 Click a resource group where the new VMs will reside.

- 3 In the left pane, select **Access control (IAM)**.
- 4 In the right pane, locate the **Add a role assignment** box and click **Add**.
- 5 In the **Add role assignment** dialog, select **Contributor** in the **Role** drop-down list.
- 6 In the **Assign access to** field, select **Azure AD user, group, or service principle**.
- 7 In the **Select** field, begin typing the name of the app that you created earlier. Once the app is found, select it.
- 8 Click **Save**.

To give the app read access to the resource group:

- 1 Repeat steps 1-4 from the list above.
- 2 In the **Add role assignment** dialog, select **Reader** in the **Role** drop-down list.
- 3 Select the application from the list (use the **Search** field to search for the application).
- 4 Click **Save**.

Note: If you would like to give the application read access to your entire subscription (not just a specific resource groups), select **All services** in the Azure portal menu, then navigate to **Categories > All > Subscriptions** and select your subscription. Select **Access control (IAM)** in the middle pane and click **Add** in the **Add a role assignment** box. Repeat steps 2-4 from the list above.

Finding your Microsoft Azure subscription ID

When you'll be adding Microsoft Azure as a VDI provider in the RAS Console, you will need to specify your Azure subscription ID. If you don't remember it, here's how to find it in the Microsoft Azure portal:

- 1 In the portal menu, choose **All services**.
- 2 In the **Categories** list, click **All**.
- 3 In the right pane, click **Subscriptions**.
- 4 Click a subscription and then copy and save the value from the **Subscription ID** field.

Summary

When you complete all of the above steps, you should have the following values saved and ready to be used to add Microsoft Azure as a VDI provider in the RAS Console:

- **App (client) ID:** Application ID.
- **Directory (tenant) ID:** Tenant ID.
- **Client secret:** Client secret (application key).
- **Subscription ID:** Your Microsoft Azure subscription ID.

Read on to learn how to add Microsoft Azure as a VDI provider in the RAS Console.

Add Microsoft Azure as a VDI provider

To add Microsoft Azure as a VDI provider:

- 1 In the RAS Console, navigate to **Farm > Site > VDI**.
- 2 On the **Providers** tab, click **Tasks > Add**.
- 3 The **Add VDI Provider** wizard opens.
- 4 On the **Select VDI provider type** page, select **Cloud computing**.
- 5 Click **Next**.
- 6 The page opens where you configure Microsoft Azure as a VDI provider. The properties that you need to specify on this page are described below.

General properties

At the top of the page, you need to type a name for the new host and an optional description if you wish.

Subscription details

In the **Subscription details** section, specify the following:

- **Authentication URL:** Prepopulated with the Microsoft authentication site URL. Unless otherwise required or indicated, keep the default value provided.
- **Management URL:** Prepopulated with the Microsoft Azure management site URL. Unless otherwise required or indicated, keep the default value provided.
- **Resource URI:** Prepopulated with the Microsoft Azure resource URI. Unless otherwise required or indicated, keep the default value provided.
- **Tenant ID:** The "Directory (tenant) ID" value of the Azure AD app that you created earlier.
- **Subscription ID:** Your Microsoft subscription ID.

Service principle details

In the **Service principle details section**, specify the following:

- **Application ID:** The "App (client) ID" value of the Azure AD app that you created earlier (p. 132).
- **Application key:** The "Client secret" value of the Azure AD app that you created earlier (p. 132).

Advanced Settings

Click the **Advanced Settings** link to open a dialog where you can configure the following optional settings:

- **Use dedicated VDI Agent:** When this option is cleared (default), the built-in RAS VDI Agent will be used. If you want to use a dedicated RAS VDI Agent, select this option and specify the server FQDN or IP address.
- **Preferred Publishing Agent:** Select a RAS Publishing Agent to be the preferred agent for this VDI provider. Select **Automatic** to let the system select an agent (this option is enabled and selected by default if you have at least three Publishing Agents installed). The automatic selection is especially important to ensure that the VDI provider and guest VMs always have a VDI Agent they communicate with in case of other Publishing Agent(s) failure. For more info, see **Enabling high availability for VDI** (p. 142).

Click **OK** to close the **Advanced VDI Provider Settings** dialog.

Complete the wizard

When done entering the Microsoft Azure information, click **Next** in the **Add VDI Provider** wizard. The wizard will display the new VDI provider information and will indicate the RAS VDI Agent status. If everything is OK, click **Finish** to exit the wizard. If something is not as expected, click **Back** and correct any mistakes if necessary.

The new VDI provider will now appear on the **Providers** tab in the RAS Console. Complete the VDI provider addition as follows:

- 1 Click **Apply** to apply the changes.
- 2 Verify the value of the **Status** column. If it's anything other than **OK**, right-click the VDI provider and choose **Troubleshooting > Check agent**. Verify the agent status and install it if necessary, then click **OK**. The **Status** column on the **Providers** tab should now say **OK**.

Modifying the VDI provider configuration

To view and modify the VDI provider configuration, right-click it and choose **Properties**. In the dialog that opens, view and modify the VDI provider properties.

Microsoft Azure and templates

When creating a template for cloning VMs in Microsoft Azure, you need to select an Azure resource group where VM clones will be created. Note that this must be a group to which you granted permissions to the Azure AD application. You also need to select a VM size and disk type to be used for cloned VMs. These settings are specified on the **Advanced** page of the **Create Template Wizard**.

Both Virtual Desktop and RD Session Host templates can be created with Microsoft Azure as a VDI provider. When VMs are cloned, you will see them appear in the RAS Console. At the same time, you can also see them in the Microsoft Azure portal.

Note: If there are multiple RAS installations using the same subscription, then the workaround is to change the VDI agent application read access from subscription level to resource group level or a set of resource groups. This is necessary to avoid a situation when a given VDI Agent intersects with the set of resource groups of another VDI agent application.

For complete information about creating and using templates, including Microsoft Azure specifics, please see the **Templates** section (p. 147).

Installing RAS VDI Agent using the installer

This topic describes a workaround for a situation when you want to install a dedicated RAS VDI Agent on a specific server, but the push installation from the RAS Console cannot be performed for any reason. If that happens, you can install the agent by running the installer directly on the target server.

Note: You can only use these instructions to install RAS VDI Agent in Windows.

To install the dedicated RAS VDI Agent:

- 1 Log in to the server where you want RAS VDI Agent installed using an administrator account and close all other applications.

Copy the standard Parallels RAS installer (RASInstaller.msi) to the server and run it:

- 1 When you get to the **Select Installation Type** page, select **Custom** and click **Next**.
- 2 Click on **RAS VDI Agent dedicated** and select **Entire Feature will be installed on local hard drive** from the drop-down menu.
- 3 Ensure that all other components are cleared (excluded from the installation) and click **Next**.
- 4 Click **Install** and follow the onscreen instruction to install the agent.

The RAS VDI Agent does not require any configuration. Once it is installed, go back to the RAS Console, highlight the server name and click **Troubleshooting > Check Agent**. If the agent is installed properly, the status should change to **Agent Installed**.

To uninstall the RAS VDI Agent from a server:

- 1 Navigate to **Start > Control Panel > Programs > Uninstall a Program**.
- 2 Find **Parallels Remote Application Server** in the list of installed programs.
- 3 If you don't have any other Parallels RAS components on the server that you want to keep, right-click **Parallels Remote Application Server** and then click **Uninstall**. Follow the instructions to uninstall the program. You may skip the rest of these instructions.

- 4 If you have other RAS components that you want to keep on the server, right-click **Parallels Remote Application Server** and then click **Change**.
- 5 Click **Next** on the Welcome page.
- 6 On the **Change, repair, or remove** page, select **Change**.
- 7 On the next page, select **Custom**.
- 8 Select **RAS VDI Agent dedicated**, then click the drop-down menu in front of it, and click **Entire feature will be unavailable**.
- 9 Click **Next** and complete the wizard.

Checking the RAS VDI Agent status

To verify that the RAS VDI Agent is installed and functions properly, do the following:

- 1 First, you can look at the **Status** column in the **Farm > Site > VDI > Providers** list. If there's a problem with the agent, the column will display an appropriate description. Note that in addition to the description, the **Status** column uses a color code to indicate the agent status as follows:
 - Red — Not Verified
 - Orange — Needs Update
 - Green — Verified
- 2 Right-click a host and then click **Troubleshooting > Check agent** in the context menu.
- 3 The **VDI Agent Information** dialog opens displaying the information about the VDI Agent, VDI Services, and other related info.
- 4 If the VDI Agent is not installed, click the **Install** button and follow the onscreen instructions. See **RAS VDI Agent installation options (p. 128)** for more info.

Modifying VDI Provider configuration

Read this section to learn how to modify the configuration of a VDI provider in Parallels RAS.

Configure a VDI provider

To configure an existing VDI provider:

- 1 In the RAS Console, navigate to **Farm > <Site> > VDI**.
- 2 Select the **Providers** tab in the right pane.
- 3 Select a VDI provider and click **Tasks > Properties**. The **Properties** dialog opens.

Note: Some of the properties described below may be unavailable on some servers. This depends on the VDI provider type.

Enable or disable a VDI provider in Site

By default a VDI provider is enabled. To enable or disable a VDI provider, use the **Enable provider in site** option on the **Properties** tab.

Properties: configure VDI provider connection settings

The **Properties** tab has different properties depending on whether it's a hypervisor-based or cloud-based provider.

Hypervisor VDI provider:

- **Type:** VDI Provider type.
- **Subtype:** Hypervisor version. If the hypervisor version that you are using is not listed, select **Other**.
- **Host:** The VDI provider host IP address.
- **Port:** Port number on which the VDI provider listens for incoming connections.
- **Resource pool:** This field is enabled for VMware vCenter only. If you've specified a vCenter resource pool while adding a VDI provider, the pool will be displayed here. The [...] button allows you to specify a different pool (or select one if the field is empty), but only if no guest VMs from the current pool have been created or used in Parallels RAS in any way. If Parallels RAS detects any current usage, you will see a warning message and will not be able to change it. If you still want to select a different resource pool, you'll have manually do a full clean up in the RAS Console, so that no usage of any kind exists.
- **Description:** An optional description.
- **Dedicated VDI Agent:** Select this option if you have a dedicated RAS VDI Agent installed on a different server. Enter the server FQDN or IP address in the **Agent address** field.

Cloud-based VDI provider:

- **Type:** Cloud-based VDI provider type (e.g. Microsoft Azure).
- **Name:** VDI provider name.
- **Description:** An optional description.
- For description of the remaining properties, please see **Add Microsoft Azure as a VDI provider** (p. 135).

Credentials: configure username and password

The **Credentials** tab has different properties depending on whether it's a hypervisor-based or cloud-based host.

Hypervisor VDI provider:

- Specify the username and password to log in to the VDI provider. Click the **Check Credentials** button to verify the credentials that you've entered.

Cloud-based VDI provider:

- For description of the Microsoft Azure properties, see **Add Microsoft Azure as a VDI provider (p. 135)**.

Agent settings: configure RAS VDI Agent

RAS VDI Agent can be configured on the **Agent Settings** tab:

- **Preferred Publishing Agent:** Select a preferred Publishing Agent to which this VDI provider should be assigned. This can be helpful when Site components are installed in multiple physical locations communicating through WAN. You can decrease network traffic by specifying a more appropriate Publishing Agent.
- **Allow URL/Mail redirection:** When a user tries to open a URL or an HTML Mailto link in a remote application, the link can be redirected to the client computer and open in a local default application (a web browser or email client) instead of an application on the remote host. This option allows you to enable or disable the redirection. Click the **Configure** button to choose from the following options:
 - a Replace registered application** — this option uses an alternative method of redirecting a link. It replaces the default web browser and mail client with "dummy" apps on the remote server side. By doing so, it can intercept an attempt to open a link and redirect it to the client computer.
 - b Support Windows Shell URL namespace objects** — the Shell URL namespace objects support means that Parallels RAS can intercept actions in published applications that use Shell namespace API to open links, which is a standard behavior in most applications. The ability to disable the support for Shell URL namespace objects is for compatibility with older versions of Parallels RAS.
- **Enable drag and drop:** Allows you to set how the drag and drop functionality works in Parallels Clients. Click **Configure** and choose from "Disabled" (no drag and drop functionality), "Sever to client only" (drag and drop to a local application only), "Client to server only" (drag and drop to a remote application only), "Bidirectional" (drag and drop in both directions).

Note: At the time of this writing, the drag and drop functionality is only supported on Parallels Client for Windows and Parallels Client for Mac.

- **Allow file transfer command (HTML5 and Chrome clients):** Enables file transfer in a remote session. To enable file transfer, select this option and click the **Configure** button. For more information, see **Configuring remote file transfer (p. 408)**.
- **Enable drive redirection cache:** Improves user experience by making file browsing and navigation on redirected drives much faster. For details, see **Drive redirection cache (p. 107)**.

RDP printer

The **RDP Printer** tab allows you to configure the renaming format of redirected printers. The format may vary depending on which version and language of the server you are using. Select the **RDP Printer Name Format** option specifically for the configured server:

- **Printrname (from Computername) in Session no.**
- **Session no. (computername from) Printrname**
- **Printrname (redirected Session no)**

The other RDP Printing option available is **Remove session number from printer name**, which will do what it says.

Scheduler: configure VDI provider maintenance time window

The **Scheduler** tab page allows you to create a maintenance time window for the server. During this time, published resources hosted on the VDI provider will not be available to end users.

Note: When the scheduled maintenance is triggered, the server is disabled in Parallels RAS and its status on the **VDI > Providers** tab is displayed as "Disabled (Scheduler)". You can cancel the disabled state at any time without waiting for the maintenance time window to end. To do so, on the **VDI > Providers** tab, select the server, click **Tasks** (or right-click) and then choose **Cancel disabled state (scheduler)**.

To configure maintenance time window click **Tasks > Add** and then set the following options:

- **Start date**
- **Time**
- **Duration**
- **Repeat**

The **On disable** option allows you to specify what should happen to current sessions when a scheduled task triggers.

MAC addresses

This tab is only displayed for Microsoft Hyper-V Failover Cluster as the VDI provider. It is used to enable or disable MAC address management for guest VMs. For more information, please see **Add a hypervisor host** (p. 129) (read the description of the step where the MAC address management is configured).

Please note that MAC address management is available in Parallels RAS since version 18. In new Parallels RAS 18 installations, this functionality is enabled by default when a VDI provider of type Microsoft Hyper-V Failover Cluster is added to a farm. In older Parallels RAS versions, the functionality is disabled for existing VDI providers, but is enabled by default when a new VDI provider is added.

Configure logging

To configure logging and retrieve or clear existing log files, right-click a VDI provider, choose **Troubleshooting > Logging** in the context menu, and then click one of the following, depending on what you would like to do: **Configure**, **Retrieve**, or **Clear**. For the information on how to perform these tasks, see the **Logging** (p. 451) section. Please also read the important information below.

Note that logging of VDI provider operations is performed on the RAS VDI Agent level. When you configure logging for a VDI provider, you are essentially configuring it for the RAS VDI Agent that services this VDI provider. This means that if you are using the built-in RAS VDI Agent, its logging configuration applies to all VDI providers that it services. Consider the following scenarios:

- When you retrieve log files for a specific VDI provider serviced by the built-in VDI Agent, the files will contain logs for all VDI providers serviced by the same agent.
- If you clear log files for a particular VDI provider, you should be careful because the logs will be cleared for all VDI providers if they are serviced by the same built-in VDI agent. The RAS Console will prompt you if you try to delete such a shared log.

If a VDI provider has a dedicated VDI Agent, which services this host only, none of the above applies.

Enabling high availability for VDI

High availability for VDI means that a VDI provider must never lose a connection with a VDI Agent. If the connection is lost, the guest VMs will become unavailable for user connections. High availability for VDI is accomplished by installing at least three RAS Publishing Agents and configuring a VDI provider to automatically select one of them. This way, if one of the Publishing Agents goes offline (and with it the built-in VDI Agent), the VDI provider will be automatically assigned to the VDI Agent running in the next available Publishing Agent.

To configure high availability for VDI, use the information and instructions below.

At least three Publishing Agents are required

Make sure you have at least three RAS Publishing Agents installed and running. You may also have additional Publishing Agents in standby mode, but you must have at least three agents in the active state for the high availability functionality to work. All Publishing Agents must be able to communicate with each other.

An odd number of agents is recommended

To properly control a possible split-brain situation, strictly more than half of all available Publishing Agents should be able to communicate with each other at any given time. Consider the following examples:

- Let's say there are three Publishing Agents in a Site. All of them can communicate with each other. If one of the agents suddenly loses a connection with the other two, the two agents will know that they are in the majority and will take over the VDI provider hosts that are currently managed by the first agent.
- Let's now say that there are four Publishing Agents. If one of them loses a connection to the remaining three, the same scenario will occur as in the example above. But if two agents simultaneously lose a connection to the other two, none of the two groups will be in the majority and therefore none will be able to make a decision who should take over the VDI provider hosts. In a situation like this, steps must be taken to prevent a split-brain scenario, which will happen if the agents continue to operate independently from each other. As a solution to this problem, all agents will simply abandon all VDI providers at the same time, so no data loss or any other problem can possibly happen.

For the reasons explained above, you should always install an odd number of Publishing Agents. This way, one of the groups of agents will always be in the majority and will continue to handle all VDI providers. Please note that the general recommendation (regardless of the high availability functionality described here) is to have three RAS Publishing Agents running in a Site. For details, see **Secondary Publishing Agents** (p. 59).

Please also note that Publishing Agents in standby mode (p. 57) don't participate in the high availability operations. These agents stay inactive until one of the active Publishing Agent goes completely offline. When that happens, an agent in standby mode is activated and takes place of the lost agent. From this point forward, it is considered a part of the high availability setup. When the lost agent is brought back online, everything goes back to what it was before.

Configuring a VDI provider for high availability

Use one of the following to configure a VDI provider for high availability:

- For an existing VDI provider, open the **Properties** dialog, select the **Agent Settings** tab and in the **Preferred Publishing Agent** field, select **Automatic**.

- When adding a new VDI provider, on the second wizard page where you specify the host type and address, click the **Advanced Settings** link and then select **Automatically** in the **Preferred Publishing Agent** drop-down box. Note that the **Automatic** option is selected by default when there are three or more Publishing Agents available.

If initially you have less than three RAS Publishing Agents in a Site, a VDI provider is assigned a specific preferred Publishing Agent. When at some point you increase the number of Publishing Agents to three or more and want to enable high availability for one or more VDI providers, you need to reconfigure each host so that the **Preferred Publishing Agent** property is set to **Automatic**.

Change VDI provider Site assignment

You can assign a VDI provider to a different Site in your Farm if needed. Please note that this functionality is only available if you have more than one Site in your Farm.

Note: You cannot assign a VDI provider to a different Site when there are templates, pools, or guest VMs that are in use on the current Site. If you try to do so, you will get an error and will not be able to proceed. To assign such a VDI provider to a different Site, you need to remove all dependencies in the current Site first.

To change the Site assignment:

- 1 Right-click a VDI provider and then click **Change Site** in the context menu. The **Change Site** dialog opens.
- 2 Select a Site in the list and click **OK**. The server will be moved to the **Providers** list on the target Site (**Farm** > <new-site-name> > **VDI / Providers**).

Site defaults (VDI)

Site defaults are settings that are defined on a Site level and can be used by templates and guest VMs (both template-based and non-template based). By default, templates (described later in this chapter) inherit Site default settings, but you can override them if needed when you configure a template. Non-template based guest VMs also use Site default settings by default and you can also override them if needed when you configure these VMs.

To view and modify Site defaults, do the following:

- 1 Navigate to **Farm** > <Site> > **VDI**.
- 2 Select the **Desktops** tab in the right pane.
- 3 Click **Tasks** > **Site defaults**. This opens the **Site Default Properties** dialog, which is described below.

Note that any modifications you make to Site defaults are immediately applied to all guest VMs in the current Site that use them.

General

The **General** tab contains the following properties:

- **Session readiness timeout:** The maximum amount of time it should require to establish a session. If the specified timeout is reached, and the session is still not ready, the user will see an error message and will have to try to log in again.
- **Protocol:** Specifies a protocol that Parallels RAS uses to communicate with a guest VM.
- **Auto remove guest VMs which failed preparation after:** If a guest VM encounters a problem during the preparation stage (for any reason), it remains on the server but cannot be used. You can identify such VMs by the "Failed to create" value in the **Guest VM state** column (**Farm** > <Site> > **VDI** > **Desktops**). Unless a VM like this is repaired, it will be automatically removed after the time period specified in this field. You can set any of the available time periods by selecting it from the drop-down list or you can type a desired value, such as "8 days" or "12 hours".
- **Auto remove persistence if guest was not used for:** The time period after which persistence should be automatically removed. You can also type any desired time period, such as "1 week 3 days".

Note: Beginning with RAS 17, the default setting for this option is **Never**. Please keep that in mind.

Settings

The **Settings** tab contains the following:

- **Logoff active session after:** The amount of time a session remains logged in after the user closes a published application. The default timeout is 25 seconds. Note that this only works for applications, but not published desktops (when a user closes a desktop, the session is logged off). This timeout is used to avoid unnecessary logins when a user closes one application and then opens another.
- **Actions:** The two drop-down lists here specify an action to perform on session disconnect or logoff.

Note for Nutanix Acropolis users: Nutanix Acropolis does not support the suspend operation for its VMs. If **Suspend** is selected in the **Perform action** field, no action will be applied to a Nutanix Acropolis VM when a session disconnect occurs (a corresponding error will be recorded in the VDI Agent log).

Security

On the **Security** tab, you can specify whether to automatically grant users Remote Desktop connection permissions on guest VMs. Here's how it works. Instead of manually adding each user to the Remote Desktop Users (or Administrators) group, you can enable this option to do it automatically. When a user logs on, he/she will be automatically added to the specified group and will therefore have the Remote Desktop connection (or full Administrator) permissions on the server. When the user logs off, they will be removed from the group (i.e. the group membership will only exist for the duration of the session).

The more important benefits of this feature are as follows:

- You don't have to permanently add your users to the Remote Desktop Users groups. This way, a user will never be able to establish a Remote Desktop session with a server outside of Parallels Client.
- By automatically adding a user to the Administrators group, you can give them rights to install applications and perform other administrative tasks. Once again, the user will only be able to do it from Parallels Client but never by connecting to the server using standard Remote Desktop tools.

Viewing guest VMs on a VDI provider

To view all guest VMs that exist on a VDI provider, including VMs that are not currently managed in Parallels RAS, do the following:

- 1 Navigate to **Farm** > <Site> > **VDI** > **Providers**.
- 2 Select a VDI provider for which you want to see the guest VM information and click the **Guest VM list** button (at the bottom of the window, below the VDI provider list).
- 3 The **Guest VM List** dialog opens listing guest VMs.

The dialog displays all guest VMs that exist on the selected VDI provider. This includes template-based and non template-based VMs (which are already managed in Parallels RAS) but also includes virtual machines that were created using native VDI provider tools outside Parallels RAS. The main purpose of this list is to give a RAS administrator a convenient overview of which virtual machines are available on the VDI provider.

While in the **Guest VM List** dialog, you can perform the following tasks on guest VMs:

- 1 Install or update the RAS Guest Agent in any guest VM, including unmanaged VMs.
- 2 Perform power operations on guest VMs (start, stop, restart, suspend, reset). Note that when restart (graceful) is used, there's a 10 min timeout. If the operation cannot be completed during this time, the reset operation (forced) will be used. Please also note that if you are using **Nutanix Acropolis**, the suspend operation is not available. The reason for this is Nutanix Acropolis does not support the suspend operation on its virtual machines.

- 3 Use the provided tools to perform standard computer management tasks on a guest VM, such as establishing a remote desktop connection, pinging, and others.
- 4 Viewing guest VM properties.

Please note that the same tasks (and more) can be performed on managed guest VMs in the **VDI > Desktops** tab. For more information, see **Managing guest VMs** (p. 172).

Templates

Templates are used to automate the creation and deployment of guest VMs in Parallels RAS. A template is based on an existing virtual machine created with one of the hypervisors supported by Parallels RAS. Once a template is ready, it can be used to create clones (guest VMs) that will inherit all properties of the template. The resulting guest VMs can then be used to host published resources.

Read the following topics to learn how to create and use a template:

- Template types and guest OS requirements (p. 147)
- Creating a template (p. 150)
- How guest VMs are created from a template (p. 161)
- Manually adding a guest VM (p. 162)
- Template maintenance (p. 163)
- Template status (p. 166)
- Managing multi-provider template distribution (p. 168)
- Managing template-based guest VMs (p. 169)

Template types

There are two types of templates in RAS VDI: **Virtual desktop** and **RD Session Host**. They are described in the following subsections.

Virtual desktop templates

Virtual desktop templates are the essential part of the Parallels RAS VDI. They are used to create guest VMs for publishing of desktops, applications, documents, etc. The guest OS support is the same as for RAS Guest Agent (which must be installed in a VM). See **Software requirements** (p. 20).

Guest VMs created from a Virtual desktop template normally serve a single user. They are managed entirely from within the RAS VDI, which includes such features as creating persistent VMs, managing VDI sessions, publishing resources from a specific Virtual desktop template, and others.

RD Session Host templates

RD Session Host templates are designed specifically to give you the ability to replicate RD Session Hosts running in virtual machines. Guest VMs created from an RD Session Host template are treated by Parallels RAS almost like regular RD Session Hosts. The main difference is, you can create as many guest VMs from a single template as you require, thus automating RD Session Host provisioning according to your needs.

RD Session Host templates are supported on the following VDI platforms:

- Microsoft Hyper-V
- Microsoft Hyper-V Failover Cluster
- VMware VCenter
- VMware ESXi
- Scale Computing HC3
- Nutanix Acropolis
- Microsoft Azure

RD Session Host templates support Windows Server 2008 R2 up to Windows Server 2022 as a guest OS. Compared to regular RD Session Hosts (p. 83), servers created from an RD Session Host template do not support earlier versions of Windows Server. The reason is, these servers run in VMs and require the RAS Guest Agent installed in them, so the guest OS requirements are limited by Windows Server versions supported by RAS Guest Agent.

Please note that the following standard RAS VDI features are not available when using RD Session Host templates:

- Pool management
- Persistent guest VMs
- Session management
- Publishing from a specific Template
- Some other strictly RAS VDI specific features.

For the information on how to provision RD Session Hosts created from a template, see **Grouping and cloning RD Session Hosts** (p. 108).

Multi-provider template distribution

When you create a template, it is normally managed by a single VDI provider, which is the same provider to which the source virtual machine belongs. Clones are deployed from the template and run on the same VDI provider, which was used to create the template. Cloned guest VMs are usually stored on a centrally shared storage, such as Storage Area Network (SAN).

Beginning with Parallels RAS 18, admins have the ability to create a template and distribute it to multiple Microsoft Hyper-V hosts. As a result, the template configuration, as seen in the Parallels RAS Console, is shared among multiple Microsoft Hyper-V hosts, while each host has its own copy of the template residing on its local storage. This makes it possible to deploy clones not only to a centrally shared storage, but also to local disks of multiple independent Microsoft Hyper-V hosts. Scaling out is easily carried out by adding as many Microsoft Hyper-V hosts to the template distribution list as necessary.

Template distribution is configured on the **Distribution** page of the **Create Parallels Template Wizard**, which is described in the subsequent sections. If you are planning on using the template distribution functionality, please read the **Prerequisites** subsection below before running the wizard.

For more information about managing multi-provider distribution for a template, see also **Managing multi-provider template distribution** (p. 168).

Prerequisites

- Template distribution is supported on a standalone Microsoft Hyper-V Server 2012R, 2016, 2019, 2022.
- All target VDI providers must have identical:
 - Provider type and subtype.
 - Folder path where guest VMs will reside.
 - Virtual switch name to which guest VMs are connected.
- Hyper-V hosts must be domain-joined. The current implementation uses a full VM copy of the template to distribute the template to other hosts (local storage) via the Hyper-V Live Migration mechanism.

Note: Full clones can also be moved to other hosts via Live Migration, but the process is time-consuming (equal to the first copying of the template).

- The Microsoft Hyper-V server hosting the source VM may also be used as a target host.
- Always ensure that enough storage space is available prior to choosing target hosts to which the template will be distributed and on which clones will be created.
- Hyper-V settings must have Kerberos authentication enabled and appropriate delegations configured in AD:
 - Go to Hyper-V settings for the host machines and enable Live Migration using Kerberos.
 - Go to Active Directory Users and Computers and for each Hyper-V host server enable delegation for "cifs" and "Microsoft Virtual System Migration Service" for all servers you want to migrate To and From.

Note: If authentication isn't working, try changing the "Use any authentication protocol" option.

Creating a template

Requirements

To complete the tasks described in this section, the following requirements must be met:

- For hypervisor-based hosts, make sure the hypervisor tools are installed and running in the guest VM.
- Make sure you know account credentials that will allow you to push install the agent software on a VM. If you run the Parallels RAS console using such credentials (e.g. a domain admin), you will not be asked to enter them during the agent installation. If you run the console using a different account, you'll be asked to enter credentials when you install the agent.
- The guest OS (Windows) running in the VM must be configured to obtain an IP address from a DHCP server.
- For users to access published resources in a guest VM, the RDP port must be open locally or via Group Policy in Windows running in the VM. The default RDP port is 3389.
- For RD Session Host templates, Network Discovery UDP port 137 must be enabled for a domain firewall profile in the guest OS. This can be done via domain group policies or manually in the guest OS.

Manual agent installation

Normally, you will push install the necessary agent software in a source VM right from the Parallels RAS console (as described later in this section). However, you can also install the software manually by running the Parallels RAS installer in Windows in the VM. When doing so, use the **Custom** installation option and select the following agent components depending on the type of the template (p. 147) that you are creating:

- **Virtual desktop.** This template type requires RAS Guest Agent to be installed in the source VM.
- **RD Session Host.** This template type requires RAS Guest Agent and RAS RD Session Host Agent to be installed in the source VM.

Create a template

The process of creating a template consists of two stages:

Stage 1 (p. 151): During the first stage, the wizard will verify if the agent software is installed and will give you an option to install it if needed.

Stage 2 (p. 152): During the second stage, you configure the template.

Each stage is described in detail in the sections that follow this one (or follow the links above).

Stage 1: Check and install the Agent

To begin creating a template:

- 1 In the RAS Console, navigate to **Farm** > <Site> > **VDI**.
- 2 Select the **Templates** tab in the right pane.
- 3 In the **Tasks** drop-down menu, click **Add** (or click the "+" icon)
- 4 In the dialog that opens, select a guest VM from which you would like to create a template and click **OK**.
- 5 The **Create Parallels Template Wizard** opens. Each wizard page is described below in the order they appear on the screen.

Template type

On the first page of the wizard, select a template type to create: **Virtual desktop** or **RD session host**. For details, see **Template types** (p. 147).

Check agent

This step only appears if an on-premise VDI provider is used. It will not appear for Azure Virtual Desktop and Azure VDI.

In this step, the wizard will check if the selected VM has the RAS Guest Agent installed. Wait for it to finish and then examine the **Status** field (closer to the bottom of the page). Depending on the result, do one of the following:

- If the agent is installed, click **Next** to continue. You may stop reading here and jump to **Stage 2: Configure the template** (p. 152).
- If the agent is not installed, you need to install it as described below.

To install the agent, first click the **Customize Guest Agent deployment settings** link and specify the options in the dialog that opens. None of the options are forced, so you can select or clear them depending on your needs. Note that depending on the template type, the options are different, as described below.

Virtual desktop:

- **Add firewall rules:** Automatically configure firewall rules in the guest VM.
- **Allow remote desktop connections:** Select to automatically configure remote desktop access in the VM.
- **Specify users or groups to be added to the Remote Desktop Users group:** Select this option and then click the **[+]** icon to add specific users to the group.

RD Session Host:

- **Add firewall rules:** Automatically configure firewall rules in the guest VM.

Note: Network Discovery UDP port 137 must be enabled for a domain firewall profile in the guest OS as a separate step. This can be done via domain group policies or manually in the guest OS.

- **Install RDS role:** Install the RDS role in the guest VM.
- **Enable Desktop Experience:** Enable the Desktop Experience feature in Windows.
- **Restart server if required:** Restart the VM if required.
- **Specify users or groups to be added to the Remote Desktop Users group:** Select this option and then click the **[+]** icon to add specific users to the group.

When done specifying the options, click **OK** to close the dialog.

Now click the **Install** button and follow the onscreen instructions to install the agent software.

Hint: If the guest VM cannot be reached by its name specified as hostname, double-click the guest VM name and change it to the correct IP address.

Once done, verify that the agent is installed by looking at the **Status** field on the **Check Agent** wizard page. If so, continue to the next section that describes **Stage 2: Configure the template** (p. 152).

Stage 2: Configure the template

Once the agent is installed, and the **Status** field on the **Check Agent** wizard page confirms this, click **Next**. The VM will now be powered off (wait for the power off operation to finish). Once the VM is powered off, the template configuration stage begins.

The subsequent wizard pages are described in the sections that follow this one. Please note that many of the wizard pages inherit the information from Site default settings, but you can override it if needed. To specify your own settings, clear the **Inherit default settings** option. To see and edit default settings, click the **Edit Defaults** link. For more information, see **Site defaults** (p. 144).

User profile

Note: The **User profile** page appears only for templates of type **Virtual desktop. RD session host** templates don't have it. Instead, the User profile functionality for RD Session Hosts created from a template is configured on the Group level: **Farm > Site > RD Session Hosts > Groups**.

This page allows you to configure FSLogix Profile Container for the template. FSLogix is a remote profile solution that maintains user context in non-persistent environments, minimizes sign-in times and provides native profile experience eliminating compatibility issues.

For the complete description of how to configure FSLogix, please see **FSLogix Profile Container** (p. 97).

After reading the **FSLogix Profile Container** section mentioned above, please note the following VDI specifics when configuring FSLogix for VDI:

- To configure FSLogix in Site defaults, navigate to **Farm > Site** and click **Tasks > Site defaults > VDI > User profile**. The settings are inherited by a VDI template and standalone desktops by default.
- When FSLogix is enabled for the template, the FSLogix software is installed automatically into the guest VM used as a source for the template. The installation is performed before the guest VM leaves the maintenance mode.
- For existing templates, FSLogix settings can be edited and an FSLogix upgrade can be performed only while the template is in maintenance. All necessary changes are performed before it leaves the maintenance mode. If necessary, the template (the virtual machine) is restarted automatically.
- If a running guest VM that has FSLogix configured has issues with the FSLogix agent, it has a status **FSLogix not available**. Hosts with this status are evicted from load balancing until the issue is resolved.

Properties

On the **Properties** page, specify the following options:

- **Template name:** Choose and type a template name.
- **Maximum guest VMs:** Specify the maximum number of guest VMs that can be created from this template.
- **Number of guest VMs deployed on wizard completion:** The number of guest VMs to deploy once the template is created. Please keep in mind that this will take some time because the VMs will be created one at a time.
- **Guest VM name:** A pattern to use when naming new guest VMs. For details, see **Guest VM naming** (p. 159).
- **Clone method:** Whether to create linked or full clones. A full clone is a complete copy of a template. As such, it occupies as much space on the physical hard drive as the source template and takes a significant time to create. A linked clone is a copy of a template made from a snapshot that shares virtual disk with the source template, therefore it takes much less space on the physical hard drive and it takes only a couple of minutes to create.

You should use full clones if your application and OS updates are too slow (full clones take longer to create, but they provide the best possible performance). Otherwise if your updates are fast enough, use linked clones as it takes much less time to create them.

Note: If the **Create a linked clone** option is grayed out, it means that the current version of Parallels RAS does not support linked clones with the VDI provider that you are using. At the time of this writing, support for linked clones is available for VMware, Microsoft Hyper-V, Scale Computing HC3, and Nutanix Acropolis.

Additional properties

This page allows you to optimize the VDI provider resources usage by controlling the power state of guest VMs after they are prepared. When you deploy a large number of guest VMs and run them automatically after they are ready, it can lead to the VDI provider resources overcommitment. Additionally, in Microsoft Azure, needlessly allocated guest VMs can also lead to unnecessary fees. You can control the guest VM power state and other related settings by specify the following options:

- **Keep available buffer:** The minimum number of guest VMs to always keep unassigned and session free for the template. As soon as the number of free and unassigned desktops drops below the setting value, it forces the template to create another guest VM. The template uses its own settings for guest VM creation including initial power state.
- **Guest VM state after the preparation:** Select the power state that should be applied to a guest VM after it is prepared. Choose from **Powered on**, **Powered off**, or **Suspended**. Note that when the power state is set to **Power off** or **Suspended**, the number of running (fully ready and waiting for incoming connections) guest VMs is controlled by the **Keep available buffer** setting (see above). For example, let's say the maximum guest VMs value is set at 200, the number of guest VM deployed on wizard completion is 100, and the power state after preparation is "Powered off". The result of such a configuration will be 100 clones deployed and powered off.
- **Delete unused guest VMs after:** Select what to do with unused guest VMs to save resources. Choose whether to never delete them or specify the time period after which they should be deleted.

Distribution

This page is used to configure template distribution to multiple Microsoft Hyper-V hosts. Note that this page will only appear if the source VM is a Microsoft Hyper-V machine. For the description of this feature and requirements, please see **Multi-provider template distribution** (p. 148).

To configure template distribution:

- 1** Select the **Enable multi-provider template distribution** option.
- 2** In the **Available** list, select one or more providers and click **Add** (or **Add all** to add all available providers). Note that only providers of the same type and subtype as the source VM are displayed in this list.
- 3** In the **Number of providers for concurrent distribution** field, specify the number of concurrent distribution operations. The template is distributed to target hosts using Hyper-V Live Migration, which first exports the virtual machine to a file and then moves it to the destination host. For each host in the **Target** list, a Live Migration operation must be performed. The number specified here dictates how many network copy operations should be started at the same time. The larger the number, the more network resources will be required. Note that virtual machine exports (the first step of Live Migration) are always done one VM at a time, so the number you specify here affect only the copy operations.

Note: The **Enable multi-provider template distribution** setting cannot be modified (selected or cleared) once the template is created. If later you decide to turn it on or off (enable or disable the feature), you will need to delete and re-create the entire template. You can, however, add or remove VDI providers to/from an existing template.

When done, click **Next** to proceed to the next wizard page.

Additional information

Managing multi-provider template distribution (p. 168)

Advanced

The **Advanced** page has different properties for different types of VDI providers. The differences are described below.

Hypervisor-based VDI providers:

- **Cluster Shared Volume (CSV), Network share:** These two options appear if you are using Hyper-V Failover Cluster. They allow you to select a type of storage where guest VMs will be created. Select a desired option and then click the [...] button next to the edit field. Depending on the option selected, specify a Cluster Shared Volume or network folder. Note that a shared folder must be compatible with SMB 3.0. Please also note that the same credentials used to register Microsoft Hyper-V host as a VDI provider will be used to access the SMB file share for Guest VMs.

Please also read the important note below.

Note: To use this functionality, you need to set SMB constrained delegation (resource-based) using Windows PowerShell. **Important:** Windows Server 2012 forest functional level is required.

On a server running Windows 2012 R2 and above install the Active Directory PowerShell module using Powershell. Note that you don't need the module on a Hyper-V host or SMB file servers.

Run the following cmdlet:

```
Install-WindowsFeature RSAT-AD-PowerShell
```

Delegate SMB delegation on a file server (cluster) for every node of Hyper-V cluster. For example if you are running a four-node Hyper-V cluster and you use a Scale-Out File Server cluster FS-CL01 as virtual machine storage:

```
Enable-SmbDelegation -SmbServer FS-CL01 -SmbClient Hyperv-01
```

```
Enable-SmbDelegation -SmbServer FS-CL01 -SmbClient Hyperv-02
```

```
Enable-SmbDelegation -SmbServer FS-CL01 -SmbClient Hyperv-03
```

```
Enable-SmbDelegation -SmbServer FS-CL01 -SmbClient Hyperv-04
```

Mandatory: verify applied settings (the actual delegations) as follows:

```
Get-SmbDelegation -SmbServer FS-CL01
```

- **Folder:** This option is available if you are using Hyper-V, VMware vCenter, or Nutanix Acropolis. It specifies a folder where guest VMs will be created.
- **Resource pool:** Specifies a VMware resource pool.
- **Physical Host:** Available for VMware vCenter. Specifies a physical host where guest VMs will be created.

Microsoft Azure VDI provider:

- **Resource group:** Select an Azure resource group where the cloned VM will be created. Note that this must be a group to which you granted permissions to the Azure AD app. For details, see [Create a Microsoft Azure AD application](#) (p. 132).
- **Size:** Select a VM size to be used for cloned VMs.
- **OS disk type:** Select a disk type to be used for cloned VMs.

Preparation

Use the **Preparation** page to select and configure an image preparation tool.

Note: When you specify properties on this page, they are remembered in your personal configuration file on the local machine. The next time you decide to create another template, the fields here will be populated automatically using the values you used the last time.

First, select whether you want to use RASprep or Sysprep. The advantages of using RASprep and the differences between the two tools are described below.

RASprep is the Parallels RAS tool for preparing Windows in a VM after cloning it from a base image. RASprep performs the following tasks during the initial startup of each new VM:

- Creates a new computer account in Active Directory for each guest VM.
- Gives the guest VM a new name.
- Joins the guest VM to the Active Directory domain.

Compared to Sysprep, RASprep works much faster because it modifies a lower number of configurable parameters and requires less reboots.

Note: Due to API limitations, RASprep cannot be used on Windows Server 2008 machines.

The following table lists the main differences between RASprep and Sysprep:

Operation	RASprep	Sysprep
Delete local accounts	No	Yes
Generate new SIDs	No	Yes
Unjoin the parent guest VM from the domain	No	Yes
Change computer name	Yes	Yes

Join the new instance to the domain	Yes	Yes
Language, regional settings, date and time customization	No	Yes
Number of reboots	1	2 (seal, mini-setup and domain joining)

After selecting the preparation tool, specify the following options:

- **Computer name:** A name pattern that should be used to assign a computer name. For example, Windows10-RAS-%ID%.
- **Owner name:** Owner name (optional).
- **Organization:** Organization name (optional).
- **Administrative password:** Local Windows administrator password.
- **Join domain:** A domain name for the VM to join.
- **Administrator:** Domain account.
- **Password:** Domain account password.
- **Target OU:** Full DN of an organizational unit. Click the [...] button to browse Active Directory and select an OU.

Optimization

The **Optimization** page allows you to specify settings that will be used to optimize Windows running in guest VMs for best performance in a Parallels RAS environment. You can select Windows components, services, and other options that will be disabled, removed, or optimized to ensure a more efficient, streamlined, and improved delivery of virtual apps and desktops. For the complete description, please see **Optimization** (p. 102).

After reading the **Optimization** section mentioned above, please also note the following VDI specifics:

- Optimization is disabled by default when you create a new template. If you plan to enable it, you should make a backup (create a full clone) of the source guest VM before doing so. You can also create a template with optimization disabled, then create a snapshot, and only then enable optimization. Making a backup is a good idea because once optimization settings are applied, they cannot be rolled back.
- To enable optimization for an existing template, the template must be in maintenance. A template in the "Ready" state has the **Optimization** tab disabled.
- When optimization is enabled or modified for an existing template and the template exits the maintenance mode, the administrator will be asked to recreate existing guest VMs, so that optimization settings are applied to them. Note that guest VMs must be recreated (now or later) to receive optimization settings.
- When optimization is applied to a template, its status changes to **Optimization in progress** (among others). At this stage, you can select the template in the list and click **Tasks > Stop optimization**, which will cancel the operation.

License keys

On the **License Keys** page, specify the license key information that will be used to activate virtual machines created from this template.

First, select the license key management type that you are using in your organization (KMS or MAK). Parallels recommend to use KMS because MAK has limited activations.

Key Management Service (KMS): If you are using KMS, click the **Finish** button to save the template configuration information. Virtual machines that will be created from this template will look for KMS in DNS (at the end of the OS mini-setup and domain joining) and will be activated accordingly.

Note: If you are using KMS activation and RASPrep, the source guest VM must be activated using KMS before you create a template from it. If the guest VM has already been activated using another method (retail key or MAK), you need to convert it to KMS activation. For the information on how to do it, please read the following article from Microsoft: <https://technet.microsoft.com/en-us/library/ff793406.aspx>.

Multiple Activation Keys (MAK): If you are using MAK, do the following:

- 1 Click the **Add** button and type a valid key in the **License key** field.
- 2 In the **Max. guests** field, specify the key limit. The limit should be greater than or equal to the max guests in the template (which you set on the first page of the wizard)
- 3 Click **OK**.

Note: Parallels RAS does not keep the old MAK key in guest VMs if it was updated in the Parallels template properties.

Settings

Specify the following additional settings on this page:

- **Publishing session timeout:** The amount of time a session remains logged in after the user closes a published application. The default timeout is 25 seconds. Note that this only works for applications, but not published desktops (when a user closes a desktop, the session is logged off). This timeout is used to avoid unnecessary logins when a user closes one application and then opens another.
- **Actions:** Allows you to select an action that will be performed on a selected session event. The **After** field, specifies the time period after which the action will be initiated.

Security

On the **Security** page, set the following option:

- **Grant users remote desktop connection permissions:** This option allows you to automatically grant a user permissions for a remote desktop connection. This is achieved by temporarily adding a user to a local group on connect and then removing him/her on logoff or disconnect. To enable this option, select the checkbox and then select one of the available local groups (Remote Desktop Users or Administrators) to which users will be added.

Summary

On the **Summary** page, review the template summary information. You can click the **Back** button to correct some of the information if needed.

Finally, click **Finish** to create the template and close the wizard.

Guest VM naming

This section describes the guest VM naming pattern that you specify on the **Properties** page (p. 153) of the template creation wizard.

Each time a new guest VM is created, a name for it is generated automatically based on the pattern that you specify in the **Guest VM name** field (p. 152). The complete name format is as follows:

```
<prefix>%ID:N:S%<ending>
```

where:

- **<prefix>** is an alphanumeric string that must begin with a letter (not a digit).
- **%ID:N:S%** is a numeric pattern used to automatically generate a unique guest VM ID. See the **Numeric pattern** subsection below.
- **<ending>** is a free-form alphanumeric string.

Numeric pattern

The numeric pattern in the VM name has the following format:

```
%ID:N:S%
```

The elements in the pattern above are:

- **ID** — Must be included as is.
- **N** — The number of digits to use, including leading zeros. Use "0" if you don't want to insert leading zeros.

- **S** — The starting number. This element is optional. If you don't include it, the number will start with 1.

Examples:

- `%ID:3%` — This pattern will generate 3 digit numbers with leading zeros, such as "001", "002", "003"... "998", "999".
- `%ID:3:200%` — This example will generate 3 digit numbers starting from 200, such as "200", "201", "202"... "998", "999".
- `%ID:0%` — This pattern will generate numbers with no leading zeros. It will start at 1 and will (theoretically) go up until the length limit is reached, which is 15 characters for the entire name.
- `VDI-R1-%ID:3:100%` — This is a complete name with an alphanumeric prefix and a numeric pattern. The resulting names will look like the following: "VDI-R1-100", "VDI-R1-101", etc.

When creating a name pattern, follow the rules listed below. If any of these rules are not observed, you will see an error message and will have to correct it:

- The name must start with a letter. A digit is not allowed as the first character.
- The alphanumeric part of a name can contain letters, digits, and a hyphen. No other characters are allowed.
- The total length of the name must not exceed 15 characters.
- The name can include just one numerical pattern (`%ID:N:S%`), which must be placed at the end or in the middle of the name.

The pattern that you specify is also validated against the value of the **Maximum guest VMs** field. If the pattern doesn't cover the maximum number of guest VMs, you will get an error and will have to correct it.

Reusing VM numbers in a name

When you delete a guest VM, the number that was assigned to it becomes unused. The next guest VM that is created will be given this number, so there are no gaps in numbering.

Parallels Test Template Wizard

The **Parallels RAS Test Template Wizard** is used to test the health of a template. The wizard allows you to see that all post-prep activities for a template complete correctly. This includes checking DHCP settings, DNS registration, correct VLAN, joining the AD domain, correct target OU, etc.

To open the wizard, right-click a template in the Parallels RAS console and choose **Test**. The test procedure consists of the following steps:

- 1 The template is switched temporarily to the "Test" mode designed specifically for this purpose. Please note that while the template is in this mode, all other operations are blocked until the test is finished and the template exits the test mode.

- 2 A guest VM is cloned from it to be used for testing. The VM is kept on the server for the duration of the test and will be deleted afterwards.
- 3 A series of tests is then run on the guest VM to test the template from which it was created.
- 4 Once the test is complete, a report is displayed on the screen showing the test results.

When the wizard starts:

- 1 The **Welcome** page opens. Read the info that it contains and click **Next** when ready.
- 2 The next page displays the list of individual tests that will be performed, including:
 - **Check guest VM Agent:** This test tries to communicate with the RAS Guest Agent installed in the VM. If the agent responds, it means that the VM has been created and started successfully.
 - **Check domain membership:** Checks that the computer has joined the AD domain.
 - **Check target OU:** Checks that the RDP connection to the computer is possible with domain credentials.
 - **Launch Parallels Client:** This test launches Parallels Client and establishes a connection with the guest VM.
- 3 While the test is running, the progress indicator is displayed on the screen. If needed, you can cancel the test at any time by clicking the **Cancel** button.
- 4 Once all tests are completed, you will see a page displaying the test results:
 - **Success:** If all tests complete successfully, the temporary guest VM will be marked for deletion and the template will be switched back to the normal operation mode.
 - **Failure:** If one or more tests fail, you will see the corresponding info and will be able to download the log file by clicking the **Download log file** link. You will also have an option to switch the template to maintenance mode, which will prevent creating guest VMs from it until it is fixed.
- 5 Click **Finish** to close the wizard.

Modifying template properties

If you need to change the configuration of an existing template, select it in the **Templates** list and click **Tasks > Properties**. This opens the **Template Properties** dialog, which consists of tabs containing the same properties as the wizard pages described in **Stage 2: Configure the template** (p. 152).

How guest VMs are created from a template

After a template is created, Parallels RAS begins creating guest VMs from it, one virtual machine at a time. The number of VMs created at this time is determined by the **Number of guest VMs deployed on the wizard completion** property (all property names here and later refer to the **Create Template Wizard** described earlier).

The number of VMs available at any time will never go below the number specified in the **Keep available buffer** property. To comply with this rule, a new VM is automatically created when needed. At the same time, the total number of VMs will never exceed the number specified in the **Maximum guest VMs** property.

Please note that creating a new guest VM from a template takes some time, especially when a template is configured to create full clones (linked clones are created much faster). If a guest VM is in the middle of being created, and no other VMs are available, a user (or users) who need it will have to wait until the VM is ready.

If a guest VM encounters a problem during the preparation stage, it will remain on the server in unusable state. You can identify such VMs by the "Failed to create" value in the **Guest VM State** column. Unless a VM like this is repaired or recreated, it will be automatically removed after the time period specified in the **Auto remove guest VMs which failed preparation after** field in Site defaults (Farm > <Site> > **VDI > Desktops > Tasks > Site defaults**). You can view the details of the failure to create a guest VM by right-clicking on the **Guest VM State** column and choosing **Details**. You can also choose **Recreate** or **Delete** from the same menu. For more information on how to recreate a guest VM, please see the **Template maintenance** section (p. 163).

Auto-deletion of guest VMs

A guest VM is automatically deleted when it is not used longer than specified in the **Delete unused guest VMs after** field in template properties.

Manually adding a guest VM

Guest VMs are created from a template automatically. In a situation when one or more additional guest VMs are required, you can add (create) them manually.

To add a guest VM:

- 1 In the RAS Console, navigate to **Farm > <Site> > VDI > Desktops**.
- 2 Click the **[+]** icon at the top of the list.
- 3 In the **Add Guest VMs** dialog that opens, select a template from which to create a new guest VM,
- 4 Specify the number of guest VMs to create. If the number you specify exceeds the "Maximum guest VMs" value set in template properties (taking into account the number of VMs that already exist), you'll see a warning message and will need use a lower number or change the max number in template properties.
- 5 Click **OK** to close the dialog.
- 6 After you click **Apply** in the RAS Console, the new guest VMs will appear in the list on the **Desktop** tab with the **Guest VM state** column saying "Cloning". Once the cloning is complete, the new guest VMs become available to users.

Assigning a template to a pool or group

When you create VDI host pools (VDI templates) or RD Session Host groups (RDSH templates), you can assign a template to a pool or group. This can be done when you create or modify a pool or group, or it can be done from the **Templates** tab.

To assign a template to a pool or group:

- 1 On the **Templates** tab, select a template. Depending on the selected template type (VDI or RDSH) do one of the following:
 - VDI: Click **Tasks > Assign to pool**.
 - RDSH: Click **Tasks > Assign to group**.
- 2 A dialog opens listing existing VDI host pools or RD Session Host groups. Host pools and groups that already have a template assigned are not shown in this list by default. To display them, select the **Show groups/pools with assigned template** option. The template that they are currently using is displayed in the **Template** column.
- 3 Select one or more host pools or groups and click **OK**.

To remove a template from a host pool or group:

- 1 Select a template and click **Tasks > Remove from pool** or **Remove from group**.
- 2 A dialog opens listing all host pools or groups to which this template is assigned.
- 3 Select the host pools or groups to remove the template from and click **OK**.

Note that if a host pool or group has guest VMs created from the template that you are removing, they will be removed as well. A message is displayed where you need to confirm the removal.

Template maintenance

A template can be put into a special mode called "maintenance", which is primarily used to update or install software in the guest operating system. While in this mode, the template becomes unavailable for all normal tasks, including creating new guest VMs, and it becomes possible to start it as a regular virtual machine. Once the virtual machine is running, you can install or update software in the guest OS or perform administrative tasks in the operating system.

Depending on whether a template is configured for full or linked clones, the maintenance mode is used slightly differently, as described below.

Full clones

If your template is configured to create full clones, do the following:

- 1 Select a template and click **Tasks > Maintenance**. The template becomes disabled (grayed out) and all operations on it are suspended. The status of template in the **Status** column changes to "Entering maintenance" and when completed, changes to "Maintenance".
- 2 Using native tools of the hypervisor, start the template as a normal virtual machine.
- 3 Install Windows updates or software as necessary.
- 4 When done, shut down the virtual machine.
- 5 Back in the RAS Console, select the template and click **Tasks > Maintenance** again to exit the maintenance mode. The status of exiting maintenance is also displayed at the bottom of the window.
- 6 At this point, you may see a message asking whether you would like to recreate existing guest VMs. The message is displayed when there's one or more existing guest VMs that were already created from this template. When you update a full clone template, the changes will only affect future clones. For existing clones to have these updates, they must be recreated. You can choose to recreate existing guest VMs now or you can postpone it. Please note that recreating a full clone is a time consuming process. Also, a new app may be installed in a full clone VM or a user profile may be changed while the recreation is in progress, all of which will be lost. To minimize impact on users, it makes sense to schedule a maintenance window during which the clones can be recreated.

Linked clones

Since linked clones share the virtual hard disk with a snapshot of a template, you need to take additional steps compared to full clones.

First, you need to notify guest VM users to save their data and log off. This is necessary for existing guest VMs to include the updates that you will install in the template. Once all users are logged off, do the following:

- 1 Select a template and click **Tasks > Maintenance**. The template becomes disabled (grayed out) and all operations on it are suspended. The status of the operation is displayed at the bottom of the window.
- 2 Using native tools of the hypervisor, start the template as a normal virtual machine.
- 3 Install Windows updates or software as necessary.
- 4 When done, shut down the virtual machine.
- 5 Back in the RAS Console, select the template and click **Tasks > Maintenance** again to exit the maintenance mode. A dialog is displayed asking if you would like to recreate existing guest VMs. If you click **No**, then the dialog is closed and the existing guest VMs are left in their current state, which means that the updates that you installed will not appear in the existing VMs. If you click **Yes**, read on.

- 6** If you click **Yes** in the previous step, existing guest VMs will be examined for active connections. If an active connection exists, another dialog opens asking if you want to proceed:
- If you click **Yes**, all active sessions are forcibly logged off and existing guest VMs (linked clones), together with the corresponding snapshot, are deleted and a new snapshot and VMs are created from the updated template.
 - If you click **No**, the **Template Guest VMs List** dialog opens where you can view the current state of each available guest VM. The dialog gives you control over a guest VM. You can send a message to the user and you can log the user off. Once all active sessions are logged off, click **OK**. The existing guest VMs and the corresponding snapshot are deleted and a new snapshot and VMs are created from the updated template.

When you are done configuring a template, click the **Apply** button on the main RAS Console window to commit the changes to Parallels RAS.

Please note that if you leave the maintenance mode without recreating linking clones, you will have to enter the maintenance mode again to apply the updates.

Updating RAS Guest Agent inside a template

A template must have the latest version of RAS Guest Agent installed in it. The agent is installed when you create a template. When a new version of RAS Guest Agent becomes available, it should be updated. To update the agent, the maintenance mode must be used as described above. To simplify agent updates, Parallels RAS monitors all installed agents and notifies the administrator when an update is available.

When the RAS Console starts, all installed agents are checked and a message is displayed if one or more agents need to be updated. This applies to servers in the RAS infrastructure and the templates. The message will ask if you want to update all agents. If you click **Yes**, you are presented with a dialog listing all servers and templates on which an agent needs to be updated. You can select or un-select a server/template to include it in the bulk update procedure or exclude it. Once you've made your selection, click **OK** to start the update. Follow the onscreen instructions and update the agents.

Full vs. linked clone templates: When you update RAS Guest Agent in a template, you also need to update Agents in guest VMs that were created from this template. This update is done differently for full and linked clone templates. Please read the instructions below for the explanation.

When you update the Agent in a linked clone template, you'll be asked if you want to recreate all guest VMs that were created from this template. You can click **Yes** and they will be automatically recreated to match the template.

When you update the Agent in a full clone template, full clone guest VMs are not automatically recreated. You will be asked if you want to recreate them. If you decide to do so, please note that full clone VMs are complete machines, so recreating them is a time-consuming process. Alternatively, you can update the agent in these VMs by push-installing it from the RAS Console. This can be done by clicking **Tasks > Upgrade all Agents** while on the **VDI > Desktops** tab.

To manually check the RAS Guest Agent status in a template, click **Tasks > Check agent**. If the agent is up to date, a message box is displayed confirming this. If a newer version of RAS Guest Agent is available, you'll see a dialog asking you to update it. Please note that the difference in updating full and linked clone templates (as described above) applies to this scenario as well.

Maintaining RD Session Hosts based on a template

If you need to do a scheduled maintenance of RD Session Hosts that were created from a template, please follow these steps:

- 1 Create a schedule that fits your maintenance window to drain a desired RD Session Host group.
- 2 During maintenance (or right before it) switch the template into maintenance mode. Then apply the necessary changes.
- 3 The schedule disables groups provisioned by the template (while the maintenance window lasts) which leads to removing (unassigning) all guest VMs from them.
- 4 Release the template from maintenance and click **Yes** when asked whether to recreate all clones.
- 5 Enable groups which were disabled in step 3 (above). At this point, the groups will begin receiving guest VMs to comply with **Keep Available Buffer** setting
- 6 From this point forward, groups are provisioned with VMs on demand.

Template status

To verify that a template is functioning as intended, you can examine its status in the main template list in the RAS Console (the **Status** column). When a template is functioning properly, the **Status** column displays "Ready", which means that guest VMs can be created from it as needed. When a template is being created or when it's in maintenance, or when it's being removed, the status will change accordingly.

Note that one of the other columns in the table is **Agent status**, which is the status of the RAS Guest Agent installed in a template. Compared to servers in the RAS infrastructure (Publishing Agent, Gateways, RD Session Hosts, etc), the agent status is not as important in a template as the template status. This is because a template is not a regular virtual machine and is not normally running, so checking the agent status in a stopped VM doesn't make much sense. This is why the agent status for a template in the RAS Console is usually **Not Available**, which is perfectly normal. The only situation when the **Agent status** displays a meaningful value is when the template is in Maintenance and running like a regular VM, in which case the agent is also running and its status can be verified.

The tables below describe what the **Status** and **Agent status** columns will show for various template states or transitions.

Template creation

Status color	Template status	Agent status	Description
Gray	Not available	Not available	When a provider is disabled or VDI Agent is disconnected or the template does not exist.
Gray	Not applied	Not applied	Awaiting admin to click Apply in the RAS Console after the wizard completion.
Orange	Creating	Not available	When using Azure Gallery as a source (no pre-created host available).
Orange	Agent installation	Not available then OK	Deploying the agent to a newly created or available VM to be used as a template.
Orange	Deployment in progress	Optimization pending	When the optimization is waiting to be applied. The admin can stop the optimization at this point.
		Optimization in progress	When the optimization is in progress. The admin can still stop the optimization at this point.
		OK then Not available	Internal procedure when converting a VM to a template. Once the conversion is completed, the template status changes to "Ready".
Red	Creation failed	Not available	A problem has occurred. For example, a quota hit or a resource creation issue in Azure. The admin can retry the action by clicking Tasks > Retry last action.
Red	Agent installation failed	Not available	Possible network issues, file share limitations, or an issue with admin rights. The admin can try Tasks > Retry last action.
Red	Deployment failed	The actual status (OK, Not available, etc.)	A problem has occurred. For example, a quota hit, storage space, or snapshot creation from provider issue. The admin can try Tasks > Retry last action.
		FSLogix not available	FSLogix agent not found.
		FSLogix not updated	FSLogix agent needs updating.

Template in production

Status color	Template status	Agent status	Description
Green	Ready	The actual status (OK, Not available, etc.)	Template is ready.

Green	Cloning	The actual status (OK, Not available, etc.)	A guest VM is being cloned from the template.
Orange	Needs update	Needs update	RAS Guest Agent needs updating.

Template in maintenance

Status color	Template status	Agent status	Description
Orange	Maintenance	The actual status (OK, Not verified, etc.)	Guest VM used as the template is up and running,
		Optimization pending	When the optimization is waiting to be applied. The admin can stop the optimization at this point.
		Optimization in progress	When the optimization is in progress. The admin can still stop the optimization at this point.
		Needs update	Guest VM used as the template is up and running but RAS Guest Agent needs updating.

Template removal

Status color	Template status	Agent status	Description
Gray	Marked for deletion	The actual status while the guest VM used as the template is still running (OK, Not verified, etc.)	The template is in the process of being deleted.

Managing multi-provider template distribution

For the description of the Multi-Provider Template Distribution feature, please see the **Multi-provider template distribution** section (p. 148).

Adding or removing VDI providers from the distribution list

You can add or remove a VDI provider to/from a distribution list at any time using the template **Properties** dialog. To open the dialog, right-click a template on the **VDI > Templates** tab and choose **Properties**.

Template distribution status

After you complete the **Create Parallels Template Wizard** and create a template, or when you add/remove a VDI provider to/from an existing template, you can monitor the template distribution status on the **Templates** tab. The status is displayed in the **Distribution** column and may have the following values:

- **Distributing** — the distribution is in progress (the template is being distributed to target hosts).
- **OK** — the template has been successfully distributed to all specified hosts.
- **Removing / Adding provider** — A VDI provider is being added or removed.
- **Failed to distribute** — indicates that an error has occurred during the distribution operation.

Distribution details

The **Tasks > Distribution details** menu on the **Templates** tab opens a dialog where you can view the current distribution **State** and **Progress** indicators for the VDI providers that use this template.

The **Progress** column displays the same values as the **Distribution** column in the main template list (see above).

The **State** column may display one of the following:

- **Ready** — The VDI provider is ready.
- **Not available** — The VDI provider is not responding.
- **Needs update** — The template distribution operation may need to be performed again. You can click the **Retry** button to retry the template distribution operation for this host.

Leaving maintenance mode

When a template leaves the maintenance mode, a prompt is usually displayed saying that "All guest VMs must be recreated because the template has been modified. Do you want to recreate them now?". If the administrator clicks **Yes** and the template uses multi-provider distribution, Parallels RAS verifies the status of each provider. If a VDI provider is not responding, a message is displayed, asking the administrator to check the provider status. You can bring the provider back online and try recreating the guest VMs again. If this cannot be done at this time, you can recreate the guest VMs later.

Managing template-based guest VMs

Guest VMs and other desktops are managed on the **VDI > Desktops** tab, where you can perform all of the standard desktop management operations from the **Tasks** menu. The operations include Recreate, Delete, Upgrade all Agents, Assign, Unassign, Show sessions, Start, Stop, Restart, Suspend, Reset, and others. The Restart operation (graceful) has a 10 min timeout. If not completed during this time, the Reset operation (forced) will be used.

By default, the **Desktops** tab displays all of the desktop available in the Farm (you may need to scroll the list to see all available desktops). To see just the guest VMs that belong to a specific template, select a template in the **Templates** tab and click **Tasks > Show guest VMs**. This will switch you to the **Desktops** tab where the list will be automatically filtered to include only the VMs that belong to the selected template.

For more information, see **Managing guest VMs** (p. 172).

VDI pool management

Pools offer administrators more flexibility when managing an extensive number of guest VMs, especially when they are implemented in large company infrastructures. The RAS Console provides you with the framework and tools needed to create a complete pool management foundation. To manage pools, in the RAS Console, navigate to **Farm > <Site> > VDI** and then click the **Pools** tab.

Read on to learn how to:

- Add and delete pools (p. 170)
- Add and delete pool members (p. 170)
- Use a wildcard to filter VMs (p. 171)
- Manage guest VMs in a pool (p. 171)

Adding and deleting pools

To add a pool, click the **Tasks** drop-down menu above the **Pools** list and then click **Add** (or click the plus-sign icon). Type a pool name and then click anywhere outside the edit field.

To delete a pool, right-click it and then click **Delete** (or click the minus-sign icon, or **Tasks > Delete**).

Adding and deleting pool members

A VDI pool can contain different types of members. These could be all available guest VMs, specific guest VMs, and guest VMs created from a template.

To add a member to a pool:

- 1 Select a pool in the **Pools** list.
- 2 In the **Tasks** drop-down menu above the **Members** list (the upper right-hand corner of the **Members** area), click **Add** and choose a member type from the following list:
 - **All guest VMs in Host**. All guest VMs that are located on a given VDI provider. After clicking this options, you'll be able to select a VDI provider.

Note: Parallels does not recommend to use this type because there's a possibility that guest VMs with unsupported OS will be added (e.g. Linux, HALB etc). If you need to use this type, please do it carefully or use a wildcard with appropriate guest VMs names (p. 171).

- **Guest VM.** A specific guest VM located in the Farm. After clicking this options, you'll be able to select a guest VM from the list.
 - **Resource pool.** A group of guest VMs that were natively configured in the hypervisor as a pool. Please note that a hypervisor may use a different term for pools (e.g. "resource pools"). After clicking this option, you'll be able to select a resource pool from the list, if any are available.
 - **Template.** Guest VMs that are automatically created from a template. After selecting this option, you'll be able to select a template. For more information about templates, refer to **Templates** (p. 147).
- 3** After you click one of the above menu items, you will be presented with the list of the available hosts, guest VMs, or templates from which you can make your selection.

Note: To avoid issues with overlapping members, a given pool can have members of the same type only. For example, if the first member that you add to a pool is a guest VM, any additional member can be a guest VM, but not a template, Resource pool, or all guest VMs on a specified host. If you want to use members of different types, you must create a separate pool for each member type (i.e. one pool for guest VMs, another pool for templates, etc.). This requirement is enforced in the UI by disabling the member type choices once the first member is added to a pool.

To delete a member from a pool, select the pool, then select a pool member you wish to delete, and then click **Tasks > Delete**.

Using a wildcard to filter VMs

Use the **Wildcard** input field at the bottom of the **Pools** tab to specify a wildcard to indicate which guest VMs should be available for users. If a VM name matches the wildcard, it will be available. If not, the users will not be able to use it. Use the asterisk operator (*) to specify a wildcard (e.g. ABC*, *ABC*).

Managing guest VMs in pools

Guest VMs that belong to a pool (and other guest VMs and desktops) are managed on the **VDI > Desktops** tab, where you can perform all of the standard desktop management operations from the **Tasks** menu. The operations include Recreate, Delete, Upgrade all Agents, Assign, Unassign, Show sessions, Start, Stop, Restart, Suspend, Reset, and others. The Restart operation (graceful) has a 10 min timeout. If not completed during this time, the Reset operation (forced) will be used.

By default, the **Desktops** tab displays all of the desktop available in the Farm (you may need to scroll the list to see all available desktops). To see just the guest VMs that belong to a specific pool, select a pool in the **Pools** tab and click **Tasks > Show guest VMs in Pool**. This will switch you to the **Desktops** tab where the list will be automatically filtered to include only the VMs that belong to the selected pool.

Managing guest VMs

There are two basic types of guest VMs when using Parallels RAS VDI: template-based and non-template based. This topic describes management tasks for both guest VM types, indicating whether a task applies to a particular guest VM type.

Viewing guest VM list

To view the list of non-template based guest VMs, select **Farm > <Site> > VDI > Desktops**. If you have a filter applied to the list, remove it by click the magnifying glass icon. Without the filter, the list shows all desktops available in this RAS Farm, including guest VMs (both template-based and non-template based), guest VMs from a pool (RAS or native), and pool-based Remote PCs. Therefore, the **Desktops** tab is a location where you can view all of your desktops in one place. Here you can perform all of the standard desktop management tasks accessible from the **Tasks** menu, including Recreate, Delete, Assign, Unassign, Start, Stop, Restart, Suspend, Reset, Show sessions, and others. The Restart operation (graceful) has a 10 min timeout. If not completed during this time, the Reset operation (forced) will be used.

To view the list of guest VMs created from a template, select **Farm > <Site> > VDI > Templates**. Select a template and click **Tasks > Show guest VMs**. You will be switched to the **Desktops** tab where the list of desktops will be filtered to include only those that belong to the template. As was mentioned above, you can perform all of the standard desktop management operations on this tab, including power operations, which are described in detail later in this section.

For the list to include only the guest VMs from a particular pool, select a pool in the **Pools** tab and click **Tasks > Show guest VMs in Pool**.

The filter in the **Desktop** tab can also be applied manually by clicking the magnifying glass icon and entering the filter criteria in the fields that appear at the top of the list.

Site defaults

Guest VMs created from a template inherit the template settings. To view the settings, note on which template a guest VM is based and then view properties of that template, specifically the **Settings** and **Security** tabs. For more information, see **Site defaults** (p. 144). Note that you a template can inherit Site default settings or you can specify your own custom settings for it.

Non-template based guest VMs have their own settings, some of which (specifically Settings and Security) are inherited from Site defaults (p. 144). To see settings for a non template-based VM, navigate to **Farm** > <Site> > **VDI** > **Desktops**. A guest VM that doesn't belong to a template is identified by an empty value in the **Template** column. Right-click a template and choose **Properties** (note that template-based guest VMs do not have this menu option).

Checking the RAS Guest Agent status

A guest VM must have the RAS Guest Agent installed and the agent must match the Parallels RAS version. The agent is installed by default when a guest VM is created from a template. If a guest VM was created using the native hypervisor tools, it may not have the agent installed in it. In such a case, the guest VM will be able to serve only the remote desktop. To enable it to server applications or documents, you'll need to install the agent yourself.

Note: Guest VMs based on an RD Session Host template must also have the RAS RD Session Host Agent installed. The functionality described here does not verify if this agent is installed. If needed, you can use **Tasks** > **Check Agent** on the template itself.

To check if the RAS Guest Agent is installed in a guest VM and is up to date:

- 1 Select a guest VM in the list and then click **Tasks** > **Troubleshooting** > **Check agent**.
- 2 The **Guest Agent Information** dialog opens displaying the information about the RAS Guest Agent.
- 3 If the agent is not installed, click the **Install** button and follow the instructions. The agent will be push installed in Windows running inside the guest VM.

Deleting a guest VM

To delete a template-based guest VM, select it and then click the **Tasks** > **Delete**.

Important: You should delete a guest VM only from the RAS Console. You should not try to delete a guest VM using the hypervisor's native client or web interface. If you do, it may delete not only the VM but its parent template as well (which will also invalidate all other guest VMs created as linked clones from this template). The reason for this is some native hypervisor clients treat linked clones as standalone VMs. Parallels RAS treats linked clones as clones, not as standalone VMs.

Managing guest VMs that failed preparation

If a template-based guest VM encounters a problem during the preparation stage, it remains on the server but cannot be used. You can identify such VMs by the "Failed to create" value in the **Guest VM State** column. Unless a VM like this is repaired, it will be automatically removed after the time period specified in the Site defaults (p. 144). To see Site defaults:

- 1 Select **Farm** > <Site> > **VDI** > **Desktop** and then click **Tasks** > **Site defaults**.

- 2 In the dialog that opens, on the **General** tab, view or modify (if needed) the **Auto remove guest VMs which failed preparation after** option. You can set any of the available time periods by selecting it from the drop-down list or you can type a desired value, such as "8 days" or "12 hours".

Recreating a guest VM

If something happens to a template-based guest VM and it becomes unusable, you don't have to delete it and create a new one. Instead, you can recreate it keeping its name and MAC address (to guarantee that VM will get the same IP address from the DHCP server). This way none of the other Site settings, which may rely on a broken guest VM, will be affected. Another reason for recreating a guest VM is to apply changes made to the template (when you exit from maintenance without executing the Recreate command). Please note that keeping the MAC address is supported on ESXi, vCenter, Hyper-v and Hyper-v Failover Cluster only.

Note: If a guest VM was created from an RD Session Host template and was already assigned to an RD Session Host group, it cannot be recreated.

To recreate one or more guest VMs:

- 1 In the Parallels RAS Console, navigate to **Farm \ <Site> \ VDI \ Templates**.
- 2 To recreate all deployed guest VMs, click the **Tasks** drop-down menu and choose **Recreate all guest VMs**.
- 3 To recreate a specific guest VM (or multiple guest VMs), click **Tasks > Show guest VMs**. This will open the **Desktops** tab, which will list guest VMs. Select one or more guest VMs and then click the **Tasks > Recreate**.

When you recreate a guest VM:

- The procedure deletes a VM and creates a new one from the same template.
- The new guest VM retains the same computer name as the one it replaces.
- If a guest VM is running, all unsaved data in its memory will be lost. For this reason, an important data should be saved to an external storage.

Persistent guest VMs

A guest VM is called persistent when it is assigned to a particular user. To make a guest VM persistent, do the following:

- 1 Begin publishing a desktop or a resource from a guest VM.
- 2 When specifying **Virtual Guest Settings** options, select **Persistent**.
- 3 Complete the publishing wizard.
- 4 As a result, the first user who uses a desktop or a resource from this guest VM will become the owner of the VM (i.e. the VM will be assigned to the user).

You can also manually assign a guest VM to a user. To do so:

- 1 Navigate to **Farm** > <Site> > **VDI** > **Desktops**.
- 2 Select a guest VM and click **Tasks** > **Assign**.
- 3 Select a user (enter your network credentials if asked) and click **OK**.
- 4 As a result, the guest VM will be assigned to the selected user.

To view persistent guest VMs, navigate to **Farm** > <Site> > **VDI** > **Desktops**. A persistent guest VM is identified by the "Persistent" value in the **Assignment** column.

To remove persistence from a guest VM, do one of the following:

- Select a guest VM on the **Desktops** tab and then click **Tasks** > **Unassign**.
- Navigate to **Farm** > <Site> > **VDI** > **Desktops** and click **Tasks** > **Site defaults**. In the dialog that opens, use the **Auto remove persistence if guest VM was not used for** option to select the time period after which persistence should be automatically removed. You can also type any desired time period, such as "1 week 3 days".

Using a VDI provider in multiple farms

This topic describes how to use the same VDI provider in multiple RAS Farms simultaneously. To better understand the problem and the solution, consider the following hypothetical example:

- Let's say we have a hypervisor with two available virtual machines.
- We also have two Farms (1 and 2).
- Our intention is to use the first guest VM to host resources in Farm 1 and the other to be a template in Farm 2. Both guest VMs will run simultaneously on the hypervisor, but each one will be available in its respective Farm only.

The problem is, RAS Guest Agent can normally communicate with one RAS VDI Agent, but since each Farm has its own RAS VDI Agent, this will not work out of the box. The solution is to make the RAS Guest Agent running in a guest VM to be aware of only one specific RAS VDI Agent with the ability to change the assignment as needed.

The assignment is done via Windows registry. All guest VMs belonging to VDI pools and guest VM clones created from a template need to have a new String value `2XVDIAgent` specifying the RAS VDI Agent name or address. To add the value:

- 1 Log in to Windows running in the virtual machine, open the registry editor (regedit) and locate the following keys:
 - 32-bit systems: `HLKM\Software\Parallels\GuestAgent`
 - 64-bit systems: `HLKM\Software\WOW6432Node\Parallels\GuestAgent`

- 2 Add a String value named `2XVDIAgent`. The value data should be specified as follows:
 - If a dedicated RAS VDI Agent is used, the value must be set to the FQDN or IP address of the server where the agent is installed.
 - If the built-in RAS VDI Agent is used with manual agent selection, the value must be set to the FQDN or IP address of the RAS Publishing Agent.
 - If the built-in RAS VDI Agent is used and the agent is selected automatically (high availability), the string must contain FQDNs or IP addresses of all RAS Publishing Agents separated by a semicolon (i.e. `<PA1 address>;<PA2 address>;<PA3 address>`).

Note that you can include names or IP addresses of multiple Publishing Agents for the manual agent selection scenario as well (the second bullet in the list above). This way you will not need to change the value every time you switch the preferred Publishing Agent for a VDI provider.

Using computer management tools

You can perform standard computer management tasks on a server right from the RAS Console. The tasks include Remote Desktop Connection, PowerShell, Computer Management, Service Management, Event Viewer, IPconfig, Reboot, and others. To access the **Tools** menu, select a server, click **Tasks > Tools** and choose a desired tool. For the complete description, please see **Computer management tools** (p. 431).

Publishing from a guest VM

This section describes how to publish resources hosted by a guest VM. The publishing functionality described here is accessed from the **Publishing** category in the RAS Console.

Publishing a desktop

Note: When you publish one or more applications from a pool or template, you cannot publish a desktop from it. Likewise, if you publish a desktop from a VDI pool or specific template, you cannot publish individual applications from it.

To publish a remote desktop from a guest VM, follow the below procedure:

- 1 In the RAS Console, select the **Publishing** category and click the **Add** icon below the **Published Resources** tree. This will launch the publishing wizard.
- 2 In the first step of the wizard select **Desktop** and click **Next**.
- 3 On the **Select Desktop Type** page, select **Virtual Desktop** and click **Next**.
- 4 On the **Virtual Desktop** page, enter a desktop name, an optional description, and change the icon if needed.

- 5 In the **Guest VM settings** section, specify from where the desktop should be published. First, you need to select an option in the **Connect to** drop-down list and then specify an additional parameter in the field below it as follows:
 - **Any guest VM.** Use the **from Pool** drop-down list to specify a pool.
 - **Specific Template.** Select a template by expanding the template drop-down list.
- 6 Select the **Persistent** option to mark a guest VM as persistent the first time a user connects to it.
- 7 In the **Desktop Size** section, specify the desktop screen resolution and size.
- 8 In the **Multi-Monitor** field, select the desired option (enable, disable, use client settings).
- 9 If needed, select the **Persistent** option (lower right) to make a guest VM persistent. For more information, see **Persistent guest VMs** (p. 174).
- 10 On the next page, specify the initial status of the resource. Choose from **Enabled** (end users can launch the resource), **Disabled** (the resource will not appear in Parallels Client), **In maintenance** (the resource will appear in Parallels Client but users will not be able to launch it). When a resource is in maintenance and a user tries to launch it, they will see a message. To customize the message, click the **Configure** button. For more info, see **Site defaults (Publishing)** (p. 245).
- 11 Click **Finish** when done.

Publishing an application

Note: When you publish a desktop from a VDI pool or specific template, you cannot publish individual applications from it. Likewise, if you publish one or more applications from a pool or template, you cannot publish a desktop from it.

To publish an application from a guest VM or guest VM clone:

- 1 In the RAS Console, select the **Publishing** category and then click the **Add** icon below the **Published Resources** tree (or right-click inside the **Published Resources** box and click **Add** in the context menu). This will launch the publishing wizard.
- 2 On the **Select Item Type** wizard page, select **Application** and click **Next**.
- 3 On the **Select Server Type** page, select **Virtual Guest** and click **Next**.
- 4 On the **Select Application Type** page, select **Single application** and click **Next**. The **Application** page opens.
- 5 Enter a name and an optional description.
- 6 In the **Run** drop-down menu, specify if the application should run in a normal window, maximized, or minimized.
- 7 In the **Target** field, specify the application that you want to publish. You may click the **[...]** button to browse for it.

- 8 In the **Start in** field, specify (or browse for) the "start in" folder. Use Windows environment variables if you are manually entering the path.
- 9 (Optional) In the **Parameters** input field, specify the parameters to pass to the application when it starts.
- 10 In the **Virtual Guest Settings** section, specify from where the application should be published. First, you need to select an option in the **Connect to** drop-down list and then specify an additional parameter in the field below it, as explained below:
 - **Any guest VM.** Use the **from Pool** drop-down list to specify a pool.
 - **Specific Template.** Select a template by expanding the template drop-down list.
- 11 If needed, select the **Persistent** option to make a guest VM persistent. For more info, see **Persistent guest VMs** (p. 174).
- 12 On the next page, specify the initial status of the resource. Choose from **Enabled** (end users can launch the resource), **Disabled** (the resource will not appear in Parallels Client), **In maintenance** (the resource will appear in Parallels Client but users will not be able to launch it). When a resource is in maintenance and a user tries to launch it, they will see a message. To customize the message, click the **Configure** button. For more info, see **Site defaults (Publishing)** (p. 245).
- 13 When done, click **Finish** to publish the application.

Publishing a web application

A web application is like any other application that you can publish using the standard application publishing functionality. However, to simplify publishing of straight URL links to web applications, a separate publishing item type is available that allows you to accomplish this task with minimal number of steps.

To publish a web application:

- 1 In the RAS Console, select the **Publishing** category and then click the **Add** icon below the **Published Resources** tree (or right-click inside the **Published Resources** box and click **Add** in the context menu). This will launch the publishing wizard.
- 2 On the **Select Item Type** page, select **Web application** and click **Next**.
- 3 On the **Select Server Type** page, select **Virtual Guest** and click **Next**.
- 4 On the **Virtual Desktop Web Application** page, specify the web application name, description, window state, and the URL. Select the **Force to use Internet Explorer** option if needed. To browse for a specific application icon, click **Change Icon**.
- 5 Use the **Virtual Guest Settings** section to specify from where the application should be published.

The options are:

- **Any guest VM.** Publish the application from any guest VM in the selected pool. Select this option and then select a pool in the **from Pool** drop-down list.

- **Specific template.** Publish the application from a specific template. Select this option and then select a template in the **Template** drop-down list.
- 6 Select the **Persistent** option to make a guest VM persistent. For more info, see **Persistent guest VMs** (p. 174).
 - 7 On the next page, specify the initial status of the resource. Choose from **Enabled** (end users can launch the resource), **Disabled** (the resource will not appear in Parallels Client), **In maintenance** (the resource will appear in Parallels Client but users will not be able to launch it). When a resource is in maintenance and a user tries to launch it, they will see a message. To customize the message, click the **Configure** button. For more info, see **Site defaults (Publishing)** (p. 245).
 - 8 When done, click **Finish** to publish the application.

Publishing a network folder

You can publish a filesystem folder via UNC path to open in Windows explorer. To minimize the number of configuration steps, a special publishing item is available that allows you to publish a network folder from a guest VM.

To publish a network folder:

- 1 In the RAS Console, select the **Publishing** category and then click the **Add** icon below the **Published Resources** tree (or right-click inside the **Published Resources** box and click **Add** in the context menu). This will launch the publishing wizard.
- 2 On the **Select Item Type** page, select **Folder on the file system** and click **Next**.
- 3 On the **Select Server Type** page, select **Virtual Guest** and click **Next**.
- 4 On the **Virtual Desktop UNC Folder** page, specify the usual application properties.
- 5 In the **UNC path** field, enter the UNC path of the folder you wish to publish. Click the **[...]** button to browse for a folder (it may take some time for the **Browse for Folder** dialog to open).
- 6 In the **Virtual Guest Settings** section, specify from where the virtual desktop should be published. First, you need to select an option in the **Connect to** drop-down list and then specify an additional parameter in the field below it, as explained below:
 - **Any guest VM.** Use the from Pool drop-down list to specify a pool.
 - **Specific Template.** Select a template by expanding the template drop-down list.
- 7 Select the **Persistent** option to make a guest VM persistent. For more info, see **Persistent guest VMs** (p. 174).
- 8 On the next page, specify the initial status of the resource. Choose from **Enabled** (end users can launch the resource), **Disabled** (the resource will not appear in Parallels Client), **In maintenance** (the resource will appear in Parallels Client but users will not be able to launch it). When a resource is in maintenance and a user tries to launch it, they will see a message. To customize the message, click the **Configure** button. For more info, see **Site defaults (Publishing)** (p. 245).

9 Click **Finish** to publish the folder and close the wizard.

When published, the network folder will appear in the **Publishing > Published resources** list, just like any other application. To view its properties, select it and then click the **Virtual Desktop Application** tab:

- The **Target** property will always be set to `PublishedExplorer.exe`. This binary is created automatically (via agents pushing) and is simply a copy of the standard `explorer.exe` executable.
- The **Parameters** property specifies the network folder that we want to publish. The folder path can be in any format that the `explorer.exe` can handle.

Publishing a document

To publish a document from a guest VM or guest VM clone:

- 1 In the RAS Console, select the **Publishing** category and then click the **Add** icon below the **Published Resources** tree (or right-click inside the **Published Resources** box and click **Add** in the context menu). This will launch the publishing wizard.
- 2 On the **Select Item Type** wizard page, select **Document** and click **Next**.
- 3 Select **Virtual Guest** and click **Next**.
- 4 Specify the content type of the document you want to publish. You can select the content type from the predefined list or specify a custom content type in the **Custom content types** input field.
- 5 Click **Next**.
- 6 On the **Virtual Desktop Application** page, enter a name, optional description, Window state, and an icon if needed.
- 7 Use the [...] button next to the **Target** input field to browse for the document. All other fields will be automatically populated. To edit any of the auto populated fields, highlight them and enter the required details.
- 8 (Optional) In the **Parameters** input field, specify the parameters to pass to the application when it starts.

Note: Use the **Server(s)** drop down list to specify different document settings for a specific server in case the document is configured differently on that particular server. The settings will be saved for each server you select individually.

9 In the **Virtual Guest Settings** section, specify from where the virtual desktop should be published.

First, you need to select an option in the **Connect to** drop-down list and then specify an additional parameter in the field below it, as explained below:

- **Any guest VM.** Use the **from Pool** drop-down list to specify a pool.
- **Specific Template.** Select a template by expanding the template drop-down list.

- 10 Select the **Persistent** option to make a guest VM persistent. For more info, see **Persistent guest VMs** (p. 174).
- 11 On the next page, specify the initial status of the resource. Choose from **Enabled** (end users can launch the resource), **Disabled** (the resource will not appear in Parallels Client), **In maintenance** (the resource will appear in Parallels Client but users will not be able to launch it). When a resource is in maintenance and a user tries to launch it, they will see a message. To customize the message, click the **Configure** button. For more info, see **Site defaults (Publishing)** (p. 245).
- 12 Click **Finish** to publish the document.

Viewing VDI provider summary

In addition to the VDI provider editor described earlier in this chapter, you can also see summary about the available VDI providers. To do so:

- 1 In the RAS Console, select the **Farm** category and then select the **Site** node in the middle pane.
- 2 The available servers are displayed in the **VDI** section in the right pane.
- 3 To go to the VDI provider editor, right-click a server and choose **Show in the Editor**.

For additional info, see **Sites in the RAS Console** (p. 45).

Managing VDI sessions

Please see **Session Management** (p. 227).

Remote PC pools

Remote PC pools is a Parallels RAS feature that allows you to create pools of standalone (preferably domain-joined) PCs and optionally assign them to specific users. The Remote PC pools functionality is integrated into RAS VDI to take advantage of the infrastructure that already handles host pools.

Remote PC pools vs. Remote PCs

Remote PCs are standalone machines (physical or virtual) that can be used to host published resources in Parallels RAS. Remote PCs are managed in the Parallels RAS Console in the **Farm** > <Site> > **Remote PCs** section. The **Remote PCs** chapter (p. 216) describes this functionality in detail. Remote PC pools described here are handled separately and differently from standalone Remote PCs. They are managed in the **Farm** > <Site> > **VDI** section of the RAS Console.

In this section:

- Adding a VDI provider (p. 182)
- Configuring the VDI provider (p. 183)
- Adding Remote PCs to a pool (p. 184)
- Managing Remote PCs in a pool (p. 185)
- Persistent Remote PCs (p. 186)
- RAS Guest Agent installation options (p. 186)
- Publishing from a pool-based Remote PC (p. 187)

Adding a VDI provider

To set up a Remote PC pool in the RAS Console, you first need to add a VDI provider of type **Remote PC**. This is a special type that exists specifically for the purpose of creating and managing Remote PC pools. It is not a real VDI provider, so it doesn't need a hypervisor installed. It simply uses the existing VDI functionality to create and manage computer pools. Note that when you add a VDI provider of this type, you can manage it like any other real VDI provider with some limitations, such as you cannot create templates and use some other strictly VDI (hypervisor)-specific functionality.

To add a VDI provider of type **Remote PC**:

- 1** Navigate to **Farm** > <Site> > **VDI**.
- 2** On the **Providers** tab, click **Tasks** > **Add**.
- 3** On the first page of the **Add VDI Provider** wizard, select **Remote PC** in the **Type** drop-down list.
- 4** In the **Address** field, specify FQDN or IP address of a server that will manage Remote PC pools. This must be a server with RAS VDI Agent installed. You can use the RAS Publishing Agent server, since it has the RAS VDI Agent built in, but it can be any other server running a dedicated RAS VDI Agent.
- 5** Enter the account name in the UPN format (e.g. `administrator@domain.local`). This must be a domain user account with administrative rights on the server specified in the step 4 above. Using a local Windows account is also possible with some limitations and only when using the static PC assignment (see **Configuring a VDI provider** (p. 183)). Using a domain account is recommended.
- 6** Enter the account password and an optional description.
- 7** Click **Next** and then click **Finish**.

The wizard closes and the server **Properties** dialog opens where you need to configure the new VDI provider. You can configure it now or you can configure it later by right-clicking the host on the **Providers** tab and choosing **Properties**. The configuration steps are described in detail in the section that follows this one.

Configuring a VDI provider

If you don't have the VDI provider properties dialog open, right click a VDI provider that you created and choose **Properties**.

On the **Properties** tab of the dialog, select **Enable host in site** to enable the host.

In the **VDI subtype** drop-down list, select how Remote PCs will be assigned to this VDI provider, so they can be later added to a Remote PC pool. The following options are available:

- **Static:** Using this approach, Remote PCs are assigned to the VDI provider by entering their FQDN or IP address (one by one) or by importing a list from a CSV file.
- **Dynamic:** This approach assigns PCs using the information from Active Directory. All you have to do is specify an organizational unit (OU) containing computer accounts to be assigned to the host.

Please note that once you choose one of the options above and assign PCs to the host later, you cannot switch from static to dynamic or vice versa later.

Depending on which **VDI subtype** you select on the **Properties** tab (see above), the **Remote PCs** tab will look differently. The subsections below describe **Static** and **Dynamic** scenarios respectively.

Static (VDI subtype)

Using this approach, Remote PCs are assigned to the VDI provider manually one by one or are imported from a CSV file. Parallels RAS 18 (and newer) supports a maximum of 1000 Remote PCs per VDI provider.

Note: To be manageable, Remote PCs should be domain-joined. In case of static assignment described here, it is possible to add non-domain joined PCs, but you will have to create the same local user account on each and everyone of them. Using a domain account and domain-joined PCs is recommended.

To add a PC, select the **Remote PCs** tab and do one of the following:

- Click **Tasks > Add** and type FQDN or IP address of a PC you want to add. You can click the **[...]** button to search for it. Next, enter the MAC address of the computer you are adding. Note that both fields are mandatory.
- Click **Tasks > Import from CSV file** and then select a CSV file containing the list of computers. The CSV file must contain two columns: (1) FQDN or IP address; (2) MAC address. Once again, both columns are mandatory and must contain valid values.

Dynamic (VDI subtype)

To use dynamic assignment, you need to specify an organizational unit (OU) containing computer accounts to be assigned to the host. To do so, select the **Remote PCs** tab and specify the organizational unit in the **Target OU** field. You can click the [...] button to browse Active Directory. Note that a maximum of 1000 Remote PCs can be returned in a single AD/OU search result.

Note: When using dynamic assignment, Remote PCs must be domain-joined. You cannot manage such PCs using a local Windows user account.

When you use the dynamic assignment, you have an option to install RAS Guest Agent on every PC by adding a Group Policy to the organizational unit with a script to deploy RAS Guest Agent. The following is an example of such script:

```
msiexec /i RASInstaller-<version & build>.msi ADDLOCAL=F_GuestAgent /qn+ /norestart
```

Other agent installation options are described in **RAS Guest Agent installation options** (p. 186).

Adding Remote PCs to a pool

Note: To be managed in a Remote PC pool, a Remote PC must have RAS Guest Agent installed. For more information, see the **RAS Guest Agent installation options** (p. 186).

Once you assigned PCs to a VDI provider, you can add them to a Remote PC pool as follows:

- 1 In **Farm** > <Site> > **VDI**, select the **Pools** tab.
- 2 Add a new pool by clicking **Tasks** > **Add** in the **Pools** pane.
- 3 Select the pool that you've created and then in the **Members** pane, click **Tasks** > **Add** and choose one of the following:
 - **All Guest VMs in Host:** Adds all Remote PCs assigned to the VDI provider. When you click this option, a dialog opens allowing you to select a VDI provider. Select the host and click **OK**.
 - **Guest VM:** Adds an individual Remote PC. In the dialog that opens, select a desired Remote PC and click **OK**. Another dialog may open asking you to upgrade RAS Guest Agent on a Remote PC (the agent is required for a PC to be managed in a pool). Click **OK** to upgrade (or install) the agent. You can also upgrade the RAS Guest Agent on one or more PCs at another time as described in **RAS Guest Agent installation options** (p. 186).

Once you add one or more Remote PCs to a pool, they will appear in the **Pool management** tab and in the **Desktops** tab.

Tip: If you need to disable the pool for maintenance, you can do so by clearing the checkbox in front of the pool name.

Managing Remote PCs in a pool

Management of pool-based Remote PCs includes assigning a PC to a specific user, upgrading the RAS Guest Agent, viewing and modifying PC properties, performing some standard administrative tasks, and some others.

To manage Remote PCs in a pool:

- 1 In **Farm** > <Site> > **VDI**, select the **Desktops** tab.
- 2 Note that the list on this tab includes all managed desktops, including guest VMs and pool-based Remote PCs. You can order the list by the **Pool** column to see Remote PCs assigned to a particular pool.
- 3 Select a Remote PC, click the **Tasks** drop-down menu and choose one of the options described below. Note that not all options available in the **Tasks** menu are applicable to Remote PCs. The list below describes only the options that you can use with pool-based Remote PCs.

The **Tasks** menu options that apply to Remote PCs are:

- **Upgrade all Agents.** Upgrade RAS Guest Agent in all Remote PCs (and guest VMs) in the list.
- **Assign.** Assign a Remote PC to a specific user (make a PC persistent). Click the menu option and specify a user.
- **Unassign.** Remove the user assignment (persistence) from a Remote PC.
- **Show sessions.** Switches the view to the **Sessions** tab and displays the session information.
- **Tools.** Allows to perform a set of standard operations, such as establishing a remote desktop connection, pinging, rebooting/shutting down a Remote PC, and others. For the description of power operations, please see **Performing power operations** below.
- **Troubleshooting.** Check and install/upgrade the RAS Guest Agent in a Remote PC.
- **Reset properties.** Resets Remote PC properties to their default values. See **Properties** below.
- **Properties.** Opens a dialog where you can view and modify Remote PC settings. The **General** tab allows you to temporarily disable the Remote PC in a pool (use the **Do not use this guest VM** option). This is specifically useful when you need to perform maintenance tasks on a PC. You can also view and modify the Remote PC display name, computer name, and the port number on which it communicates with the VDI provider. For the description of **Settings** and **Security** tabs, see **Site defaults** (p. 144).

Performing power operations

For remote power operations to work, WMI must be enabled in Windows running in a VM and TCP ports 30004 and 30005 must be open. At the time of this writing, this is not automated in Parallels RAS but will be in future versions.

To perform power operations on a guest VM (start, stop, restart, suspend, reset), open the **VDI > Desktops** tab, select a guest VM, then click **Tasks** and choose an operation that you want to perform (for start and stop operations, you can click the corresponding icons at the top). The restart operation (graceful) has a 10 min timeout. If not completed during this time, the reset operation (forced) will be used.

Please note if you are using **Nutanix Acropolis**, the suspend operation is not available (the **Suspend** icon is disabled). The reason for this is Nutanix Acropolis does not support the suspend operation on its virtual machines.

Persistent Remote PCs

A persistent Remote PC is a PC assigned to a particular user. Once a PC is assigned, no other user can connect to it.

There are two ways to make a Remote PC persistent:

- When you publish a resource (application, desktop, etc.) from a pool-based Remote PC using the publishing wizard, you can select the **Persistent** option in the **Virtual Guest Settings** section. This way, a Remote PC in a pool will be assigned to the first user that opens the published resource. For more info, see **Publishing from a pool-based Remote PC** (p. 187).
- You can also assign a Remote PC to a user manually. To do so, navigate to **Farm > <Site> > VDI**, select the **Desktops** tab, then select a Remote PC in the list and click **Tasks > Assign**. In the dialog that opens, specify the target user.

To remove persistence from a Remote PC, select it in the **Desktops** tab and click **Tasks > Unassign**.

RAS Guest Agent installation options

To be managed in a Remote PC pool, a Remote PC must have RAS Guest Agent installed. This can be done using one of the following options:

- When you add an individual Remote PC to a pool, you'll be asked to upgrade the agent. Follow the onscreen instructions and install or upgrade it.
- When you add all Remote PCs in a host to a pool at once, you can add them first and then use the **Tasks > Upgrade all Agents** menu option in the **Desktops** tab.
- When you assign Remote PCs to a VDI provider via Active Directory, you can have a Group Policy in the OU with a script to deploy the agent. See **Configuring the VDI provider > Dynamic (VDI subtype)** (p. 183).
- To install or upgrade the agent on an individual Remote PC, select it in the **Desktops** tab and click **Tasks > Troubleshooting > Check agent** option. In the dialog that opens, click **Install**.
- Finally, you can install RAS Guest Agent manually by running the Parallels RAS installer on a Remote PC and selecting to install the RAS Guest Agent component.

Publishing from a pool-based Remote PC

To publish a resource (application, desktop, etc.) from a pool-based Remote PC, do the following:

- 1** In the RAS Console, select the **Publishing** category.
- 2** Click the **Add** icon (lower left-hand corner).
- 3** Select a resource type and click **Next**.
- 4** On the **Server Type** page, select **Virtual Guest**. The **Remote PC** option there is for standalone Remote PCs. To publish from a pool-based Remote PC, the **Virtual Guest** is the correct type.
- 5** Advance to the page where you specify the resource name and properties (e.g. **Virtual Desktop Application** page).
- 6** In the **Properties** section, select **Any Guest VM** in the **Connect to** field and then select a Remote PC pool in the **from Pool** field.
- 7** To make a Remote PC persistent, select the **Persistent** option. This will assign a Remote PC in a pool to the first user who will use this published resource.
- 8** Populate the rest of the fields as usual and click **Finish** to publish the resource.

CHAPTER 9

Azure Virtual Desktop

Azure Virtual Desktop (formerly known as Windows Virtual Desktop) is a desktop and app virtualization service running on Microsoft Azure, providing access to RD Session Hosts and VDI, including the new offering of Windows 10 and Windows 11 Enterprise multi-session hosts. Parallels RAS 18 provides the ability to integrate, configure, maintain, support and access Azure Virtual Desktop workloads on top of the existing technical capabilities of Parallels RAS.

In This Chapter

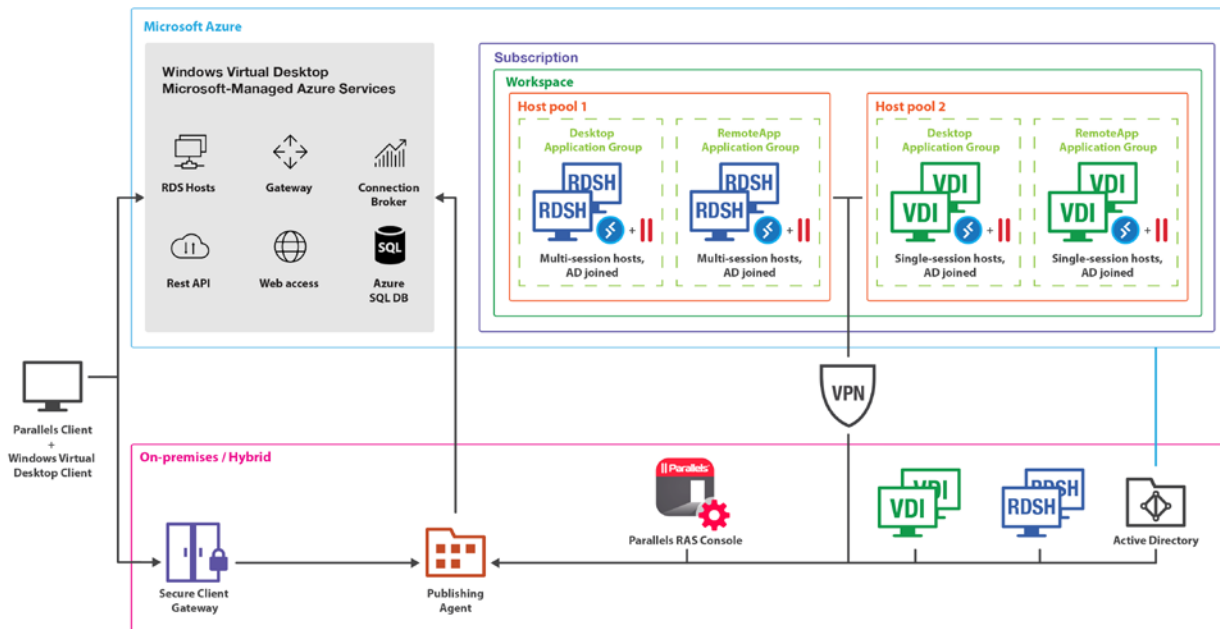
Introduction	188
Prerequisites	190
Deploy Azure Virtual Desktop	192
Manage Azure Virtual Desktop.....	197
Site defaults (Azure Virtual Desktop)	208
Verify the deployment.....	213
Publish resources	213
Using Parallels Client with Azure Virtual Desktop.....	214

Introduction

The diagram below illustrates a hybrid deployment of Parallels RAS and Azure Virtual Desktop with the following characteristics:

- Workload hosts are available both on-premises through standard Parallels RAS deployment and on Microsoft Azure through the service.
- Azure Virtual Desktop objects such as workspaces, host pools, desktop and RemoteApp groups are created and configured from the Parallels RAS Console.
- Azure Virtual Desktop hosts (multi-session or single-session) contain both Azure Virtual Desktop Agent and RAS Agent for management and configuration purposes.

- Parallels Client for Windows is connecting to both Parallels RAS Secure Client Gateway and Azure Virtual Desktop service providing resource availability to end-users from a single interface.



Extended values and capabilities

- Simplify and enhance Azure Virtual Desktop deployment and management.
- Unify administration and user experience – single pane of glass – Parallels Clients and Parallels RAS Console.
- Extend reach with flexibility to use hybrid and multi-cloud deployments.
- Automate and streamline administrative routines, provisioning, and management of Azure Virtual Desktop workloads.
- Built in Auto-scale capability on Microsoft Azure and/or on-premises.
- Management of users, sessions, and processes.
- Utilize RAS Universal Printing and Scanning.
- Utilize AI based session prelaunch for ultra-fast logons.
- Accelerated file redirection with the use of enable drive cache redirection.
- Integrated automatic image optimizations and FSLogix Profile Containers.
- Client management.
- Security policies for clients.
- Leverage RAS Reporting and Monitoring from the RAS Console.

Prerequisites

The below highlights the prerequisites required to use Azure Virtual Desktop and configuration in Parallels RAS environment.

Microsoft Azure subscription

You need a Microsoft Azure subscription, including:

- An Azure Tenant ID.
- An Azure subscription with sufficient credit.

Azure Virtual Desktop user license entitlement

Customers with the licenses listed below are entitled to use Azure Virtual Desktop at no additional charge apart from Azure compute, storage, and network usage billing.

To run Windows 10 and Windows 11 with Azure Virtual Desktop you need to have one of the following per user license:

- Microsoft 365 F3, E3, E5, A3, A5, Student Use Benefits or Business Premium
- Windows 10 Enterprise E3, E5
- Windows 10 Education A3, A5
- Windows 10 VDA per user

To run Windows Server 2012 R2, 2016, 2019, 2022:

- Per user or per device Remote Desktop Services (RDS) Client Access License (CAL) with active Software Assurance (SA).

For further information, please refer to Microsoft licensing requirements at <https://docs.microsoft.com/en-us/azure/virtual-desktop/overview>.

Permissions and Azure resource providers

The below highlights permissions and resource providers to be registered in the subscription:

- Permissions to enable resource providers on your Azure subscription and create virtual machines (VMs).
- The necessary Microsoft Azure resource providers (**Azure Portal > Subscription > Resource Providers**) must be enabled, including **Microsoft.ResourceGraph**, **Microsoft.Resources**, **Microsoft.Compute**, **Microsoft.Network**, **Microsoft.DesktopVirtualization**.

Microsoft Azure AD application

For a detailed information about creating an Azure AD application, please see **Create a Microsoft Azure AD application** (p. 132).

Once an Azure AD Application is created, give the application the following API permissions in the Microsoft Azure Portal (**Azure Active Directory > App Registrations > API permissions > Add a permissions > Microsoft.Graph > Application permission**):

- **Group > Group.Read.All**
- **User > User.Read.All**

Note: Please make sure that when adding Graph API permissions, User and Group, the permission type is "Application" not "Delegated".

Give the application read and write access to resources:

- The Azure AD application that you created must have read and write access to Azure resources as described in **Create a Microsoft Azure AD application** (p. 132). Look for "Give the application read and write access to resources".

Roles and permissions for the application should include:

- "User Access Administrator" role for the application from **Subscription > Access Control (IAM)**.
- "Contributor" role at the Resource group level from **Resource group > Access Control (IAM)**.

If a resource group creation is required, also assign contributor role at the subscription level **Subscription > Access Control (IAM)**.

Note: If you would like to also view/read resources outside the resource group make sure that the application is also given read access at the subscription level.

Active Directory

- A Server Active Directory environment or Azure Active Directory Domain Services (AADDs). See <https://azure.microsoft.com/services/active-directory-ds/>.
- Azure AD Connect — AD must be in sync with your Azure AD, so users can be associated between the two.
- The user must be sourced from the same Active Directory that's connected to Azure AD. Azure Virtual Desktop does not support B2B or MSA accounts.
- The user configured in the Parallels client with access to Azure Virtual Desktop resources must exist in the Active Directory domain the session host it is joined to.

Other

- Azure Virtual Network providing session hosts connection to the domain.
- Session hosts must be domain-joined to Active Directory.
- (optional) Site-to-site VPN or ExpressRoute is required if hybrid Parallels RAS deployment is used.
- (optional) Shared network location to be used for FSLogix Profile Containers which may run on Azure Files or Azure NetApp Files.

Note: At the time of writing, Windows 7 is not supported by Parallels RAS as an Azure Virtual Desktop session host.

Additional notes

Please also note the following Provider and Azure Application requirements for different RAS Farm and RAS Site scenarios:

- Same RAS Farm, same RAS Site: The same Farm, Site, and Application ID is possible to be used for both VDI and Azure Virtual Desktop. Build the guest VM list with Azure Virtual Desktop tags for Azure Virtual Desktop provider and guest VMs with VDI tags (or no tags) for Azure VDI provider.
- Same RAS Farm, same RAS Site: It is recommended to use different Azure Applications for multiple providers of the same type. For example, multiple Azure Virtual Desktop or multiple VDI providers but not mixed.
- Same RAS Farm, different RAS Sites or different RAS Farms: The point above applies. Alternatively, different RAS Farms or Sites can (and must in this case) reside in different virtual networks with no communication to common set of VMs.

Important: It is recommended that Parallels RAS managed Azure Virtual Desktop objects are managed through the Parallels RAS console. Configuration changes outside Parallels RAS console may result in a broken state of Azure Virtual Desktop objects. For such cases, Parallels RAS provides the ability to repair objects. For example, auto created friendly names and associated tags for workspaces and host pools can also be viewed from the Microsoft Azure portal, however they are not to be edited as these are used to ensure proper functionality.

Deploy Azure Virtual Desktop

Azure Virtual Desktop deployment in Parallels RAS is done by completing a series of wizards, including:

- 1** Enable Feature and Add Azure Virtual Desktop Provider.
- 2** Add an Azure Virtual Desktop Workspace.

- 3 Add an Azure Virtual Desktop host pool and then add standalone or template-based hosts to the host pool.
- 4 Publish Azure Virtual Desktop resources.

You can run all from the **Start** category as part of a single deployment procedure. Read on to learn how to do it.

Enable Azure Virtual Desktop and add a provider

Azure Virtual Desktop integration must first be enabled in the RAS Farm. This can be done from two places in the RAS Console:

- Using the **Deploy Azure Virtual Desktop** wizard in the **Start** category.
- By going to **Farm > Site > Settings** and selecting the **Features** tab.

The instructions below are for enabling and deploying Azure Virtual Desktop from the **Start** category. The **Features** tab in **Site > Settings** has the same elements as the **Enable Feature** page described below.

Note: If you haven't enabled Azure Virtual Desktop in the RAS Farm yet, the wizard pages will open in the order described below. If Azure Virtual Desktop is already enabled (e.g. you ran the wizard before or enabled Azure Virtual Desktop from Site settings), the first two pages will be skipped and the first page you'll see is **Add Azure Virtual Desktop Provider** where you need to enter the provider information.

To begin the deployment:

- 1 In the Parallels RAS Console, select the **Start** category and launch the **Deploy Azure Virtual Desktop** wizard.
- 2 **System and user requirements:** On the first page, read system and user requirements. Click a link at the bottom of the page to read a Parallels KB article for more information. Click **Next**.
- 3 **Enable Feature:** This page allows you to enable Azure Virtual Desktop in the RAS Farm. First, select where to store the Azure Virtual Desktop agent and bootloader from the following options:
 - **Publishing Agent:** Store on the RAS Publishing Agent server.
 - **Network share:** Specify or select a network share.
- 4 Click the **Download agent and bootloader** button. Wait for the download to complete and examine the **Status** section, which should indicate "Available" and display the version number. When a new version of the Azure Virtual Desktop agent is available, "Needs update" is shown so new servers deployed from Parallels RAS will use the updated version.
- 5 The **Client feature set** selection specifies which client features will be available when you open a published resource in Parallels Client. Select from the following options:

- **Standard:** Standard feature set. This is identical to opening and running a published resource using the Microsoft Windows Desktop client, also known as Remote Desktop (MSRDC) client, which is the client used to access apps and desktops from Azure Virtual Desktop.
 - **Advanced:** This option also uses the Windows Desktop client but adds advanced Parallels Client features, such as RAS Universal Printing and Scanning, URL redirection, drag and drop, and others.
 - **Advanced with fallback:** This option first tries to open a published resource using the Advanced feature set. If Advanced doesn't work for any reason, it will try to open the resource using the Standard option.
- 6** This completes the task of enabling Azure Virtual Desktop in the RAS Farm. Click **Next** to advance to the next page.
- 7 Add Azure Virtual Desktop Provider:** On this page you need to specify your Microsoft Azure Tenant ID, Subscription ID, Application ID, and a secret key. This is similar to setting up Microsoft Azure as a VDI provider in Parallels RAS. For the explanation of how to specify these properties, please see **Add Microsoft Azure as a VDI provider** (p. 135). Please note that under subscription details, URIs/URLs may be edited during creation of a provider. The **Feed URL** setting, which by default is <https://rdweb.wvd.microsoft.com/api/arm/feeddiscovery>, may also be edited once an Azure Virtual Desktop provider is created.
- 8** Click **Next**, review the summary and click **Finish**. Note that changing Microsoft Azure app permissions after a provider is created may require a restart of the Parallels RAS redundancy service so new permissions are loaded and used.

This completes the first wizard in the series. On the last page, the **Launch Azure Virtual Desktop Workspace wizard** option is enabled by default. This will automatically open the next wizard where you can add an Azure Virtual Desktop workspace.

Add a workspace

A workspace is a logical grouping of application groups in Azure Virtual Desktop. Each Azure Virtual Desktop application group must be associated with a workspace for users to see the remote apps and desktops published to them.

To add a workspace:

- 1** Select whether you want to create a new workspace or select an existing one:
 - To select an existing workspace, click the [...] button next to the **Name** field.
 - To create a new workspace, type a name and optional description. Select an existing or create a new resource group. Specify a location. Note that the location that you select here will be used for all Azure Virtual Desktop objects, including workspaces, host pools, and application groups.

The **Friendly name** value will be generated automatically. It is used by Parallels RAS and is required for Azure Virtual Desktop to work properly in Parallels RAS.

- 2 Click **Next**, review the summary and click **Finish**.

The last page of the wizard has the **Launch Azure Virtual Desktop host pool wizard** option selected by default. This will automatically open the **Add Azure Virtual Desktop Host Pool** wizard when you click **Finish**.

Add a host pool

A host pool is a collection of one or more identical virtual machines (VMs) within an Azure Virtual Desktop environment. Each host pool contains an application group that users can access.

To add a host pool:

- 1 Select whether you want to create new or select an existing host pool:
 - To select an existing host pool, click the [...] button next to the **Name** field. Note that adding an existing host pool will not add any existing application groups, so this, along with application publishing, will need to be reconfigured in the RAS console.
 - To create a new host pool, type a name and optional description and select a resource group.
- 2 Click **Next**.
- 3 On the **Configuration** page, specify the following:
 - **Host pool type:** Select from Pooled (multi-session hosts) or Personal (single-session hosts).
 - **Publishing type:** Select from Application or Desktop depending on what you want to use the pool for.
 - **Load balancer:** Select a load balancer type. Breadth-first load balancing allows you to evenly distribute user sessions across the session hosts in a host pool. Depth-first load balancing allows you to saturate a session host with user sessions in a host pool. Once the first session host reaches its session limit threshold, the load balancer directs any new user connections to the next session host in the host pool until it reaches its limit, and so on.
 - **Limit number of sessions on host:** For a pooled (multi-session) pool type, specify the maximum allowed number of sessions on a host.
 - **Power on host on-demand:** Specify whether a powered down host should be powered on when a user tries to connect to it. Note that this applies only if all session hosts in the host pool are powered off.
 - **Service updates validation:** Select the **Validation environment** option if you want to make this host pool a validation environment for Microsoft service updates.
- 4 Click **Next**.
- 5 On the **Provisioning** page, select whether this host pool will contain template-based or standalone hosts:

- **Template:** Hosts will be created dynamically from a template. You will need to create or select an existing template in the next step or later. Choosing **Template** as the provisioning type ensures a homogeneous host pool, which is recommended to provide consistent user experience across the host pool.
- **Standalone:** Select one or more hosts that already exist. You'll be able to do it in the next step or you can do it later. Prior to adding hosts to host pools, ensure that hosts are domain joined and have network access to the domain environment. Note that the Standalone provisioning is considered "unmanaged" as it lacks some of the functionality, such as Autoscaling.

6 Click **Next**.

7 Depending on the selection made on the **Provisioning** page (above), do one of the following:

- **Standalone:** Select one or more hosts from the list to be included in the host pool (you can also add hosts to the pool later).
- **Template:** On the **Template** page, select a template from the list or click **Create new** to create a new template. If you select a template from the list, specify additional options if needed. Select **Enable autoscale** (multi-session hosts) and specify the initial number of hosts to add to the pool. You can also overwrite the virtual machine size, which is normally set on the template level. To do so, select the **Overwrite the size specified in template properties** option and select a desired size from the **Size** drop-down list. The size that you specify here will be used by this host pool only. Other host pools using the same template will be unaffected. Note that if a VM is later taken out of such a host pool due to autoscale settings, the VM will retain the last known size and may join another host pool with the new size specified. Also note that available sizes may depend on the location, size, and power state of the host pool members and the template. For instructions on how to create a new template, see **Manage templates** (p. 202).

8 Click **Next**.

9 On the **Assignment** page, specify users or groups to be assigned to the application group in the host pool. This is necessary for users to have access to published apps or desktops. Click **Tasks > Add** and specify a user or group. An application group of type **Desktop** or **RemoteApp** (whichever is appropriate) will be created and associated with the host pool automatically on wizard completion.

10 Follow the onscreen instructions and complete the wizard.

Note: In the case of using the Advanced Client Feature Set, RemoteApp groups are not required for publishing applications since the Desktop App Group with Parallels seamless technology will be used to provide application publishing from configured desktop app groups.

Next step

Verify the Azure Virtual Desktop deployment (p. 213)

Manage Azure Virtual Desktop

Read this section to learn how manage Azure Virtual Desktop components in Parallels RAS.

In his section:

- Manage providers (p. 197)
- Manage workspaces (p. 198)
- Manage host pools (p. 198)
- Manage templates (p. 202)
- Manage hosts (p. 205)
- Manage sessions (p. 207)
- Using scheduler (p. 207)

Manage providers

An Azure Virtual Desktop provider in Parallels RAS is a collection of IDs and other properties that give you access to Azure resources. Properties include Tenant ID and Subscription ID, among others. Normally, an organization is given one Tenant ID by Microsoft, but there could be multiple subscription IDs owned by the same organization. For each Tenant ID and subscription ID combination, a provider must be configured in Parallels RAS.

To manage Azure Virtual Desktop providers, navigate to **Farm > Site > Azure Virtual Desktop** and select the **Providers** tab.

To add a new provider, click **Tasks > Add** to open the **Azure Virtual Desktop Provider** wizard. For the information on how to complete the wizard, please see **Enable Azure Virtual Desktop and add a provider** (p. 193). Look for the **Add Azure Virtual Desktop Provider** step, which describes the wizard.

To view and modify some of the existing provider properties, right-click a provider in the list and choose **Properties**.

Other provider management tasks can be accessed from the **Tasks** menu, including:

- Checking provider status: **Tasks > Troubleshooting > Check status**.
- Configure and manage logging: **Tasks > Troubleshooting > Logging**.

Choosing the Active Directory Domain Services type

Parallels RAS 18.3 supports virtual machines joined to Windows Server Active Directory Domain Services and Azure Active Directory Domain Services. By default, Parallels RAS is configured to work with Windows Server Active Directory Domain Services, but you can change this if needed.

To choose the type of Active Directory Domain Services for a provider:

- 1 Right-click a provider in the list and choose **Properties**.
- 2 In the provider properties window, select the **Credentials** tab.
- 3 In the **Active Directory Domain Services type** drop-down list, select the type of Active Directory Domain Services:
 - If your users are created with Windows Server Active Directory and your virtual machines are joined to Windows Server Active Directory Domain Services, select **Windows Server AD DS** (selected by default).
 - If your users are created with Windows Server Active Directory and your virtual machines are joined to Azure Active Directory Domain Services, select **Azure AD DS**.
 - If your users are created with Azure Active Directory and your virtual machines are joined to Azure Active Directory Domain Services, select **Azure AD DS**.

Manage workspaces

A workspace is a logical grouping of application groups in Azure Virtual Desktop. Each Azure Virtual Desktop application group must be associated with a workspace for users to see published remote apps and desktops.

To manage Azure Virtual Desktop workspaces, navigate to **Farm > Site > Azure Virtual Desktop** and select the **Workspaces** tab.

To add a workspace:

- 1 Click **Tasks > Add** to open the **Add Azure Virtual Desktop Workspace** wizard.
- 2 Select a provider at the top of the wizard page (if you have more than one). You can also create a new provider right from this page. If you wish to do so, click the **New provider** button to open another wizard. For details, see **Manage providers** (p. 197).
- 3 After selecting (or creating) a provider, complete the workspace wizard as described in **Add an Azure Virtual Desktop workspace** (p. 194).

To view properties of an existing workspace, right-click it and choose **Properties**. You can enable or disable the workspace and modify the workspace description; other properties are read-only. Note that if you disable the workspace, all associated objects, including host pools and published resources will also be disabled.

Manage host pools

A host pool is a collection of one or more identical virtual machines (VMs) within an Azure Virtual Desktop environment. Each host pool contains an application group that users can access.

Host pools can be configured a number of different ways depending on the intended purpose. The following table describes different options that you can choose when creating a host pool.

Option	Description
Personal vs. pooled	<ul style="list-style-type: none"> • Personal host pools contain single session hosts, each of which is assigned to a single user. The assignment is persisted even after the user logs off or the host is powered off. You can unassign the host from a user and assign it to a different user if needed. • Pooled host pools contain multi-user session hosts (RD Session Hosts or multi-session Windows 10 machines), which are not assigned to any particular user. Each host in a pool can serve multiple users (multi-session).
Application vs. desktop	<p>A host pool can only publish applications or desktops, but not both at the same time. When you create a host pool, you choose a publishing type from Desktop or Application. An application group of the appropriate type (Desktop or RemoteApp) for the host pool is created automatically. Note that you cannot change the publishing type later. If you decide that you want to change it, you'll have to delete the existing host pool and create a new one.</p>
Template vs. standalone	<p>When you create a host pool, you need to select from Template or Standalone. A host pool can contain hosts that already exist (Standalone) or it can use a template which in turn could be based on an existing guest VM or chosen to be created on the fly from images in Azure Marketplace or in your Shared Image Gallery.</p> <ul style="list-style-type: none"> • Template: Hosts can be created from the template by the administrator manually or they can be created automatically when there's a demand. Automatic host creation (called Autoscale in Parallels RAS) can be turned on or off in the host pool properties. • Standalone: Hosts are added and removed to/from a host pool by the administrator. Hosts (virtual machines) must already exist in Azure and must be domain joined.

To manage Azure Virtual Desktop host pools, navigate to **Farm > Site > Azure Virtual Desktop** and select the **Host pools** tab.

To add a new host pool:

- 1 Click **Tasks > Add** to open the **Add Azure Virtual Desktop Host Pool** wizard.
- 2 Select a provider and workspace at the top of the page. You can also create a new provider and/or workspace right from this page by clicking a corresponding **Create new** button, which will open a wizard.
- 3 After selecting (or creating) a provider and workspace, complete the wizard as described in **Add an Azure Virtual Desktop host pool** (p. 195).

To view and modify host pool properties, right-click it and choose **Properties**. In the dialog that opens, select tabs and view or modify host pool properties as described below.

General

On the **General** tab, you can enable or disable the host pool. Note that if you disable it, all hosts and published resources will also be disabled.

You can also modify the host pool description and view general host pool properties.

Note the application group settings at the bottom of the tab page. Depending on the host pool publishing type, you'll see the following:

- **Desktop:** If the host pool is configured to publish desktops, you will see just the **Desktop application group** section.
- **Application:** If the host pool is configured to publish applications, you will see two groups: **Desktop** and **RemoteApp**. The **RemoteApp** is the standard Azure group type for publishing remote applications. The **Desktop** group is added by Azure automatically, but is not used in Parallels RAS if the host pool publishing type is **Application**.

The **Friendly name** property value is generated automatically by Parallels RAS and is required for host pools to function properly.

Configuration

On the **Configuration** tab, examine the host pool configuration properties. You should be familiar with them from when you created a host pool.

You can modify the following properties on this page:

- **Load balancer**
- **Limit number of sessions on host**
- **Power on host on-demand**
- **Validation environment**

For the explanation of configuration properties, see **Add an Azure Virtual Desktop host pool** (p. 195).

Autoscale

This tab is shown only for host pools with **Template** as a provisioning type. Here you can select a template if one has not been specified for the host pool yet. You can also create a new template by clicking the **Create new** button, which will open a wizard. If you don't have any templates, the only selection available is **None**, which means that there's no template to create hosts from. If that's the case, you need to create a template first and then select it here. See also **Manage templates** (p. 202).

The **Autoscale settings** section contains settings that determine how hosts (virtual machines) are created from the specified template. These settings work the same as Autoscale settings for RD Session Host groups. The only difference is, in Azure Virtual Desktop we deal with hosts and host pools, while in RD Session Host groups we talk about servers and groups, otherwise the settings work in a similar manner. For details, please see the **Autoscale** subsection in **Grouping and cloning RD Session Hosts** (p. 108).

The **Specifications section** allows you to overwrite the virtual machine size specified in the settings of the template used by the given host pool. Select the **Overwrite the size specified in template properties** option and select a desired size from the drop-down list. The selected size will only be used by this host pool. Other host pools using the same template will be unaffected. Note that available sizes may depend on the location, size, and power state of the host pool members and the template. Also note that overwriting the size requires a host reboot.

Hosts

The **Hosts** tab lists hosts from this host pool. You can examine the status of a host and other properties by looking at the values in the table.

The **Registration** column should indicate "OK" if a host is operating normally. To verify the agent status, right-click a host and choose **Check agent**. If you see a message that "Agent did not reply", click **Install** to try and install the agent. If everything goes well, the agent will be updated and the **Registration** column should say "OK".

To add a new host to the pool:

- 1 Click **Tasks > Add**.
- 2 Depending on the host pool provisioning type, do one of the following:
 - If the host pool provisioning is configured as **Standalone**, select one or more hosts from the list. You can also select the **Show hosts in existing host pools not managed by RAS** option to show hosts that exist in other host pools on Azure.
 - If the host pool provisioning is configured as **Template**, you cannot manually add hosts here. Instead, use the **Hosts** tab in the main Azure Virtual Desktop view. (p. 205)
- 3 Click **OK**.

Assignment

The **Assignment** tab displays Active Directory users and groups assigned to Azure Active Directory objects. In order for users to see published desktops and applications, they must be assigned to the application group available in the host pool.

To create a new assignment:

- 1 Click **Tasks > Add**.
- 2 In the **Select User or Group** dialog, specify a user or group and click **OK**.
- 3 Follow the onscreen instructions and complete the assignment. Note that additional filtering in the **Publishing** category may be used to control Azure Virtual Desktop resource availability in Parallels Client. For details, see **Publish resources** (p. 213).

User profile

By default, this tab inherits its settings from Site defaults. If you wish to specify custom settings, clear the **Inherit default settings** option. For the information about configuring user profile, please see **Site defaults (Azure Virtual Desktop)** (p. 208).

Optimization

The **Optimization** tab allows you to specify settings that will be used to optimize session hosts for best performance in a Parallels RAS environment. You can select Windows components, services, and other options that will be disabled, removed, or optimized to ensure a more efficient, streamlined, and improved delivery of virtual apps and desktops. By default, this tab inherits its settings from Site defaults. If you wish to specify custom settings, clear the **Inherit default settings** option. For the information about configuring optimization options, please see **Site defaults (Azure Virtual Desktop)** (p. 208).

Host pool settings

This tab allows you to configure settings such as sessions timeouts, client URL/Mail redirection, drag and drop and others. By default, this tab inherits its settings from Site defaults. If you wish to specify custom settings, clear the **Inherit default settings** option. For the information about configuring host pool settings, please see **Site defaults (Azure Virtual Desktop)** (p. 208).

RDP printer

The **RDP Printer** tab allows you to configure the renaming format of redirected printers. By default, this tab inherits its settings from Site defaults. If you wish to specify custom settings, clear the **Inherit default settings** option. For the information about configuring user profile, please see **Site defaults (Azure Virtual Desktop)** (p. 208).

Manage templates

An Azure Virtual Desktop template is a virtual machine from which other virtual machines are created as clones of the original VM and added to a host pool as session hosts.

To manage Azure Virtual Desktop templates, navigate to **Farm > Site > Azure Virtual Desktop** and select the **Templates** tab.

In this section:

- Create a template (p. 203)
- Manage existing templates (p. 204)

Create a template

To create a template:

- 1 On the **Templates** tab, click **Tasks > Add**. This opens the **Create Parallels Template Wizard**.
- 2 On the first page, select an Azure Virtual Desktop provider (if you have more than one).
- 3 Select a template type from the following:
 - **Multi-session:** Multiple concurrent user sessions are allowed on a single host running a Windows Server operating system, Windows 10 or Windows 11 Enterprise multi-session.
 - **Single-session:** A single user session is allowed on a single session host.
- 4 On the **Template Source** page, select a source from the following:
 - **Custom host:** Displays a list of existing virtual machines.
 - **Azure Gallery:** Allows you to select an image and create a new virtual machine from it. Depending on the template type, multi-session or single-session, commonly used marketplace images such as Windows 10 Enterprise multi-session, are predefined to be easily chosen and created as a template. Select a location and specify the local administrator username and password. The **Browse all images** button opens a dialog where you can choose any other image from the Marketplace or Shared Image Gallery. When choosing an image from the Shared Image Gallery, select from a list of publishers, SKUs, offers, and other options.
- 5 On the **Properties** page, type a template name and specify the following options:
 - **Number of hosts to deploy on wizard completion:** How many hosts to create as soon as the template is created.
 - **Maximum number of hosts:** The maximum allowed number of hosts that can be created from the template.
 - **Create an availability set:** If selected, hosts will be deployed from the template in an availability set. Note that the maximum number of hosts that can be deployed in an availability set is 200 (this is an Azure limitation). If you require more than 200 hosts, clear this option and specify your own value in the **Maximum number of hosts** field.
 - **Host prefix:** Specify a host prefix that will be used when naming hosts.
- 6 On the **Settings** tab, specify the following:
 - **Keep available buffer:** The minimum number of hosts to always keep unassigned and session free for this template. As soon as the number of free and unassigned hosts drops below the setting value, it forces the template to create another host. The template uses its own settings for hosts creation including the initial power state.
 - **Host power state after the preparation:** Select whether a host should be powered on or off after it is ready.
 - **Delete unused hosts after:** Select a time period after which an unused host should be deleted.

7 On the **Hosts** page, select virtual machine properties from the predefined Azure values according to your needs:

- First, specify an Azure resource group.
- Select a virtual machine size.

Note: The virtual machine size can be overridden in the settings of the host pool using a given template. This gives you the ability to specify a different VM size at the host pool level.

- Select a disk type.
- Select a virtual network and subnet.

Note: In case using Accelerated networking for the Template, make sure you select the appropriate host size for session hosts that support accelerated networking.

- 1** On the **Optimization** page, configure optimization settings. These settings are inherited from Site defaults but custom settings can be specified if needed. For details, please see **Site defaults (Azure Virtual Desktop)** (p. 208).
- 2** On the **Preparation** page, select an image preparation and specify the required options. This is similar to how an image is prepared for a RAS VDI template. There are some minor differences, but the configuration procedure is essentially the same. For details, please see **Preparation** (p. 156).
- 3** On the **Summary** page, review the settings and click **Finish** to create the template.

Manage existing templates

Modify a template

To modify an existing template, right-click it and choose **Properties**. Some properties cannot be modified, while many can. For the description of individual properties and settings, please refer to instructions in **Create a template** (p. 203).

To delete a template, select it in the list and click **Tasks > Delete**. Note that at the time of this writing, there's a known issue that if a template is deleted in the RAS Console, the template and associated hosts may not be completely removed from Microsoft Azure. To make sure that all such objects are removed, it is recommended to do it from the Azure portal.

Assign a template to a host pool

When you create a host pool and set its provisioning type as Template, you need to assign an existing template to it. This can be done when you create or modify a host pool, or you can assign a template to a host pool on the **Templates** tab.

To assign a template to a host pool:

- 1 Select a template and click **Tasks > Assign to host pool**.
- 2 The **Assign to Host Pool** dialog opens listing available host pools. The list contains all suitable host pools depending on the location and other relevant properties. Host pools that already have a template assigned are not included in the list. To display those host pools as well, select the **Show host pools with assigned template** option. The currently assigned template is displayed in the **Template** column.
- 3 Select one or more host pools and click **OK**.

To remove a template from a host pool:

- 1 Select a template and click **Tasks > Remove from host pool**.
- 2 The **Remove from Host Pool** dialog opens listing all host pools using the selected template.
- 3 Select one or more host pools to remove the template from and click **OK**.
- 4 Note that if a host pool has guest VMs, they will be removed. You will see a message and need to confirm the removal.
- 5 If one or more host pools are being locked by another administrator, you will also see a message and will have to repeat the operation later when a pool is unlocked.

Manage hosts

To manage Azure Virtual Desktop hosts, navigate to **Farm > Site > Azure Virtual Desktop** and select the **Hosts** tab.

The list displays hosts from all available host pools. You can apply a filter to the table to see hosts from a particular pool or using other criteria. To apply a filter, click the magnifying glass icon and specify the filter in a column (or columns) of interest.

Tasks that you can perform on a host are accessible from the **Tasks** menu and include the following:

- **Add:** Add a host to one of the available host pools. See the **Add a host** subsection below.
- **Assign:** This option is enabled for hosts from a Personal host pool. It allows you to assign the selected host to a user. If a host is already assigned to another user, you'll be asked if you want to change the assignment. Select an Azure Active Directory user when asked. The assignment is done in Azure, so the host status will change to "Assigning" for the duration of the operation.
- **Unassign:** Removes the user assignment from the selected host, see **Assign** above. This menu option is enabled for hosts that are currently assigned to a user. The status of the host changes to "Unassigning" for the duration of the operation.
- **Search:** Allows you to search for a host in the list by applying a filter.
- **Show sessions:** Switches to the **Sessions** tab with a filter applied to show the selected host sessions.

- **Show published resources:** Displays a list of resources published from the selected host.
- **Control:** Control options, including enable or disable logons on the selected host, cancel a pending reboot (originated by scheduler), cancel a disabled state (originated by scheduler). See **Using scheduler** (p. 207) for details.
- **Start, Stop, Reset, Restart:** Power operations that can be performed on the selected host. The Restart operation (graceful) has a 10 min timeout. If not completed during this time, the Reset operation (forced) will be used.
- **Upgrade all Agents:** Upgrades agents on every host in the list (if necessary).
- **Stop optimization:** When an optimization is applied to a host, it can be canceled in the beginning stages. For more information, see **Optimization** (p. 102).
- **Tools:** Standard RAS tools, including Remote Desktop, computer management, service management, event viewer, Powershell, and others. For the complete description, please see **Computer management tools** (p. 431).
- **Troubleshooting:** Allows you to check the agent status and update it if necessary. Also allows you to manage logging.
- **Details:** Shows details when a host was not created due to a failure. The option opens a dialog describing the reason of failure and some additional information.
- **Recreate:** Recreates a host.
- **Delete:** Deletes a host from the list and from the host pool to which it belongs. The host (virtual machine) itself is kept or deleted depending on the host pool provisioning type. A host created from a template will be completely removed. A standalone host is not deleted, which means that the virtual machine stays intact.
- **Refresh:** Refreshes the list.

Add a host

You can add a host to a host pool from the **Hosts** tab. To do so:

- 1 Click **Tasks > Add**.
- 2 In the **Add Hosts** dialog, select a target host pool. Depending on the provisioning type configured for the selected host pool, do the following:
 - **Standalone:** Select one or more hosts from the list. You can also select the option at the bottom to show hosts in other existing host pools on Azure that are not managed by Parallels RAS.
 - **Template:** Specify the number of hosts to add to the pool from the template.
- 3 Click **OK**.

Manage sessions

To view and manage Azure Virtual Desktop sessions, navigate to **Farm > Site > Azure Virtual Desktop** and select the **Sessions** tab. Sessions from all hosts in all host pools are displayed in the list.

For a detailed information about managing sessions, please see **Session Management** (p. 227).

Using scheduler

The **Scheduler** tab allows you to schedule a maintenance time window for one or more hosts or host pools by taking it offline at the specified time for a specified period of time. It also allows you to schedule a reboot of one or more hosts or entire host pools.

Note: When the scheduled event is triggered, affected hosts are disabled in Parallels RAS and their status is displayed as "Disabled (scheduler)" or "Pending reboot (scheduler)". You can cancel these states by right-clicking a host on the **Hosts** tab and choosing **Control > Cancel disabled state (scheduler)** or **Control > Cancel pending reboot (scheduler)**.

To configure a maintenance time window:

- 1 Click **Tasks > Add > Disable Host** or **Disable Host pool**. Both scenarios are configured similarly.
- 2 On the **General** tab, select the **Enable Schedule** option.
- 3 Specify a name for this schedule and an optional description.
- 4 Select a host or a host pool (depending on which action you selected) in the **Available** list and click **Add**. The host (or host pool) will appear in the **Target** list.
- 5 Select the **Trigger** tab and specify start date and time, duration, and recurrence settings for this event. To make this a one-time event, select "Never" in the **Recur** field.
- 6 Select the **Options** tab. Here you can configure a message that will be sent to users before the host goes offline. Click **Tasks > Add** and specify the message title, body, and the time period when it should be sent.
 - The **On disable** option on the **Options** tab allows you to specify what should happen to current sessions when a scheduled task triggers. Select a desired option from the drop-down list.
- 7 Click **OK** to save the schedule.

To configure a scheduled reboot:

- 1 Click **Tasks > Add > Reboot Host** or **Reboot Host pool**.
- 2 Specify the schedule properties the same way as described above for the "Disable host" task. The only difference is the **Options** tab, which contains the following additional options for the "reboot" tasks:

- **Enable Drain Mode** and **Force server reboot after**: The two options work together. If you enable the drain mode, the following will happen. When the task triggers, new connections to a host are refused but active connections will continue to run. The server will be rebooted when all active users close their sessions or when **Force server reboot after** time is reached, whichever comes first. For active users not to lose their work, create a message that will advise them to save their work and log off.
- **Enforce schedule for currently inactive hosts**: This option is enabled when the **Enable Drain Mode** option is selected. If selected, inactive hosts will be rebooted too.

3 Click **OK** to save the schedule.

For a more detailed explanation of how to use the scheduler, please see **RD Session Hosts > Using scheduler** (p. 112). The topic describes how to use scheduler with RD Sessions Hosts, but the functionality is the same.

Site defaults (Azure Virtual Desktop)

When you configure Azure Virtual Desktop components and objects in the RAS Console, some of the properties are inherited from Site defaults. When you see the **Inherit default settings** option in a dialog or tab page, it means that settings can be either inherited from Site defaults or custom values can be specified for a given object.

To view and configure Site defaults for Azure Virtual Desktop, navigate to **Farm > Site**, click the **Tasks menu** and choose one of the following:

- **AVD multi-session hosts**: Opens a dialog to configure Azure Virtual Desktop Site defaults for multi-session hosts (p. 208).
- **AVD single-session hosts**: Opens a dialog to configure Azure Virtual Desktop Site defaults for single-session hosts (p. 210).

Each dialog is described below.

Site defaults for multi-session hosts

Host pool settings

Configure the following settings:

- **Disconnect active session after**: Specifies the amount of time each session remains connected in the background after the user has closed the published application. This option is used to avoid unnecessary reconnections with the server.
- **Logoff disconnected session after**: This setting allows you to control how long it takes for a session to be logged off after it is marked as "disconnected".

- **Session readiness timeout:** The maximum amount of time it should require to establish a session. If the specified timeout is reached, and the session is still not ready, the user will see an error message and will have to try to log in again.
- **Allow URL/Mail redirection:** When a user tries to open a URL or an HTML Mailto link in a remote application, the link can be redirected to the client computer and open in a local default application (a web browser or email client) instead of an application on the remote host. This option allows you to enable or disable the redirection. Click the **Configure** button to choose from the following options:
 - a **Replace registered application** — this option uses an alternative method of redirecting a link. It replaces the default web browser and mail client with "dummy" apps on the remote server side. By doing so, it can intercept an attempt to open a link and redirect it to the client computer.
 - b **Support Windows Shell URL namespace objects** — the Shell URL namespace objects support means that Parallels RAS can intercept actions in published applications that use Shell namespace API to open links, which is a standard behavior in most applications. The ability to disable the support for Shell URL namespace objects is for compatibility with older versions of Parallels RAS.
- **Enable drag and drop:** Allows you to set how the drag and drop functionality works in Parallels Clients. Click **Configure** and choose from "Disabled" (no drag and drop functionality), "Sever to client only" (drag and drop to a local application only), "Client to server only" (drag and drop to a remote application only), "Bidirectional" (drag and drop in both directions).

Note: At the time of this writing, the drag and drop functionality is only supported on Parallels Client for Windows and Parallels Client for Mac.

- **Allow 2XRemoteExec to send command to the client:** Select this option to allow a process running on the server to instruct the client to deploy an application on the client side.
- **Use RDP Shortpath:** Enables RDP Shortpath. RDP Shortpath establishes the direct connectivity between a remote desktop client and session host. Direct connectivity reduces the dependency on the Azure Virtual Desktop gateways, improves the connection reliability, and increases the bandwidth available for each user session. A session host requires a restart for this setting to become effective. To allow access to the RDP Shortpath listener across network security boundaries, Azure Network Security Group must be configured to allow inbound UDP port 3390. VPN or ExpressRoute is required or each session host should have public IP address.
- **Enable applications monitoring:** Enable or disable monitoring of applications on the server. Disabling application monitoring stops the WMI monitoring to reduce CPU usage on the server and network usage while transferring the information to RAS Publishing Agent. If the option is enabled, the collected information will appear in a corresponding RAS report. If the option is disabled, the information from this server will be absent from a report.
- **Allow file transfer command (HTML5 and Chrome clients):** Enables file transfer in a remote session. To enable file transfer, select this option and click the **Configure** button. For more information, see **Configuring remote file transfer** (p. 408).

- **Enable drive redirection cache:** Improves user experience by making file browsing and navigation on redirected drives much faster. For details, see **Drive redirection cache explanation** (p. 107)
- **Use RemoteApp if available:** Enable this option to allow use of remote apps for shell-related issues when an app is not displayed correctly. This feature is supported on the Parallels Client for Windows only.

User profile

The **User profile** tab allows you to configure the user profile functionality. You can select from **Do not manage by RAS** (user profiles will not be managed) or **FSlogix**. Microsoft FSLogix Profile Container allows to maintain user context in non-persistent environments, minimize sign-in times and provides native profile experience eliminating compatibility issues. For complete instructions, please see **User profile** (p. 96).

Optimization

The **Optimization** tab allows you to specify settings that will be used to optimize a session host for best performance in a Parallels RAS environment. You can select Windows components, services, and other options that will be disabled, removed, or optimized to ensure a more efficient, streamlined, and improved delivery of virtual apps and desktops. For the complete description, please see **Optimization** (p. 102).

RDP printer

The **RDP Printer** tab allows you to configure the renaming format of redirected printers. The format may vary depending on which version and language of the server you are using.

Set your RDP Printer Name Format specifically for the configured server by choosing any of the below options from the RDP Printer Name Format drop down menu:

- Printername (from Computername) in Session no
- Session no. (computername from) Printername
- Printername (redirected Session no)

The **Remove session number from printer name** does just that, it removes the session number from the name, so it's not visible.

Site defaults for single-session hosts

Host pool settings

Configure the following settings:

- **Disconnect active session after:** Specifies the amount of time each session remains connected in the background after the user has closed the published application. This option is used to avoid unnecessary reconnections with the server.
 - **Logoff disconnected session after:** This setting allows you to control how long it takes for a session to be logged off after it is marked as "disconnected".
 - **Session readiness timeout:** The maximum amount of time it should require to establish a session. If the specified timeout is reached, and the session is still not ready, the user will see an error message and will have to try to log in again.
 - **Allow URL/Mail redirection:** When a user tries to open a URL or an HTML Mailto link in a remote application, the link can be redirected to the client computer and open in a local default application (a web browser or email client) instead of an application on the remote host. This option allows you to enable or disable the redirection. Click the **Configure** button to choose from the following options:
 - a Replace registered application** — this option uses an alternative method of redirecting a link. It replaces the default web browser and mail client with "dummy" apps on the remote server side. By doing so, it can intercept an attempt to open a link and redirect it to the client computer.
 - b Support Windows Shell URL namespace objects** — the Shell URL namespace objects support means that Parallels RAS can intercept actions in published applications that use Shell namespace API to open links, which is a standard behavior in most applications. The ability to disable the support for Shell URL namespace objects is for compatibility with older versions of Parallels RAS.
 - **Enable drag and drop:** Allows you to set how the drag and drop functionality works in Parallels Clients. Click **Configure** and choose from "Disabled" (no drag and drop functionality), "Server to client only" (drag and drop to a local application only), "Client to server only" (drag and drop to a remote application only), "Bidirectional" (drag and drop in both directions).
- Note:** At the time of this writing, the drag and drop functionality is only supported on Parallels Client for Windows and Parallels Client for Mac.
- **Allow 2XRemoteExec to send command to the client:** Select this option to allow a process running on the server to instruct the client to deploy an application on the client side.

- **Use RDP Shortpath:** Enables RDP Shortpath. RDP Shortpath establishes the direct connectivity between a remote desktop client and session host. Direct connectivity reduces the dependency on the Azure Virtual Desktop gateways, improves the connection reliability, and increases the bandwidth available for each user session. A session host requires a restart for this setting to become effective. To allow access to the RDP Shortpath listener across network security boundaries, Azure Network Security Group must be configured to allow inbound UDP port 3390. VPN or ExpressRoute is required or each session host should have public IP address.
- **Enable applications monitoring:** Enable or disable monitoring of applications on the server. Disabling application monitoring stops the WMI monitoring to reduce CPU usage on the server and network usage while transferring the information to RAS Publishing Agent. If the option is enabled, the collected information will appear in a corresponding RAS report. If the option is disabled, the information from this server will be absent from a report.
- **Allow file transfer command (HTML5 and Chrome clients):** Enables file transfer in a remote session. To enable file transfer, select this option and click the **Configure** button. For more information, see **Configuring remote file transfer** (p. 408).
- **Enable drive redirection cache:** Improves user experience by making file browsing and navigation on redirected drives much faster. For details, see **Drive redirection cache explanation** (p. 107)
- **Use RemoteApp if available:** Enable this option to allow use of remote apps for shell-related issues when an app is not displayed correctly. This feature is supported on the Parallels Client for Windows only.

User profile

The **User profile** tab allows you to configure the user profile functionality. You can select from **Do not manage by RAS** (user profiles will not be managed) or **FSlogix**. Microsoft FSLogix Profile Container allows to maintain user context in non-persistent environments, minimize sign-in times and provides native profile experience eliminating compatibility issues. For complete instructions, please see **User profile** (p. 96).

Optimization

The **Optimization** tab allows you to specify settings that will be used to optimize a session host for best performance in a Parallels RAS environment. You can select Windows components, services, and other options that will be disabled, removed, or optimized to ensure a more efficient, streamlined, and improved delivery of virtual apps and desktops. For the complete description, please see **Optimization** (p. 102).

RDP printer

The **RDP Printer** tab allows you to configure the renaming format of redirected printers. The format may vary depending on which version and language of the server you are using.

Set your RDP Printer Name Format specifically for the configured server by choosing any of the below options from the RDP Printer Name Format drop down menu:

- Printername (from Computername) in Session no
- Session no. (computername from) Printername
- Printername (redirected Session no)

The **Remove session number from printer name** does just that, it removes the session number from the name, so it's not visible.

Verify the deployment

To verify the Azure Virtual Desktop deployment, do the following:

- 1** Navigate to **Farm > Site > Settings** and select the **Features** tab. Verify that the **Enable Azure Virtual Desktop management** option is selected and the **Status** section says **Available** and displays the version number.
- 2** Navigate to **Farm > Site > Azure Virtual Desktop**. Select the following tabs and verify that corresponding components are configured and functioning properly:
 - **Providers**
 - **Workspaces**
 - **Host pools**
 - **Templates** (if you created a template, it should be listed on this tab)
 - **Hosts** (should list one or more session hosts)

Next step

Manage providers, workspaces, host pools, and templates (p. 197)

Publish resources

This topic describes how to publish Azure Virtual Desktop resources. Before using the instructions below, please make sure that you've read preceding sections describing how to deploy Azure Virtual Desktop in Parallels RAS.

To publish an Azure Virtual Desktop resource:

- 1** Select the **Publishing** category and click **Add** (at the bottom of the middle pane) to open the publishing wizard.
- 2** Select a resource type (e.g. **Application**) and click **Next**.

- 3 Select **Azure Virtual Desktop** as the publishing source and click **Next**.
- 4 Select **Installed Applications** to select an application from a list (or select **Single application** to specify the application properties yourself).
- 5 Select a host pool from which to publish an application. Note that depending on what is chosen to be published (Application or Desktop), the corresponding host pools with the same publishing type are shown.
- 6 Select an application to publish (or specify the application properties yourself if you selected **Single application**).
- 7 On the next page, specify the initial status of the resource. Choose from **Enabled** (end users can launch the resource), **Disabled** (the resource will not appear in Parallels Client), **In maintenance** (the resource will appear in Parallels Client but users will not be able to launch it). When a resource is in maintenance and a user tries to launch it, they will see a message. To customize the message, click the **Configure** button. For more info, see **Site defaults (Publishing)** (p. 245).
- 8 Click **Next** and then click **Finish**.

To see the properties of the application that you just published:

- 1 Select the application in the **Published Resources** list.
- 2 Review the information and verify that everything is as expected.
- 3 Select the **Filtering** tab and examine the **User Assignment** section. The list should show users or groups mapped to Azure Active Directory objects. This is important because in order for users to see and use this published resource, they must be assigned to the application group in the host pool from which this application is published. In addition to Azure Active Directory user assignment, you can also configure Parallels RAS filtering rules in the section above **User Assignment**. For details, see **Using filtering rules** (p. 247).

You can now open the application from Parallels Client. For details, see **Using Parallels Client with Azure Virtual Desktop** (p. 214).

Using Parallels Client with Azure Virtual Desktop

After you deployed Azure Virtual Desktop in Parallels RAS and published resources from it, you can use Parallels Client to access applications and/or desktops that you published. This topic describes Parallels Client requirements and provides additional information about accessing published resources.

Requirements

Parallels Client requirements for opening Azure Virtual Desktop apps and desktops are as follows:

- Parallels Client for Windows version 18 or newer (Basic or Full versions).
- Azure Virtual Desktop resources are shown in Parallels HTML5 Client, but can only be launched using Parallels Client for Windows installed on the same machine.

Note: At the time of this writing, it is not yet supported to use other platform-specific Parallels Clients to launch Azure Virtual Desktop resources. In these clients, Azure Virtual Desktop resources are filtered out and not shown.

- Supported Windows versions: Windows 7 SP1 or later. Note that Windows Server operating systems are not supported.
- Windows updates: Update for Universal C Runtime for Windows (KB2999226). Microsoft Windows 10 incorporates this by default.
- Microsoft .NET Framework 4.6.0 or later is required. Microsoft Windows 10 incorporates .NET Framework 4 and has it enabled by default.
- Microsoft Windows Desktop client, also known as Remote Desktop (MSRDC) client, must be installed. The client is downloaded and installed automatically upon launching an Azure Virtual Desktop resource from Parallels Client (if not already installed on a supported Windows client device). You may also download the client using the following link:
<https://go.microsoft.com/fwlink/?linkid=2068602>.

Note: If you are using the **Standard** client feature set option (p. 193) and Windows 10 and Windows 11 Enterprise Virtual Desktop as desktop OS where the Parallels Client is running (nested), the administrator needs to have the Windows Desktop client preinstalled using the per-device installation as highlighted by Microsoft in the following article: <https://docs.microsoft.com/en-us/windows-server/remote/remote-desktop-services/clients/windowsdesktop-admin>.

Accessing Azure Virtual Desktop resources

When you connect to Parallels RAS from Parallels Client, all published resources, including Azure Virtual Desktop resources, are listed and made available for a user to access. Note that Azure Virtual Desktop resources are only shown in Parallels Client running on supported versions of Windows (see above).

If the **Client feature set** option is set to **Advanced** or **Advanced with feedback**, you can use advanced Parallels RAS features when running a published resource, such as RAS Universal Printing and Scanning, session prelaunch, accelerated file redirection, drag and drop, and others. If the option is set to **Standard**, these features will not be available. To view and modify this setting, navigate to **Farm > Site > Settings**, select the **Features** tab and select a desired setting in the **Client feature set** drop-down list.

Remote PCs

In This Chapter

Overview	216
Adding a Remote PC to a Farm.....	216
Configuring a Remote PC.....	220
Viewing Remote PC summary	222
Using computer management tools.....	222
Publishing from a Remote PC.....	223

Overview

In addition to RD Session Hosts, Virtual desktops, and Azure Virtual Desktop, resources can also be published from a standalone Remote PC running a supported version of Windows (p. 20). A Remote PC can be a physical box or a virtual machine treated as a physical PC, but typically they are physical computers. If you have virtual machines on your network, it makes sense to use them as part of the VDI infrastructure as described in the **VDI and Virtual Desktops** chapter (p. 126). However, if you don't need the guest VM cloning functionality or, for example, if your end users require full administrative permissions on a PC for customization, you can use a virtual machine as a Remote PC.

Note: Remote PCs can also be combined into pools and managed as pool members. Remote PC pools use the RAS VDI infrastructure and work differently than individual Remote PCs described in this chapter. For more information see **Remote PC pools** (p. 181).

This chapter describes how to add a Remote PC to a Farm and how to publish remote applications and desktop from it.

Adding a Remote PC to a Farm

Remote PCs can be added to a RAS Farm using one of the following methods:

- Admin-initiated enrollment (p. 217) from the RAS Console by specifying the PC's IP and MAC addresses and installing (remotely or directly) RAS Remote PC Agent on it.
- By allowing users to use self-service Remote PC enrollment (p. 218).

Read the subsequent sections for the description of each method.

Admin-initiated Remote PC enrollment

Requirements

In order to add a Remote PC to a RAS Farm, RAS Remote PC Agent must be installed on it. The requirements to push install RAS Remote PC Agent are as follows:

- The firewall must be configured on the server to allow push installation. Standard SMB ports (139 and 445) need to be open. See also **Port reference** for the list of ports used by Parallels RAS.
- SMB access. The administrative share (\\server\c\$) must be accessible. Simple file sharing must be enabled.
- Your Parallels RAS administrator account must have permissions to perform a remote installation on the PC. If it doesn't, you'll be asked to enter credentials of an account that does.
- The PC should be joined to an AD domain. If it's not, the push installation may not work and you will have to install the Agent on it manually. Please see **Installing Remote PC Agent manually**.

Adding a Remote PC to a Farm

Follow the below procedure to add a Remote PC to a Farm:

- 1** In the RAS Console, select the **Farm** category and click the **Remote PCs** node in the navigational tree.
- 2** Click **Add** in the **Tasks** drop-down menu to launch the setup wizard.
- 3** Specify the IP address or FQDN of a Remote PC. Click the **Get MAC** button to obtain the PC's MAC address. To automatically resolve IP address to FQDN, enable the global **Name Resolution** option. For details, see **Host name resolution (p. 430)**.
- 4** Click **Next**.
- 5** In this step, the Parallels RAS checks if the Remote PC Agent is installed on the specified PC. If it's not installed, click **Install** to push install the agent on the PC. If the push installation of Remote PC Agent fails for any reason, you can install it manually. See **Installing Remote PC Agent Manually** below.
- 6** Click **Add** to add the Remote PC to the Parallels RAS Farm.

Installing Remote PC Agent manually

You may need to install the Remote PC Agent manually if the automatic push installation cannot be performed for any reason. To do so:

- 1** Log into the PC where the Remote PC Agent is to be installed using an administrator account and close all other applications.

- 2 Copy the Parallels RAS installation file (`RASInstaller.msi`) to the PC and double click it to launch the installation.
- 3 Follow the onscreen instructions and proceed to the installation type page. Select **Custom** and click **Next**.
- 4 Click on the **Remote PC Agent** and select **Entire Feature will be installed on local hard drive** from the drop down menu.
- 5 Ensure that all other components are deselected and click **Next**.
- 6 Click **Install** to start the installation. Click **Finish** once the installation is finished.

Remote PC Agent does not require any configuration. Once the agent is installed, select the Remote PC name in the Parallels RAS Console and click **Troubleshooting > Check Agent**. If the agent is installed properly, the status should change to **Agent Installed**.

Uninstalling Remote PC Agent

To uninstall Remote PC Agent from a server:

- 1 Navigate to **Start > Control Panel > Programs > Uninstall a Program**.
- 2 Find **Parallels Remote Application Server** in the list of installed programs.
- 3 If you don't have any other Parallels RAS components on the server that you want to keep, right-click **Parallels Remote Application Server** and then click **Uninstall**. Follow the instructions to uninstall the program. You may skip the rest of these instructions.
- 4 If you have other RAS components that you want to keep on the server, right-click **Parallels Remote Application Server** and then click **Change**.
- 5 Click **Next** on the Welcome page.
- 6 On the **Change, repair, or remove** page, select **Change**.
- 7 On the next page, select **Custom**.
- 8 Select **Remote PC Agent**, then click the drop-down menu in front of it, and click **Entire feature will be unavailable**.
- 9 Click **Next** and complete the wizard.

Self-service Remote PC enrollment

As an alternative to the admin-initiated Remote PC enrollment described above, end users can add their preferred PCs to a RAS Farm in a self-service manner. A computer enrolled this way is automatically added to a RAS Farm as a Remote PC and a corresponding published desktop is created with access granted (through automatic filtering) to the user who performed the self-enrollment. The Remote PC can then be accessed by the user from any device, anywhere.

Note: This feature applies to standalone Remote PCs. It is not intended for Remote PCs configured through VDI technology.

Requirements

- Users enrolling their PCs must have local administrator privileges in Windows to install Parallels RAS Remote PC Agent.
- Mailbox must be configured in a RAS Farm to send an invitation email to end users.

Configure self-service enrollment

To configure self-service Remote PC enrollment:

- 1 In the RAS Console, navigate to **Farm > Site > Remote PCs**.
- 2 In the right pane, click the **Tasks** menu and choose **Self-service enrollment**. In the dialog that opens, specify the options described below.
- 3 To enable self-service enrollment, select **Allow** or **Allow until**. For the latter, specify the date and time. To disable the functionality (e.g. temporarily), select **Do not allow**.
- 4 In the **Settings** section, specify a publishing folder in which the PCs will appear as published resources. You can select an existing folder or create a new one. Click the **[..]** button to select or create a folder.
- 5 The **Remote PC invitation hash** field contains a hash that must be specified when enrolling a PC. The hash can also be copied from here and used separately for scripting purposes. IT administrators can make use of this hash along with the `msiexec` command to silently install and configure Remote PC on users' behalf. For details, see **Enrolling a PC** below.
- 6 To send an invitation email to users that will contain instructions on how to enroll their PCs, click the **Send via email** button.
- 7 In the dialog that opens, specify the recipients by typing (or pasting) their email addresses. You can also click the **[..]** button and select the recipients.
- 8 In the **Review the invitation email** text box, review or modify (if needed) the email. The variables used in the email are set internally and are substituted with their values in the actual email. To preview the final email text, click **Tasks > Preview**.
- 9 Click **Send** to send the email. If you don't want to send the email at this time, click **Cancel** to return to the previous dialog where you can click **OK** to save the changes.

Enrolling a PC

When a user receives the invitation email, they will follow the instructions that it contains to enroll their PC. The installation process consists of the steps described below.

Log in to the Remote PC. Download or copy the RASInstaller.msi file (the Parallels RAS installer) and run the following command with administrative privileges (the invitation email will contain this command with the hash value already in it):

```
msiexec /qb /i <RAS installer> ADDLOCAL=F_PCAGENT ADDFWRULES=1 SELFENROLL=<hash key>
[OVERRIDEUSER=user@domain] [OVERRIDEPAIP=ip of PA] [OVERRIDEHOST=published name]
```

The following arguments can be used to customize the enrollment. Such arguments are required if a Remote PC is not in the Active Directory domain:

- **OVERRIDEPAIP:** The IP address of one of the Publishing Agents in the Farm Site. Use this if the standard installation fails to connect using the system detected IP address.
- **OVERRIDEUSER:** Use this argument if you don't want to register the Remote PC to the user logged in to the machine.
- **OVERRIDEHOST:** Use this argument if you want to change the published item name from the hostname of the Remote PC.

Once the installation is complete, launch Parallels Client and log in using the local machine credentials or the ones specified in the **OVERRIDEUSER** argument. Look for your Remote PC using the IP address or the name specified via **OVERRIDEHOST** in the list of published resources and launch the desktop.

Configuring a Remote PC

To view the properties of a Remote PC, highlight the computer in the navigation tree and click **Tasks > Properties**. This opens the Remote PC properties dialog.

Properties

By default, a PC is enabled in the Farm. When it is disabled, published applications and virtual desktops cannot be served from it. To enable or disable a PC in the Farm, select or clear the **Enable Remote PC** option.

If the IP or MAC address of a Remote PC has changed, modify them using the **Remote PC** and **MAC Address** input fields.

The **Change Direct Address** option allows you to specify an IP address that Parallels Client can use to connect to the PC directly. This address is only used in the Direct Connection mode and it could be an internal or external IP address.

Note: The Wake On Lan option should be enabled in BIOS so the machine could be automatically turned on. If you are using a virtual machine, the option is usually supported by a hypervisor natively or via a 3rd party software. To test if the Wake On Lan option is turned on, close the **Remote PC Properties** dialog and then click the **Test Wake on LAN** button, which is located below the **Remote PCs** list.

Agent Settings

Each Remote PC in the Farm has a RAS Remote PC Agent installed to conduct communications between Parallels RAS and the PC. The agent can be configured on the **Agent Settings** tab page.

- **Logoff active session after:** The amount of time a session remains logged in after the user closes a published application. The default timeout is 25 seconds. Note that this only works for applications, but not published desktops (when a user closes a desktop, the session is logged off). This timeout is used to avoid unnecessary logins when a user closes one application and then opens another.
 - **Session readiness timeout:** The maximum amount of time it should require to establish a session. If the specified timeout is reached, and the session is still not ready, the user will see an error message and will have to try to log in again.
 - **Port.** Specify a different remote desktop connection port number if needed.
 - **Preferred Publishing Agent:** Select a Publishing Agent with which the Remote PC Agent should communicate. This can be helpful when Site components are installed in multiple physical locations communicating through WAN. You can decrease network traffic by specifying a more appropriate Publishing Agent.
 - **Allow URL/Mail redirection.** When a user tries to open a URL or an HTML Mailto link in a remote application, the link can be redirected to the client computer and open in a local default application (a web browser or email client) instead of an application on the remote host. This option allows you to enable or disable the redirection. Click the **Configure** button to choose from the following options:
 - a **Replace registered application** — this option uses an alternative method of redirecting a link. It replaces the default web browser and mail client with "dummy" apps on the remote server side. By doing so, it can intercept an attempt to open a link and redirect it to the client computer.
 - b **Support Windows Shell URL namespace objects** — the Shell URL namespace objects support means that Parallels RAS can intercept actions in published applications that use Shell namespace API to open links, which is a standard behavior in most applications. The ability to disable the support for Shell URL namespace objects is for compatibility with older versions of Parallels RAS.
 - **Enable drag and drop:** Allows you to set how the drag and drop functionality works in Parallels Clients. Click **Configure** and choose from "Disabled" (no drag and drop functionality), "Sever to client only" (drag and drop to a local application only), "Client to server only" (drag and drop to a remote application only), "Bidirectional" (drag and drop in both directions).
- Note:** At the time of this writing, the drag and drop functionality is only supported on Parallels Client for Windows and Parallels Client for Mac.
- **Allow file transfer command (HTML5 and Chrome clients):** Enables file transfer in a remote session. To enable file transfer, select this option and click the **Configure** button. For more information, see **Configuring remote file transfer** (p. 408).
 - **Enable drive redirection cache:** Improves user experience by making file browsing and navigation on redirected drives much faster. For details, see **Drive redirection cache** (p. 107).

RDP Printer

The **RDP Printer** tab allows you to configure the renaming format of redirected printers. The format may vary depending on which version and language of the server you are using.

Set your RDP Printer Name Format specifically for the configured server by choosing any of the below options from the RDP Printer Name Format drop down menu:

- Printername (from Computername) in Session no.
- Session no. (computername from) Printername
- Printername (redirected Session no)

The other RDP Printing options available in the RDP Printer tab are:

- Remove session number from printer name
- Remove client name from printer name

Configure logging

A Remote PC is monitored and logs are created containing relevant information. To configure logging and retrieve or clear existing log files, right-click a Remote PC, choose **Troubleshooting > Logging** in the context menu, and then click **Configure**, **Retrieve**, or **Clear** depending on what you want to do. For the information on how to perform these tasks, see the **Logging** (p. 451) section.

Viewing Remote PC summary

In addition to the Remote PCs editor described in this chapter, you can also see the summary about the available Remote PCs. To do so:

- 1** In the RAS Console, select the **Farm** category and then select the **Site** node in the middle pane.
- 2** The available servers are displayed in the **Remote PCs** group in the right pane.
- 3** To go to the Remote PCs editor, right-click a server and choose **Show in the Editor**.

For additional info, see **Sites in the RAS Console** (p. 45).

Using computer management tools

You can perform standard computer management tasks on a server right from the RAS Console. These include Remote Desktop Connection, PowerShell, Computer Management, Service Management, Event Viewer, IPconfig, Reboot, and others. To access the **Tools** menu, select a server, click **Tasks > Tools** and choose a desired tool. For requirements and usage information, see **Computer management tools** (p. 431).

Publishing from a Remote PC

This section describes how to publish resources hosted by a standalone Remote PC. The publishing functionality described here is accessed from the **Publishing** category in the RAS Console.

Read on to learn how to publish resources from a Remote PC.

Publishing a desktop

To publish a desktop from an RD Session Host:

- 1 In the RAS Console, select the **Publishing** category and click the **Add** icon below the **Published Resources** tree. This will launch the publishing wizard.
- 2 In the first step of the wizard select **Desktop** and click **Next**.
- 3 On the **Select Desktop Type** page, select **Remote Desktop PC** and click **Next**. The **Remote PC Desktop** page opens.
- 4 Specify a name, an optional description, and change the icon if needed.
- 5 Click the [...] button next to the **Selected Remote PC** field to specify from which Remote PC the desktop should be published. In the box that opens, double-click a PC to select it.
- 6 Select the desired **Desktop Size** properties.
- 7 On the next page, specify the initial status of the resource. Choose from **Enabled** (end users can launch the resource), **Disabled** (the resource will not appear in Parallels Client), **In maintenance** (the resource will appear in Parallels Client but users will not be able to launch it). When a resource is in maintenance and a user tries to launch it, they will see a message. To customize the message, click the **Configure** button. For more info, see **Site defaults (Publishing)** (p. 245).
- 8 Click **Finish** to publish the desktop.

Publishing an application

To publish an application from a Remote PC:

- 1 In the RAS Console, select the **Publishing** category and then click the **Add** icon below the **Published Resources** tree (or right-click inside the **Published Resources** box and click **Add** in the context menu). This will launch the publishing wizard.
- 2 On the **Select Item Type** wizard page, select **Application** and click **Next**.
- 3 On the **Select Server Type** page, select **Remote PC** and click **Next**.
- 4 On the **Select Application Type** page, select **Single Application** and click **Next**. The **Remote PC Application** page opens.

- 5 Enter a name and an optional description.
- 6 In the **Run** drop-down menu, specify if the application should run in a normal window, maximized, or minimized.
- 7 In the **Target** field, specify the application that you want to publish. You may click the [...] button to browse for it.
- 8 In the **Start in** field, specify (or browse for) the "start in" folder. Use Windows environment variables if you are manually entering the path.
- 9 (Optional) In the **Parameters** input field, specify the parameters to pass to the application when it starts.
- 10 Click the [...] button in the **Remote PC Settings** section to select a Remote PC from which the application should be published. In the box that opens, double-click a PC to select it.
- 11 Select the **Persistent** option to mark a guest VM as persistent the first time the user connects to it.
- 12 On the next page, specify the initial status of the resource. Choose from **Enabled** (end users can launch the resource), **Disabled** (the resource will not appear in Parallels Client), **In maintenance** (the resource will appear in Parallels Client but users will not be able to launch it). When a resource is in maintenance and a user tries to launch it, they will see a message. To customize the message, click the **Configure** button. For more info, see **Site defaults (Publishing)** (p. 245).
- 13 When done, click **Finish** to publish the application.

Publishing a web application

A web application is like any other application that you can publish using the standard application publishing functionality. However, to simplify publishing of straight URL links to web applications, a separate publishing item type is available that allows you to accomplish this task with minimal number of steps.

To publish a web application:

- 1 In the RAS Console, select the **Publishing** category and then click the **Add** icon below the **Published Resources** tree (or right-click inside the **Published Resources** box and click **Add** in the context menu). This will launch the publishing wizard.
- 2 On the **Select Item Type** wizard page, select **Web Application** and click **Next**.
- 3 On the **Select Server Type** page, select **Remote PC** and click **Next**.
- 4 On the **Remote PC Web Application** wizard page that opens, specify the web application name, description, window state, and the URL. Select the **Force to use Internet Explorer** option if needed. To browse for a specific application icon, click **Change Icon**.
- 5 In the **Remote PC Settings** section, click the [...] button to select a Remote PC.

- 6 On the next page, specify the initial status of the resource. Choose from **Enabled** (end users can launch the resource), **Disabled** (the resource will not appear in Parallels Client), **In maintenance** (the resource will appear in Parallels Client but users will not be able to launch it). When a resource is in maintenance and a user tries to launch it, they will see a message. To customize the message, click the **Configure** button. For more info, see **Site defaults (Publishing)** (p. 245).
- 7 Click **Finish** to publish the application.

Publishing a network folder

You can publish a filesystem folder via UNC path to open in Windows explorer. To minimize the number of configuration steps, a special publishing item is available that allows you to publish a network folder from a PC.

To publish a network folder:

- 1 In the RAS Console, select the **Publishing** category and then click the **Add** icon below the **Published Resources** tree (or right-click inside the **Published Resources** box and click **Add** in the context menu). This will launch the publishing wizard.
- 2 On the **Select Item Type** wizard page, select **Folder on the file system** and click **Next**.
- 3 On the **Select Server Type** page, select **Remote PC** and click **Next**.
- 4 On the **Remote PC UNC Folder** wizard page, specify the usual application properties.
- 5 In the **UNC path** field, enter the UNC path of the folder you wish to publish. Click the **[...]** button to browse for a folder (it may take some time for the **Browse for Folder** dialog to open).
- 6 In the **Remote PC Settings** section, select the **[...]** button and then select a Remote PC from the list.
- 7 On the next page, specify the initial status of the resource. Choose from **Enabled** (end users can launch the resource), **Disabled** (the resource will not appear in Parallels Client), **In maintenance** (the resource will appear in Parallels Client but users will not be able to launch it). When a resource is in maintenance and a user tries to launch it, they will see a message. To customize the message, click the **Configure** button. For more info, see **Site defaults (Publishing)** (p. 245).
- 8 Click **Finish** to publish the folder and close the wizard.

Publishing a document

To publish a document from a Remote PC clone:

- 1 In the RAS Console, select the **Publishing** category and then click the **Add** icon below the **Published Resources** tree (or right-click inside the **Published Resources** box and click **Add** in the context menu). This will launch the publishing wizard.
- 2 On the **Select Item Type** wizard page, select **Document** and click **Next**.

- 3** Select **Remote PC** and click **Next**.
- 4** Specify the content type of the document you want to publish. You can select the content type from the predefined list or specify a custom content type in the **Custom content types** input field.
- 5** Click **Next**.
- 6** On the **Remote PC Application** page, enter a name, an optional description, a desired window state, and an icon if needed.
- 7** Use the [...] button next to the **Target** input field to browse for the document. All other fields will be automatically populated. To edit any of the auto populated fields, highlight them and enter the required details.
- 8** (Optional) In the **Parameters** input field, specify the parameters to pass to the application when it starts.
- 9** Click the [...] button in the **Remote PC Settings** sections to browse for a Remote PC from which the document should be published. In the box that opens, double-click a PC to select it.
- 10** On the next page, specify the initial status of the resource. Choose from **Enabled** (end users can launch the resource), **Disabled** (the resource will not appear in Parallels Client), **In maintenance** (the resource will appear in Parallels Client but users will not be able to launch it). When a resource is in maintenance and a user tries to launch it, they will see a message. To customize the message, click the **Configure** button. For more info, see **Site defaults (Publishing)** (p. 245).
- 11** Click **Finish** to publish the document.

Session Management

In This Chapter

Overview	227
Session information	228
Monitoring settings	231
Managing sessions	231
The Resources tab	233

Overview

When users connect to Parallels RAS and establish a session, the session information is displayed in the Parallels RAS Console in the following locations:

- The **Sessions** category (new since Parallels RAS 18.1).
- The **Sessions** tab in RD Session Hosts, VDI, and Azure Virtual Desktop views (**Farm > Site > RD Sessions Hosts > VDI > Azure Virtual Desktop**).

The **Session** category displays user sessions for all available host types, including RD Sessions Hosts, VDI, and Azure Virtual Desktop. This is the place where you can view all current sessions irrespective of the type of a server hosting a session. Individual **Sessions** tabs display sessions for their respective host types.

Sessions category

When you select the **Session** category (in the main category list), the following two tabs are displayed in the right pane of the RAS Console:

- **Users:** Lists user sessions for all available host type.
- **Resources:** Lists currently running published resources (apps and desktops) from hosts of all types.

You can filter the lists in the **Sessions** category by clicking **Tasks > Search** (or clicking the magnifying glass icon) and specifying the criteria in one or more column headings. For example, you can filter the list on the **Users** tab by host type using the **Source** column, which can contain one of the following values:

- **RDSH**: RD Session Host
- **VDI**: Virtual desktop
- **RemotePC**: Remote PC through VDI
- **AVD**: Azure Virtual Desktop

Sessions tabs

The **Sessions** tabs display user sessions for their respective host types. To see sessions for a particular host, you can filter the list by host name.

Please note that when you open the **Sessions** category or a **Sessions** tab, some of the columns in a list may not be populated right away. This is because it takes some time to calculate these values. The examples of such columns include **Logon duration**, **UX Evaluator**, **Latency**. Simply wait a few seconds and the values will appear in the list.

Most of the information in this chapter is common to both the **Sessions** category and **Sessions** tabs. Specifics and differences are described where applicable.

Session information

To see the complete information for a specific session, right-click it and choose **Show information**. This opens the **Session Information** dialog where session properties are grouped by functionality.

The following groups are displayed:

- **Session Setup**: Contains general session information.
- **Logon Details**: Displays logon metrics that can be used to evaluate the logon process.
- **Session Details**: Displays the current session state, logon time, in/out data size, and general session information.
- **Connection Details**: Displays connection and authentication details.
- **User Experience**: Displays metrics that can be used to evaluate user experience.
- **Client Details**: Displays information about the user device and Parallels Client type and version.

Parallels RAS 18 introduces over 25 new session detail metrics available. The following tables give an overview of these new and some of the important preexisting metrics.

Session Setup

Metric	Description
Session host*	Session host name
Source*	Sessions category only. Host type: RDSH (even if its through VDI), VDI, RemotePC (through VDI only), AVD

* New since Parallels RAS 18.1

Logon details

Metric	Description
Logon duration*	Time taken to logon excluding the time waiting on UI.
Logon duration breakdown*	Connection time Authentication duration Host preparation (inc. load balancing algorithm) User profile load time RAS Policies lookup Group Policy processing Desktop loading Other
User Profile*	User Profile method in use: FSLogix, User Profile Disk, or Other (also contains additional information, such as error code).

* New since Parallels RAS 18.0

User Experience

Metric	Description
UX Evaluator*	This is the time interval measured at the client between the first step (user action) and the last step (graphical response displayed).
Connection quality*	Connection quality rating (poor – excellent)
Latency*	Network latency
Transport Protocol*	TCP or UDP (over RDP)
Bandwidth availability*	Bandwidth availability as seen from the client
Reconnects*	Number of reconnects the current session suffered from inception (excluding graceful ones)
Last Reconnects*	Number of reconnects suffered from the current device session (excluding graceful ones)
Disconnect reason*	The last session disconnect reason

* New in Parallels RAS 18.0

Session details

Metric	Description
Session State	Active, Idle, Disconnected, etc
Logon time	Time and date when the session was established
Session Length	Time the session has been established
Idle Time	Time the session has been idle
Incoming Data*	Amount of data received from the client
Outgoing Data*	Amount of data sent to the client
Resolution	Session resolution
Color Depth	Session colors depth
Bandwidth Usage*	Bandwidth used by the client

* New since Parallels RAS 18.0

Client details

Metric	Description
Device name	Name of the device from which the session was established
IP Address	Client private IP address
Client OS*	The operating system on which the client is running
Client OS version*	The operating system version on which the client is running
Client version*	The Parallels Client version is use

* New since Parallels RAS 18.0

Export session information

To export the session information to a CSV file, click the **Export** button in the **Session Information** dialog and specify the location and file name.

You can also export session information from the main session list by clicking **Tasks > Export**.

Note that depending on what is selected in the list, the following will be exported:

- A single session — the information about that session is exported.
- Multiple sessions — the information for all selected sessions is exported.
- No selection — the information about all current sessions is exported. Exported CSV includes the exported session details along with export detail in the following format:

Session details (%Server type% such as RD Session Hosts) from Parallels RAS Farm %Farm name% and Site %Site name% exported by %Administrator% on %date% at %time%

Monitoring settings

The **Monitoring Settings** functionality allows you to add colors to thresholds to identify Warning and Critical levels for better aid to Administrators or helpdesk.

To configure monitoring settings, in the **Sessions** category or **Sessions** tabs, click **Tasks > Monitoring settings**. The dialog opens where you can configure settings for various session metrics:

- 1 Select a metric for which you want to enable color coding.
- 2 Specify a threshold in the **Warning** and **Critical** columns. The Warning threshold is denoted by the orange color. The Critical threshold is denoted by the red color.
- 3 If you want to see just the Critical color (red), then set both thresholds to the same value, in which case only the red color will be used when the threshold is reached.

When a metric with color coding enabled is below any of the specified thresholds, it is highlighted with the green color in the session list. When a threshold is reached, the value of a metric is highlighted using the corresponding threshold color (orange or red). Note that critical threshold value can be greater or equal to warning threshold value. In case both warning and critical values are equal then the critical color coding is used which is red.

Monitoring Settings are set globally, which means that other RAS admins will be able to see and change them.

Managing sessions

To manage a session (or multiple sessions at the same time), select one or more sessions and then use the **Tasks** drop-down menu to choose from the following actions:

- **Refresh.** Refresh the list.
- **Disconnect.** Disconnect the selected session(s).
- **Log off.** Log off the session(s).
- **Send message.** Opens the **Send Message** dialog where you can type and send a message to the session owner(s).
- **Remote control.** Remotely control the selected user session. To establish a connection, domain or local Windows account credentials (whichever the user used to log in to this computer) of the current RAS Console administrator will be used. Note that the current user (specifically if it's the local Windows user) may not be permitted to connect to the remote computer. In such a case, use the **Remote control (prompt)** option (described below). See also the **User session remote control** subsection below for important information.
- **Remote control (prompt).** Same as above but prompts you to enter credentials. Use this option when the current user credentials cannot be used to control a session.

- **Show processes.** Display and manage running processes. See **Managing processes** below for details.

User session remote control

The **Remote Control** and **Remote control (prompt)** menu options (see above) allow you to shadow a user RDS session. There are limitations as described below:

- Parallels RAS cannot shadow RDS sessions running on Windows 7 and Windows Server 2008 R2 (plain Windows Server 2008 is fine). This doesn't work even with native tools.
- If you need to shadow a user session running on Windows Server 2008, the RAS console must also be running on Windows Server 2008. If the RAS console is installed on a later version of Windows Server, shadowing will not work. As a workaround, you can add an RD Session Host running Windows Server 2008 to the Farm, publish the Parallels RAS console from it, and then use the console remotely to manage user RDS sessions running on Windows Server 2008. Please note that to finish a remote control session, the administrator must log off from the RAS console remote session. This is a limitation of the shadow.exe utility from Microsoft that doesn't take any arguments that would allow us to add a control like a bar, a button, or a key combination.

Managing processes

The **Tasks > Show processes** option opens the **Running Processes** dialog where you can view running processes for one or more hosts.

Note: You can also open the **Running Processes** dialog by right-clicking a server in the main host list and choosing **Show Processes**. This will open the **Running Processes** dialog with a filter applied to it to display only the processes that belong to the selected host.

In the **Running Processes** dialog, use the **Show processes from** drop-down menu to filter the list using the following options:

- **Selected Session.** Displays processes for the session selected in the **Sessions** list.
- **Selected Server.** Displays all running processes for the server on which the selected session is running.
- **All Servers.** Displays all running processes for all available servers.

You can also filter the list by specifying a search criteria for one or more columns. To do so, click the magnifying glass icon (top right) and then type a desired text in one or more columns. The list is filtered as you type to match the specified criteria.

The **Tasks** drop-down menu in the **Running Processes** dialog includes the following options:

- **Refresh.** Refresh the list.
- **Kill process.** Kill the selected process.

- **Go To Published Item.** Enabled when you select a process that belongs to a running published resource. Brings up the main Parallels RAS Console window and navigates to the corresponding published resource.
- **Disconnect.** Disconnect the session.
- **Log off.** Log off the session.
- **Send message.** Send a message to the session owner.
- **Remote control.** Remotely control the selected user session.

The Resources tab

The **Resources** tab in the **Sessions** category displays currently running published resources (apps and desktops).

Some of the notable columns in the list are:

- **ID:** The published resource ID (as seen in the **Publishing** category).
- **Published name:** Published resource name (as seen in the **Publishing** category).
- **User:** Session owner.
- **Session ID:** Session ID.
- **Session Host:** Session host name.
- **Source:** Session source (RDSH, VDI, RemotePC, AVD).

To perform a task on a resource, click the **Tasks** menu. Some of the tasks include:

- **Search:** Allows you to filter the list using one or more columns (e.g. User, Session ID, Session Host, etc.).
- **Show user session view:** Switches to the **Users** tab and applies a filter to display the session to which the selected resource belongs.
- **Go to published resource:** Takes you to the **Publishing** category and displays the selected resource information.
- **Show information:** Displays the resource summary info and the session information. The session information includes the same metrics as described in **Session information** (p. 228).

Published Resources Management

In This Chapter

Overview	234
General management tasks.....	235
Manage published applications	236
Manage published desktops.....	240
Manage published documents	241
Manage folders	243
Site defaults (Publishing).....	245
Using filtering rules	247
Configuring preferred routing	250
Understanding session prelaunch.....	252
Checking effective access	253
Specifying client settings	254
Quick keypad	256

Overview

Resources that can be publish in Parallels RAS include:

- Applications
- Containerized applications
- Desktops
- Documents
- Web applications
- Network folders

For the information on how to publish resources from various types of servers, follow the links below:

- **Publishing from an RD Session Host** (p. 118)
- **Publishing from a guest VM** (p. 176)
- **Publish from Azure Virtual Desktop** (p. 213)
- **Publishing from a Remote PC** (p. 223)

This chapter describes management tasks that you can perform on resources that have been already published.

General management tasks

To view published resources, select the **Publishing** category in the Parallels RAS Console. In the middle pane, expand the **Published Resources** node (if it's collapsed) to see the resources.

Right-click a resource to open a context menu. The menu has the following options:

- **Add:** Starts the publishing wizard, which you can use to publish a resource.
- **New Folder:** Allows you to add a folder to the **Published Resources** tree. Folders are described in the **Manage folders** section (p. 243).
- **Find:** Allows you to search the list for a resource by name.
- **Duplicate:** Duplicates a selected resource. You can publish multiple resources of the same type, but configure them differently according to your needs.
- **Delete:** Deletes a published resource from the Parallels RAS Farm. This only removes the published resource item from the Farm. The actual application is not affected. To avoid accidental deletions, a dialog box is displayed asking for your confirmation.
- **Status:** Choose from **Enabled**, **Disabled**, **In maintenance**. When a resource is disabled or in maintenance, it is unavailable to users. A disabled resource will not appear in Parallels Client in the available resources list. A resource that is in maintenance will appear, but will be grayed out (in HTML5 Client it will say so in the resource name). If a user tries to open the resource, a message is displayed. You can customize this message in Site defaults for published resources (p. 245).

Note that when you set the status of a published folder, all subfolders (if any) and resources that it contains will inherit the parent folder's status.

- **Delegate Permissions:** Opens the **Delegate Permissions** dialog where you can add users and grant them publishing permissions.
- **Settings audit:** Allows you to see recent changes to published resources and revert them. The changes that can be reverted include Create, Delete, and Update.
- **Verify Target(s):** Verifies that the target specified for the selected resource is valid. To see the target, select a resource and then click the **Application** tab.
- **Convert Filters to Secure Identifiers:** If filtering for a resource is specified using WinNT or LDAP, you can use this option to convert it to SID (Secure Identifier). For more information, see **Using filtering rules** (p. 247).
- **Running Instances:** Opens the **Running Processes** dialog. For more information about the dialog, please see **Managing sessions > Managing running processes** (p. 116). When the dialog is opened, a filter is applied to the process list to include only the processes for the selected published resource (a resource ID is used as a value). You can further filter the list to include only the process for a particular user (the **Username** column).

The action items at the bottom of the screen allow you to perform the following actions:

- **Add:** Same action as the **Add** menu item described above.
- **New Folder:** Same action as the **New Folder** menu item described above.
- **Delete:** Same as the **Delete** menu item described above.
- **Move Up:** Moves a selected published resource item up the list.
- **Move Down:** Moves a selected published resource item down the list.
- **Disable:** Same as the **Disable** menu item described above.
- **Sort.** Sorts resources alphabetically. For this action item to become enabled, you must select the **Published Resources** node (the topmost one) or a folder containing individual items.
- **Find:** Same as the **Find** menu item described above.
- **Running Instances:** Same as the **Running Instances** menu item described above.
- **Effective Access:** Allows you to view which published resources are available for a specific user. For complete details, see **Checking effective access** (p. 253).

After making any changes to published resources, please don't forget to click the **Apply** button to commit them to the Parallels RAS Farm.

Manage published applications

Configuring a published application

When publishing an application using a wizard, you specify multiple application parameters, such as name, executable path, etc. You can modify these options after the application is published.

To modify published application settings:

- 1** In the RAS Console, select the **Publishing** category and then select a desired application in the **Published Resources** tree.
- 2** Use the tabs in the right pane to change the application settings as described in the following subsections.

Publish from — configure from which servers the application is published

You can specify RD Session Hosts from which an application is published on the **Publish From** tab. The following options are available:

- **All Servers in Site:** The application will be published from all servers on which it is installed.
- **Server Groups:** Select this option and then select one or more RD Session Host groups from which the application should be published.

- **Individual Servers:** Select this option and then select one or more individual RD Session Hosts.

Application – configure application and hosting server settings

The **Application** tab displays application- and server-specific settings.

You can modify the basic application settings (name, description, window mode) as needed. Other settings include:

- **Status:** Choose from **Enabled**, **Disabled**, **In maintenance**. When a resource is disabled or in maintenance, it is unavailable to users. A disabled resource will not appear in Parallels Client in the available resources list. A resource that is in maintenance will appear, but will be grayed out (in HTML5 Client it will say so in the resource name). If a user tries to open the resource, a message is displayed. You can customize this message in Site defaults for published resources (p. 245).

Note that when you set the status of a published folder, all subfolders (if any) and resources that it contains will inherit the parent folder's status.

- **Start automatically when user logs on:** Select this option to start an application as soon as a user logs in. This option works on desktop versions of Parallels Client only.
- **Exclude from session prelaunch:** For details about using this option, see **Understanding session prelaunch** (p. 252).

The **Server settings** section contains server-specific options that you can configure. If an application was published from multiple servers, the **Server(s)** drop-down list can be used to select individual servers and set **Target**, **Start in**, and **Parameters** values for a particular server. As an example, you should do this when different servers have the application installed in different folders, so that the **Target** and **Start in** field values would be valid on each server.

To save the currently displayed server settings as default, click the **Save as Default Settings** button. To apply the saved default settings to a server, click the **Use Default Settings** button. These two buttons give you the flexibility of using custom settings or defaults in different server configuration scenarios. Please note that when you save settings as default, Parallels RAS will check if this Site has applications with per-server settings and will display a message asking if you would like those servers to use the new default settings. If you say, "No", the servers will keep their unique settings. The defaults will still be saved.

To verify that the specified **Target** and **Start In** values are correct for all servers, click the **Verify Target(s)** button. The **Target Verifier** dialog opens listing each server and the verification status in the **Progress** column. If the application is installed at a different path on one of the servers, the **Progress** column will indicate an error. In such a case, close the **Target Verifier** dialog and then select the server in the **Server(s)** drop-down list. Specify new values in the **Target**, **Start In**, and (if necessary) **Parameters** fields specific for that server. Click **Apply** to save your changes.

The **Target Verifier** dialog can also be used to verify the targets for all published applications at once. To do so, right-click **Published Resources** (the root node of the **Published Resources** tree) and then click **Verify Target(s)** in the context menu. This time, the **Target Verifier** dialog will contain all published applications and their verification status.

Filtering

Please see **Using filtering rules** (p. 247).

Routing

Please see **Configuring preferred routing** (p. 250).

Shortcuts — configuring shortcut options for a published application

Click the **Shortcuts** tab to enable the creation of the application shortcut on the user's desktop and in the Start and Auto Start folders. When the **Auto Start** option is selected, the application will start automatically on computer startup. To use Site default settings, select the **Inherit default settings** option. You can view or modify Site defaults by clicking the **Site Defaults** link. See **Site defaults (Publishing)** for more info (p. 245).

Note: Shortcuts are not available on all operating systems.

File extension — configuring file extension associations

To modify file extension association for a particular published application, click the **File Extensions** tab.

To add, remove, or modify an entry, select the **Associate File Extensions** option. To add a new extension to the list, click **Add** in the **Tasks** drop-down menu (or click the + icon) and specify the desired extension.

To modify an existing association, highlight the extension and click **Properties** in the **Tasks** drop down menu (or double-click the **Parameters** column) and type the parameter.

License — configuring licensing options for published applications

Click the **Licensing** tab to configure the following licensing options:

- **Disable session sharing:** If this option is enabled, it allows you to isolate a given published application to one session. If the same application is launched more than once, the instances of the application will share the same sessions. A different application, however, will start in its own session.
- **Allow users to start only one instance of the application:** If this option is enabled, a user can only launch a single instance of the application.

- **Concurrent Licenses.** Use this option to specify the maximum number of concurrent instances the application can run. E.g. if the license of the application allows you to only run 10 instances of the application, set the Concurrent licenses option to 10 so once such limit is reached, other users cannot initiate other instances.
- **If limit is exceeded:** From this drop down menu you can specify what action should the Parallels RAS take in case any of the above licensing configured limits are exceeded.

To use Site default settings, select the **Inherit default settings** option. You can view or modify the default settings by clicking the **Site Defaults** link. See **Site defaults (Publishing)** for more info (p. 245).

Display — configuring display settings for a published application

On the **Display** tab, you can configure the following options:

- **Wait until all RAS Universal Printers are redirected before showing the application:** Enable this option to wait for printers to be redirected before the application is loaded. You can also specify the maximum wait time (in seconds) for the Universal Printers to be redirected. Please note that redirecting a printer may take some time. To avoid confusion, a progress bar is shown to the user while the printers are being redirected.
- **Color Depth, Resolution, Width, Height:** Select the desired display settings for the application.
- **Start the application as maximized when using mobile clients:** This option applies only to Parallels Client running on mobile devices. When the option is selected, the application will start on a mobile device in the maximized state. This gives users the best experience while working with a remote application. This option gives the RAS administrator an easy way to always maximize an application without taking any additional steps.
- **Start in fullscreen mode for Wyse ThinOS:** If selected, the application will start in fullscreen mode in Wyse ThinOS. In some cases, the bottom part of an application may be covered by the taskbar, not allowing to see the entire application window. When this option is used, the taskbar will be hidden and the entire application window will be visible.

Note that to specify custom display values, the **Inherit default settings** checkbox must be cleared; otherwise Site defaults settings are used. To view and modify Site defaults, click the **Site Defaults** link. See **Site defaults (Publishing)** for more info (p. 245).

Quick keypad

The **Quick Keypad** section allows you to select a Quick Keypad template that should be assigned to this application. The **Quick Keypads** link below the drop-down list takes you to the **Quick Keypad** category in the console where you can configure keypad templates. For more information, please see the **Quick keypad** section (p. 256).

Manage published desktops

Configuring a published desktop

When publishing a desktop using a wizard, you have to specify the desktop settings, such as display size, etc. You can modify these options after the desktop has been published.

To modify a published desktop, select it in the **Published Resources** tree in the **Publishing** category.

Sites — configuring from which sites a published desktop is available

By default, a published desktop is available through all of the available sites. To restrict access to a specific Site or a Site group, select a desktop in the **Published Resources** tree and then click the **Sites** tab in the right pane. Select the sites from which the desktop should be available.

Note: For the **Sites** tab to be available, you need more than Site in a farm.

Publish from — configuring from which RD Session Hosts a desktop is published

When configuring an RD Session Host desktop, you can specify from which servers it should be published. To do so, click the **Publish From** tab and select the desired servers.

Desktop — configuring desktop name, size and other properties

Depending on the desktop type, click the **Desktop**, **Remote PC Desktop**, or **Virtual Desktop** tab to configure the desktop name, description, icon, resolution, status, and other settings.

- **Status:** Choose from **Enabled**, **Disabled**, **In maintenance**. When a resource is disabled or in maintenance, it is unavailable to users. A disabled resource will not appear in Parallels Client in the available resources list. A resource that is in maintenance will appear, but will be grayed out (in HTML5 Client it will say so in the resource name). If a user tries to open the resource, a message is displayed. You can customize this message in Site defaults for published resources (p. 245).

Note that when you set the status of a published folder, all subfolders (if any) and resources that it contains will inherit the parent folder's status.

- **Connect to administrative session:** Select this option if you want users to connect to the administrative session. Note that a user connecting to a desktop with this option enabled must have administrative privileges; otherwise "Access is denied" error will be shown to the user.
- **Start automatically when user logs on:** Select this option if you want to open a desktop as soon as a user logs in. For the information about **Exclude from session prelaunch** option, see **Understanding session prelaunch** (p. 252).

- **Desktop Size:** Select a desired desktop size from the drop-down list.
- **Multi-Monitor:** Select whether the multi-monitor should be enabled, disabled, or whether the client settings should be used.

Filtering

Please see **Using filtering rules** (p. 247).

Routing

Please see **Configuring preferred routing** (p. 250).

Shortcuts — configuring shortcut options for a published desktop

Click the **Shortcuts** tab to enable the creation of a shortcuts on the user's desktop and in the Start and Auto Start folders. When the Auto Start shortcut is enabled, the application will start automatically on computer startup. To use Site default settings, select the **Inherit default settings** option. See **Site defaults (Publishing)** for more info (p. 245).

Note: This option is not available on all operating systems.

Manage published documents

Configuring a published document

When publishing a document using a wizard, you have to specify the document settings. These options can be modified after the document has been published.

To modify a published document, select it in the **Published Resources** tree in the **Publishing** category and then use the tabs in the right pane to configure the published document settings.

Sites — configuring from which sites a published document is available

By default, a published document is available through all available sites. To restrict access to a specific Site or a Site group, click the **Sites** tab in the right pane. Select the sites from which the document should be available.

Note: For the **Sites** tab to be available, you need more than one Site in a Farm.

Publish from — configuring from which servers a document is published

Click the **Publish From** tab and select the servers from which the document should be published. Please note that a server must have the application installed that can open this particular document type.

Application — configuring server-specific document settings

By default, the settings configured in the **Target** (application path), **Start In**, and **Parameters** fields apply to all servers a document is published from. If a document exists in a different folder on one (or more) of the servers, you can specify the above settings for a specific server or servers individually.

To do so:

- 1 Click the **Application** tab and.
- 2 Select a server in the **Server(s)** list.
- 3 Specify the **Target**, **Start In**, and **Parameters** (optional) properties. The values that you specify will apply to the selected server only. Repeat the steps for other servers if needed.
- 4 Click the **Verify Target(s)** button to verify the document path on all servers from which this application is published. The results are displayed in the **Target Verifier** dialog where you can see whether the target is correct or not for each server.

Filtering

Please see **Using filtering rules** (p. 247).

Routing

Please see **Configuring preferred routing** (p. 250).

Shortcuts — configuring shortcut options for a published document

Click the **Shortcuts** tab to enable the creation of shortcuts on the user desktops, shortcuts in the **Start** folder and shortcut in the **Auto Start** folder. When the **Auto Start** shortcut is enabled, the application will start when the user's computer is started.

Note: This option is not available on all operating systems.

File extension — configuring file extension associations

To modify file extension association for a particular published document, click the **File Extensions** tab. To add a new extension to the list, click **Tasks > Add** and specify the extension. To modify the extension parameters, highlight the extension and click **Tasks > Properties**.

Licensing — configuring licensing options for published documents

Click the **Licensing** tab to configure any of the below licensing options:

Select the **Inherit default settings** option to use the defaults. To specify your own settings, clear the option and set the following options:

- **Disable session sharing.** If this option is enabled, it allows you to isolate a given published application to one session. If the same application is launched more than once, the instances of the application will share the same sessions. A different application, however, will start in its own session.
- **Allow users to start only one instance of the application.** If this option is enabled, a user can only launch a single instance of the application.
- **Concurrent Licenses.** Use this option to specify the maximum number of concurrent instances the application can run. E.g. if the license of the application allows you to only run 10 instances of the application, set the Concurrent licenses option to 10 so once such limit is reached, other users cannot initiate other instances.
- **If limit is exceeded.** From this drop down menu you can specify what action should the Parallels RAS take in case any of the above licensing configured limits has been exceeded.

To use Site default settings, select the **Inherit default settings** option. See **Site defaults (Publishing)** for more info (p. 245).

Display — configuring display settings for a published document

Click the **Display** tab to configure the color depth of the published document, resolution, width and height. If these options are left at their default values, the client-specified options will take over.

You can also enable the option to wait for the Universal Printers to be redirected before the application is loaded. When enabling this option, you can also configure the maximum wait time (in seconds) for the Universal Printers to be redirected. To use Site default settings, select the **Inherit default settings** option. See **Site defaults (Publishing)** for more info (p. 245).

Manage folders

Folders are used to organize published resources and to facilitate filtering options.

There are two types of folders that you can create in the **Published Resources** tree in the Parallels RAS Console:

- **Folders for administrative purposes.** Folders of this type are intended for Parallels RAS administrators (users of the Parallels RAS Console). They are used to logically organize published resources in the Parallels RAS Console but they do not appear in the Parallels Client launchpad on user devices. These folders are used to help administrators manage published resources more efficiently.

- **Regular folders.** These folders are similar to administrative folders described above but they do appear in the launchpad on user devices. You normally use these folders to group published resources by type (e.g. office applications, specific business applications, utilities, etc.).

Creating a folder

To create a new folder:

- 1 In the RAS Console, select the **Publishing** category.
- 2 Right-click anywhere in the **Published Resources** tree and choose **New Folder** (or click the **[+] New Folder** icon at the bottom).
- 3 In the **Folder** dialog, specify a folder name and an optional description.
- 4 To make it a folder for administrative purposes, select the **Use for administrative purposes** option. To publish a regular folder, clear the option. See above for the explanation of the two folder types.
- 5 When creating a regular folder, you can change its icon by clicking the **Change icon** button. Administrative folders use a built-in icon that cannot be changed. Icons appear in the **Publishing** category in the Parallels RAS Console and in the Parallels Client launchpad (regular folders only).
- 6 On the next page, specify the initial status of the resource (the folder). Choose from the following options:
 - **Enabled:** End users can see the folder and will be able to launch published resources that it contains.
 - **Disabled:** The folder will not appear in Parallels Client.
 - **In maintenance:** The folder will appear in Parallels Client but users will not be able to launch resources that it contains. If the folder has subfolders, they inherit the status of the parent folder, which means that none of the resources contains in any of the folders in the hierarchy will be accessible to users. When a folder is in maintenance and a user tries to launch a resource from it, they will see a message. To customize the message, click the **Configure** button. For more info, see **Site defaults (Publishing)** (p. 245).
- 7 Click **Finish** to create the folder.

Managing folders

To modify an existing folder:

- 1 Select a desired folder in the **Published Resources** tree.
- 2 The **Information tab** in the right pane displays the folder information (read-only).
- 3 On the **Folder** tab, you can see and modify the folder name and description. You can also select or clear the **Use for administrative purposes** option to change the folder type (see above for an explanation). To change the folder icon, click the **Change icon** button. Note that the button is disabled if the **Use for administrative purposes** option is selected.

- 4 The **Filtering** tab specifies filtering options. Once set, these options will be inherited by all published resources in that folder. For more information, please see **Using filtering rules** (p. 247).
- 5 For the information about **Routing**, please see **Configuring preferred routing** (p. 250).

Adding published resources to a folder

To add a published resource to a folder, first add it to the root location and then drag it to the desired folder.

Delegating permissions to custom administrators

If you have custom administrators in your Farm, you can delegate permissions to them to manage a folder. This is specifically useful when a power administrator needs to grant permissions to a custom admin (they couldn't otherwise do it because they cannot manage user account directly). To grant folder rights, right-click anywhere in the **Published Resources** pane and then click **Delegate Permissions**. In the dialog that opens, select a user to grant folder permissions to. In the lower right pane of the **Delegate Permission - Publishing** dialog, select permissions (view, modify, add, delete) for a desired folder you want the user to have. Note that the custom administrator will be granted permissions to manage the folder and all its child items, including sub-folders. For more information about custom administrators, see **Managing administrator accounts** (p. 50).

Site defaults (Publishing)

The **Default Settings** dialog allows you to view and modify Site default settings for publishing. Published resources can inherit the following groups of settings from Site defaults:

- Shortcuts
- License
- Display
- Maintenance

To open the **Default Settings** dialog, navigate to **Farm > Site**. Click the **Tasks** menu and choose **Site defaults > Published resources**. The dialog consists of tabs, which are described below.

Shortcuts

In this tab specify whether and how the published resource shortcuts should be created on the user's computer. The following options are available:

- **Create shortcut on Desktop.** If selected, a shortcut will be created on the user's desktop.
- **Create shortcut in Start folder.** If selected, a shortcut will be added to the **Start** folder. You can specify the target subfolder name and path in the field provided. The default (and only) %Groups% variable will add additional subfolders as they appear on the host server where the published resource is hosted. For example, if the resource is located in "Myapps > Games" on the host server, the same folder structure will be added to the path. Note that you cannot use any custom variables.
- **Create shortcut in Auto Start folder.** If selected, the published resource will start automatically on computer startup.

License

The **License** tab contains the following options:

- **Disable session sharing.** If this option is enabled, it allows you to isolate a given published application to one session. If the same application is launched more than once, the instances of the application will share the same sessions. A different application, however, will start in its own session.
- **Allow users to start only one instance of the application.** If this option is enabled, a user can only launch a single instance of the published resource.
- **Concurrent Licenses.** Use this option to specify the maximum number of concurrent instances the published resource can run. For example, if the license of the application allows you to only run 10 instances of the application, set the **Concurrent licenses** option to 10, so once this limit is reached, other users cannot initiate other instances.
- **If limit is exceeded.** Specifies which action should Parallels RAS take in case any of the licensing limits configured above are exceeded.

Display

The **Display** tab contains the following options:

- **Wait until all RAS Universal Printers are redirected before showing the application.** Enable this option to wait for printers to be redirected before the application is loaded. You can also specify the maximum wait time (in seconds) for the Universal Printers to be redirected. Please note that redirecting a printer may take some time. To avoid confusion, a progress bar is shown to the user while the printers are being redirected.
- **Color Depth, Resolution, Width, Height.** These options specify the desired display settings for the application.

- **Start the application as maximized when using mobile clients.** This option applies only to Parallels Client running on mobile devices. When the option is selected, the application will start on a mobile device in the maximized state. This gives users the best experience while working with a remote application. This option gives the RAS administrator an easy way to always maximize an application without taking any additional steps.

Maintenance

The **Maintenance** tab allows you to specify a message that users will see when trying to launch a published resource in maintenance. When a resource is in maintenance, it will still appear in Parallels Client, but will be grayed out (in HTML5 Client, it will say so in the resource name). If a user tries to open the resource, they will see the message that you specify here. If you modified a message, but want to return the default one, select a message in the desired language and click **Tasks > Reset to default**. To reset messages in all languages, click **Tasks > Reset all to default**.

You can replicate the Site settings described above to other sites in your Parallels RAS Farm. To do so, select the **Replicate settings** option in a desired tab. All settings contained in the tab will be replicated.

Using filtering rules

Filtering is a feature that allows you to control who can access a particular published resource. You can define filtering rules based on any of the following:

- User
- Client device name
- Client device operating system
- IP address
- MAC address
- Gateway

By default, no filtering rules exist for a published resource, therefore the resource is available to anyone who is connected to the Parallels RAS Farm. Once you specify a filtering rule for a published resource, only those users/computers who satisfy the criteria will be able to use it.

To create a filtering rule, select a published resource in the **Published Resources** tree and click the **Filtering** tab. In the **Select Filtering Type** drop-down list, select criteria and then define a filtering rule as described below.

Filtering by user

To allow individual users or a user group to access the published resource:

- 1 Select **User** in the **Search Filtering Type** drop down list.
- 2 Select the **Allow the following Users** option.
- 3 Click **Tasks > Add** and specify a user or a group in the **Select Users** dialog. Click **OK** to add a user/group to the list on the **Filtering** tab.
- 4 In the **Default Object Type** drop-down list, select whether this rule will applies to users, groups, or both.
- 5 In the **Browse Mode** drop-down list, select the browsing mode you would like to use to connect to Active Directory or Windows.

The options are:

- **WinNT.** WinNT is faster than LDAP but does not support group nesting. Used only for backward compatibility.
- **LDAP.** LDAP supports group nesting but is slow. Used only for backward compatibility.
- **Secure Identifier.** This is the preferred and fastest method. It supports group nesting and renaming.

To convert users or groups specified using WinNT or LDAP, select a user entry and then click **Tasks > Convert**.

Filtering by client device name

To allow a specific client device or a list of client devices to access the published resource, follow these steps:

- 1 Select **Client** device name in the **Search Filtering Type** drop-down list.
- 2 Select the **Allow the following Clients** option. You can use the asterisk character (*) as a wildcard in a name. To include a wildcard in a name, select a client in the list and then click **Tasks > Edit**.
- 3 Click **Tasks** and choose one of the following:
 - **Add from network browse.** Opens a dialog where you can select a client from the list populated from the network.
 - **Add from Active Directory.** Opens a dialog where you can specify a computer or search the Active Directory for it.
 - **Add from known devices.** Opens a dialog where you can select a client from the list populated by previously connected clients.
 - **Add custom entry.** Allows you to type the name of a client. To modify the name, select it and then click **Tasks > Edit**.

- **Edit.** Allows you to modify the name of a selected client. If you want to include a wildcard (*) in a name, you can do it using this option. If no client is selected in the list, the option is disabled.
- **Import from CSV.** Allows you to select a CSV file containing the list of names of client devices. The file should contain a single device name on each row. The names must be unique (no duplicates) or you will see an error message.
- **Export to CSV.** Allows you to export the list of client device names to a CSV file.
- **Delete.** Allows you to delete a selected client. If no client is selected in the list, the option is disabled.

4 Click **OK** to add your selection to the **Client** list.

Filtering by client device operating system

To allow client devices running a particular operating system to access the published resource, follow these steps:

- 1 Select **Client device operating system** in the drop-down list.
- 2 Select the **Allow access to clients on the following operating system:** option to enable the filtering rule.
- 3 Select one or more operating systems.
- 4 Click **Apply** at the bottom of the RAS Console window to save the changes.

When using the **Checking effective access** (p. 253) functionality, the filtering rule information will be displayed as "Client device operating system filtering is enabled".

Filtering by IP address

To allow a specific IP address (or multiple addresses) or a range of IP addresses to access the published resource, follow these steps:

- 1 In the **Search Filtering Type** drop-down list, select **IP Address**.
- 2 Select the **Allow the following IPs** option.
- 3 Click **Tasks > Add** in the IPv4 and/or IPv6 sections to specify the IP address or a range of IP addresses and click **OK**.

Filtering by MAC address

To allow a MAC address or a specific list of MAC addresses to access the published resource, follow these steps:

- 1 In the **Select Filtering Type** drop-down list, select **MAC**.
- 2 Select the **Allow the following MACs** option.

3 Click **Tasks > Add** and choose one of the following:

- **Add.** Select clients to add to the list **OK**.
- **Import from CSV.** Select a CSV file containing the list of names of client devices. The file should contain a single MAC address on each row. The addresses must be unique (no duplicates) or you will see an error message.
- **Export to CSV.** Allows you to export the list of MAC addresses to a CSV file.

Filtering by gateway

To allow users to connect to a published resource through a specific gateway, follow these steps:

- 1** Select the **Gateway** filtering type.
- 2** Select the **Allow connections from the following gateway** option.
- 3** Click **Tasks > Add** to specify the gateway and its IP address (if it has multiple IP addresses).

Configuring multiple filtering rules

If multiple filtering rules are configured for a specific published resource, the connecting user has to match ALL of them to be allowed access to the published resource.

Please note that if you applied multiple filters, all of them will be visible in the **Information** tab of a published item.

Configuring preferred routing

Overview

Preferred routing is a useful feature when Parallels RAS users with different geo located deployments are connecting to the same Parallels RAS Farm/Site. A common access layer usage (RAS Secure Client Gateway, HALB, or a third-party load balancer) is not optimal if a resource is located in a different data center in the same RAS Farm/Site. The solution is to configure a preferred access layer server for a specific published resource, in which case any user would connect to a default Secure Client Gateway, but would be redirected using proximity rules set by the administrator. Typically, using the Secure Client Gateway closest to the session host provides improved user experience, reduced internal network traffic and associated costs along with providing better use of resources.

Note: Preferred routing doesn't apply to Azure Virtual Desktop published objects.

Here's how preferred routing works:

- 1** Parallels Client establishes a connection with a Secure Client Gateway using a standard authentication.

- 2 Through the RAS Publishing Agent, the resource's preferred route (if configured) is identified.
- 3 Parallels Client receives the preferred public address to launch the resource.
- 4 Parallels Client then tries to launch the resource through the redirected address and falls back to the original Gateway if it fails.

Configure preferred routing

To configure preferred routing for a published resource:

- 1 In the RAS Console, select the **Publishing** category.
- 2 Select an existing published resource and then select the **Routing** tab.
- 3 Select the **Enable preferred routing** option.
- 4 Click **Tasks > Add**. The **Add preferred route** dialog opens. Read on.

In the **Add preferred route** dialog:

- 1 Type a name for this route and an optional description.
- 2 From the **Type** drop-down list, select one of the following:
 - **HALB Virtual Server**: Select a RAS HALB virtual server from the list below the **Type** field. Note that for a RAS HALB virtual server to appear in the list, the HALB server must have a public address specified, as far as you cannot add HALB server here.
 - **Gateway**: Same as for HALB virtual server (described above), the **Public address** field must have a value for the Gateway to appear in the list. See the **Public address** field when you create or configure a RAS Secure Client Gateway.
 - **Custom**: A third-party load balancer will be used. Select this option, then click **Tasks > Add** and specify the server properties in the list below the field. The properties include **Name**, **Description**, **Public address**, **Port**, and **SSL port**. You can add as many servers as required and then select one of them to be used for a given published resource.

When configuring preferred routing, please also consider the following:

- If routing fails, an automatic fallback to the originating address is carried out.
- If routing is enabled in the RAS Console, but not configured, the administrator will see an error message and will have to either configure or disable it.
- It is recommended to use Folders (configured for administrative purpose) in case routing needs to be configured on many resources. Routes are not inherited from parent folder if the routing is enabled on the child object, i.e. only one set of routes is used.
- Same user credentials will be used when redirecting the RDP traffic to the same RAS Site. Users will not be asked to re-enter credentials.
- If you had existing session tunneled in a particular Secure Client Gateway and using session sharing, the same session workflow path will be used (if a published resource is also available on the same session host) irrespective of configured routing.

- Routing is supported in SAML environments.
- Supported clients are Windows, macOS, Linux, Android, iOS, HTML5.

Specifying public address when inviting users

When you use the **Invite Users** wizard, you can specify a public address on the second page where you specify target platforms and connection options. This way, you can set a preferred routing for a group of users in a specific geo location. For more information, see **Invite users** (p. 37).

Deleting or disabling Gateway or HALBs

If an administrator tries to delete a Gateway or HALB when they are being used as preferred routing, the information about objects using it will be shown on the screen, so no accidental deletion takes place.

Understanding session prelaunch

When a user opens a remote application or desktop, a session must first be launched. Launching a session can take time, which will result in the user waiting for the application to start. To improve user experience, a session can be launched ahead of time, before the user actually opens an application. When a session is prelaunched, it will all happen in the background, so the user will not see any windows or message boxes on the screen. When the user starts an application, it will open using the prelaunched session, so it will start very quickly.

For the information about how to configure session prelaunch, see **Client policies** (p. 389).

When you configure session prelaunch, the following options are available:

- **Off.** No session prelaunch is used.
- **Basic.** A session is prelaunched as soon as the user gets the application listing. The assumption is, the user will open an application within the next few minutes. The session will stay active for 10 minutes. If the user doesn't open an application during that time, the client will disconnect from the session.
- **Machine Learning.** When the application listing is acquired, a session is prelaunched based on user habits. With this option enabled, Parallels Client will record and analyze the user habits of launching applications on a given day of the week. A session is started a few minutes before the user usually opens an application.

You can configure rules when session prelaunch must not be used. The following options are available:

- Specify dates on which the prelaunch must not be used.
- Exclude a published resource from the session prelaunch decision making. If a resource is excluded from the analysis, it is never considered by Parallels Client when making a decision whether to prelaunch a session. For example, when you have a server on which you never want to prelaunch sessions, you can flag all published resources hosted by that server as to be excluded from session prelaunch. To exclude a published resource from session prelaunch, in the RAS Console, navigate to **Published Resources**, select a resource and then select the **Exclude from session prelaunch** option.

Checking effective access

Filtering rules described in the previous section (p. 247) allow you to configure who can access a particular published resource. If a Parallels RAS user cannot see one or more published resources in Parallels Client, you would normally have to check filtering settings for each resource to make sure that it is published for a given user. The Effective Access functionality simplifies this task by allowing you to view in one place which published resources are available for a user and which are not.

To open the **Effective Access** dialog, select the **Publishing** category in the Parallels RAS Console and then click the **Effective Access** item in the toolbar at the bottom of the window (if you don't see the item, maximize the console window). You can also open the dialog by right-clicking anywhere in the **Published Resources** pane and choosing **Effective Access** in the context menu.

The **Effective Access** dialog allows you to specify a user (and optionally additional criteria) and then view published resources this user is allowed to access. To choose a user, do one of the following:

- Type the user name in the **User** field, or click the [...] button next to it and use the **Select User or Group** dialog to select a user.
- Select a device owned by this user from the list of known devices. To do so, click the **Select a Device** button then select a device. Note that if a device has never been used to connect to this Parallels RAS Farm, it will not be included in the list. For more information, see the **Monitoring devices** section. (p. 375) After selecting a device, click **OK** to return to the **Effective Access** dialog. All of the fields will be automatically populated using properties of the selected device.

Once you specify a user, enter the additional criteria if needed (all fields except **User** are optional):

- **Client.** Client name assigned to a device. This could be a computer name, FQDN, or a custom name that the user could have set in Parallels Client.
- **IP Address.** Client IP address.

- **MAC.** Client MAC address.
- **Gateway.** RAS Secure Client Gateway name through which the client connects to the Farm.

The **Manage groups** button allows you to preview how user access changes if the user is added to one or more groups. When you click the button:

- 1** The **Manage Groups** dialog opens listing groups to which the user already belongs.
- 2** Click the **[+]** button to add the user to one or more additional groups. Note that this will only be a simulation; the user will not be actually added to any additional group.
- 3** To remove a "simulated" group, select it in the lower pane and click the **[-]** button.
- 4** Click **Close** to return to the **Effective Access** dialog.

Finally, to view the effective access information for the specified user, click the **View** button. This opens the **Effective Access - Summary** dialog, which displays the following information:

- The left pane contains the complete list of resources published in the current Site. To view only the resources that the specified user can access, select the **Show only allowed published resources** option. If the user is not allowed to access a resource, the resource name is highlighted in red.
- The right pane contains information whether the user is allowed to access a resource selected in the left pane and whether filtering is enabled for the selected resource. Additional information may include filtering details and extended group membership.

By looking through the resource list, you can see which resources the user can or cannot access and take appropriate actions if necessary. If needed, you can export the effective access information to a CSV file. To do so, click the **Export** button and specify a file name. The CSV file has the following columns:

- **Name.** Application name.
- **ID.** Application ID.
- **Accessible.** Whether the application is accessible to the user (Yes or No).
- **Rule.** Filtering rule. If no rules are configured for the application, the column will have no value.

Specifying client settings

To specify client settings for published resources, navigate to **Farm** > <Site> > **Settings** and select the **Client Settings** tab. On this page, you can specify how published application icons are displayed on the client side and some other options.

Select icon resolution

Published resources are displayed in Parallels Client as icons or as a list. You can specify which resolution should be used when the resources are displayed as icons. Select from the following options:

- **Send standard resolution icons.** Standard resolution icons.
- **Send high resolution icons.** High resolution icons. Please note that this option will use more network bandwidth.

Enable, disable or change the overlay icon

Note: These configuration changes are applicable to desktop clients only (Windows, Mac, Linux). They have no effect on mobile and HTML5 clients.

The other options on this tab is to enable, disable or update the overlay icon. An overlay icon is placed over a standard application icon to indicate that it's a remote application served by Parallels RAS. When you launch a remote application from Parallels Client, the application icon is displayed on the local desktop (e.g. on the taskbar in Windows or Dock in macOS). By using an overlay icon, you give the user the ability to tell at a glance which of their running applications are remote Parallels RAS applications and which are local (or any other kind).

Parallels RAS uses the Parallels logo as the overlay icon by default. However, the administrator may also change this to use the standard Microsoft RemoteApp overlay icon. When using Parallels logo as the overlay icon, an application icon on a local computer will look like the following sample icons:



As you can see, these are standard icons used by the Windows Calculator and Paint applications with the Parallels logo icon (red parallel lines in the right corner) displayed on top of them. When a user notices the overlay, they'll know right away that this is a remote application served by Parallels RAS, not a local Windows app.

Show password expiration reminder

You can automatically remind your Parallels RAS users to change their domain password when it nears the expiration date. To enable this functionality, select the **Show password expiration reminder** option. When it is enabled, a Parallels Client user whose password is about to expire will see a notification right after they connect to Parallels RAS. Note that the option is disabled by default.

Quick keypad

The **Quick Keypad** category in the Parallels RAS Console allows you to define custom keys to perform common actions in published applications running on mobile devices. Custom keys appear above the standard keyboard in iOS and Android and can be tapped just like any other key on the virtual keyboard.

This feature is designed for users who run published applications on a phone or a tablet. When a particular software requires repeated selection of certain menu or toolbar items, using custom keys can significantly improve user experience. For example, let's say a user has some data entry task which requires them to press **File > New** and **File > Save** menu items over and over again. If you define two custom keys to perform these actions, the user will see them above the standard keyboard in iOS or Android, so instead of tapping the application's native menu items (which can be cumbersome), they can tap these keys, which is much easier and quicker.

To define custom keys, select the **Quick Keypad** category in the Parallels RAS Console. The **Quick Keypads** view in the right pane allows you to create a Quick Keypad template. A template is created for a specific application (or a group of applications with the identical UI design) and contains shortcuts to perform common actions in an application. Once a template is created, you assign it to a published application or a group of applications, so each application (or a group) has its own Quick Keypad.

To create a Quick Keypad template:

- 1 Click the **Tasks** drop-down menu and choose **New Quick Keypad** (or click the **[+]** icon).
- 2 Specify a Quick Keypad template name (e.g. "Office apps").
- 3 You can organize a Quick Keypad using a multi-level menu system. If you want to do this, click the **New menu** item and specify the menu item name. You can add sub-menu items too. To move a menu item across the tree, simply drag and drop it to the desired tree node.
- 4 When you have your basic menu structure defined, you can add shortcuts (or you can do it any order you like).
- 5 To add a shortcut, click the **New shortcut** item.
- 6 In the **Label** field, enter the name (e.g. "New").
- 7 Click the **Shortcut** field and press a shortcut on the keyboard as you would in the target application. For example, the standard shortcut to create a new document in many applications is Ctrl+N, so to input this shortcut, you would press and hold Ctrl and then press N. The shortcut will appear in the field as "Ctrl+N". You can input up to three shortcuts in this field.
- 8 To add another shortcut to the template, click the **New shortcut** item again. Repeat until all desired shortcuts are defined.
- 9 Click **OK** to close the dialog. The new template will appear in the **Quick Keypads** list.

To modify the template, right-click it and choose **Properties**.

You now need to assign the template that you created to an application (or multiple applications). To do so:

- 1** Right-click a template and choose **Assign to Application** (you can also use the **Tasks** drop-down menu or click the "link" icon).
- 2** In the **Assign Quick Keypad Template** dialog, select one or more applications to which the template should be assigned.
- 3** Click **OK** when done.

When a remote user runs an application on their mobile device and opens a virtual keyboard, they will see the extra keys corresponding to shortcuts that you defined for a Quick Keypad template. Tapping a key will perform the corresponding action (e.g. Ctrl-N, which will open a new document).

Exporting and importing a Quick Keypad template

To easily move a Quick Keypad template from one Parallels RAS Farm to another, use the Import and Export functionality. To export a template, right-click a template and choose **Export**. Specify the file name and location and click **Save**. To import a template, right-click on an empty space in the **Quick Keypads** list and choose **Import**. You can also perform these actions using the **Tasks** drop-down menu.

SSL Certificate Management

The Parallels RAS Console includes a certificate management interface that allows you to manage all of your SSL certificates in one place.

Certificates are managed on a Site level. Once a certificate is added to a Site, it can be used with any RAS Secure Client Gateway or HALB that also exist in this Site.

To manage certificates, in the RAS Console, navigate to **Farm > Site > Certificates**. The **Certificates** tab in the right pane displays the existing certificates. When you install Parallels RAS, the <Default> self-signed certificate is created automatically, so you will see at least this certificate in the list. The default certificate is also automatically assigned to all new RAS Secure Client Gateways and HALB.

You can perform the following certificate management tasks in the **Certificates** sub-category:

- Generate a self-signed certificate (p. 259)
- Generate a certificate signing request (CSR) (p. 260)
- Import a certificate from file (p. 260)
- Export a certificate to a file (p. 261)
- Assigning a certificate to Gateways and HALB (p. 261)
- Audit certificates (p. 263)
- Set permissions to manage certificates (p. 263)

The subsequent sections describe certificate management tasks in detail and provide additional certificate information and instructions.

In This Chapter

Generating a self-signed certificate	259
Generating a certificate signing request (CSR)	260
Importing a certificate	260
Exporting a certificate	261
Assigning a certificate to Gateways and HALB	261
Auditing certificates	263
Permissions to manage certificates	263
Upgrading from an older RAS version.....	264

Generating a self-signed certificate

To generate a self-signed certificate, navigate to **Farm > Site > Certificates**. Click **Tasks > Generate self-signed certificate**. In the dialog that opens, specify the following options:

- **Name:** Type a name for this certificate. This field is mandatory.
- **Description:** An optional description.
- **Usage:** Specify whether the certificate should be used for RAS Secure Client Gateways or HALB, or both. This selection is mandatory.
- **Key size:** The certificate key size, in bits. Here you can select from the predefined values. The default is 2048 bit, which is the minimum required length according to current industry standards.
- **Country code:** Select your country.
- **Expire in:** The certificate expiration date.
- **Full state or province:** Your state or province info.
- **City:** City name.
- **Organization:** The name of your organization.
- **Organization unit:** Organizational unit.
- **E-mail:** Your email address. This field is mandatory.
- **Common name:** The Common Name (CN), also known as the Fully Qualified Domain Name (FQDN). This field is mandatory.
- **Alternative names:** Specify one or more subject alternative names (SANs). Click the [...] icon and then add one or more DNS or IP addresses. Note that because Parallels Client for mobile devices doesn't support the SAN field, it's safest to set your common name to the name that most mobile devices will be using.

Click **Save** to generate the certificate. When done, the certificate will appear in the **Certificates** list in the RAS Console with the **Status** column indicating **Self-signed**.

To view the certificate info, right-click it and choose **Properties**. In the dialog that opens, examine the properties and then click the **View certificate info** button to view the certificate trust information, details, certification path and the certificate status. You can also view the certificate info by right-clicking it and choosing **View certificate info**.

Generating a certificate signing request (CSR)

To generate a CSR, navigate to **Farm > Site > Certificates**. Click **Tasks > Generate a certificate request**. In the dialog that opens, specify the required information. The information is exactly the same as for the self-signed certificate described above (p. 259). If you need an explanation, please refer to the list of options described in that section.

After entering the information, click **Generate**. Another dialog will open displaying the request. Copy and paste the request into a text editor and save the file for your records. The dialog also allows you to import a public key at this time. You can submit the request to a certificate authority now, obtain the public key, and import it without closing the dialog, or you can do it later. If you close the dialog, the certificate will appear in the RAS Console with the **Status** column indicating **Requested**.

To submit the request to a certificate authority and import a public key:

- 1 If the certificate request **Properties** dialog is closed, open it by right-clicking a certificate and choosing **Properties**. In the dialog, select the **Request** tab.
- 2 Copy the request and paste it into the certificate authority web page (or email it, in which case you will need to come back to this dialog later).
- 3 Obtain the certificate file from the certificate authority.
- 4 Click the **Import public key** button and finalize the certificate registration by specifying the key file and the certificate file.

Importing a certificate

To import a certificate from a file, on the **Certificates** tab, click **Tasks > Import certificate**. In the dialog that opens, specify the following:

- **Name:** Type a name for the certificate.
- **Description:** An optional description.
- **Private key file:** Specify a file containing the private key. Click the [...] button to browse for the file.
- **Certificate file:** When you specify a private key file (above) and have a matching certificate file, it will be inserted in this field automatically. Otherwise, specify a certificate file.
- **Usage:** Specify whether the certificate will be used for RAS Secure Client Gateways or HALB, or both.

Click **OK** when done. The certificate will appear in the list in the RAS Console with the **Status** column indicating **Imported**.

To view the certificate info, right-click it and choose **Properties**. In the dialog that opens, examine the properties and then click the **View certificate info** button to view the certificate trust information, details, certification path and the certificate status. You can also view the certificate info by right-clicking it and choosing **View certificate info**.

For imported certificates, the **Properties** dialog has an additional tab **Intermediate**. If the original certificate included an intermediate certificate (in addition to the root certificate), it will be displayed here. You can paste a different intermediate certificate here if you wish.

Exporting a certificate

To export a certificate to a file, on the **Certificates** tab, click **Tasks** > Export certificate, specify a filename and click **Save**. You can later import the certificate in a different Farm or Site by clicking **Tasks** > **Import certificate** and specifying the certificate file in the **Private key file** field.

Assigning a certificate to Gateways and HALB

After you add a certificate to a Site, you can assign it to a RAS Secure Client Gateway, HALB, or both depending on the usage type that you specified when you created the certificate (described in the beginning of this chapter). More on the certificate **Usage** option below.

Certificate usage

Certificate **Usage** is an option that you specify when you create a certificate. It specifies whether the certificate should be available for RAS Secure Client Gateways, HALB, or both. When setting this option, you can choose from the following:

- **Gateway:** If selected, makes the certificate available for RAS Secure Client Gateways.
- **HALB:** If selected, makes the certificate available for HALB.

You can select one of the options above or both, in which case the certificate becomes available for both, Gateways and HALB. For details on how to create a certificate and choose these options, please see *Generating a self-signed certificate* (p. 259) and *Generating a certificate signing request (CSR)* (p. 260).

When you configure SSL for a RAS Secure Client Gateway or HALB later, you need to specify an SSL certificate. For the information on how to do this, please see SSL/TLS encryption (p. 70) and Configuring HALB in the RAS Console (p. 292). When you select a certificate, the following options will be available depending on how the **Usage** option is configured for a particular certificate:

- **<All matching usage>**: This is the default option, which is always available. It means that any certificate on which the **Usage** selection matches the object type (Gateway or HALB) will be used. For example, if you are configuring a Gateway and have a certificate that has **Usage** set to "Gateway", it will be used. If a certificate has both, Gateway and HALB usage options selected, it can also be used with the given gateway. This works the same way for HALB when you configure the LB SSL Payload. Please note that if you select this option for a Gateway or HALB, but not a single matching certificate exists, you will see a warning and will have to create a certificate first.
- Other items in the **Certificates** drop-down list are individual certificates, which will or will not be present depending on the certificate's **Usage** settings. For example, if you configure LB SSL Payload for HALB and have a certificate with the **Usage** option set to "HALB", the certificate will appear in the drop-down list. On the other hand, certificates with **Usage** set to "Gateway" will not be listed.

As another example, if you need just one certificate, which you would like to use for all of your Gateways, you need to create a certificate and set the **Usage** option to "Gateways". You can then configure each Gateway to use this specific certificate or you can keep the default **<All matching usage>** selection, in which case the certificate will be picked up by a Gateway automatically. Same exact scenario also works for HALB.

Gateways

To assign a certificate to a RAS Secure Client Gateway:

- 1 Navigate to **Farm > Site > Gateways**.
- 2 Right-click a gateway and choose **Properties**.
- 3 Select the **SSL/TLS** tab.
- 4 In the **Certificates** drop-down list, select the certificate that you created.
- 5 Click **OK**.

Please note that you can also select the **<All matching usage>** option, which will use any certificate that either has the usage set to Gateway or both Gateway and HALB.

HALB

To assign a certificate to a HALB, navigate to **Farm > Site > HALB**. Assuming that your HALB is enabled and configured, and the **LB SSL Payload** option is selected, follow the instructions below:

- 1 Click **Configure** next to the **LB SSL Payload** option.
- 2 A certificate must be used when the **Mode** option is set to **SSL Offloading**. Once again, assuming it is selected, continue to the next step.
- 3 Click **Configure**.
- 4 In the **SSL** dialog, select the certificate in the **Certificates** drop-down list.

As with gateways, you can also select the **<All matching usage>** option, which will use any certificate that has the usage set to HALB or both HALB and Gateway.

Auditing certificates

All actions that you perform on certificates are audited and can be viewed later. Note that reverting certificate changes is not possible. If you need to revert to a previous state, you'll have to delete a certificate and create a new one.

To audit certificates:

- 1 In the RAS Console, navigate to **Farm > Site > Certificates**.
- 2 Click **Tasks > Settings audit**.
- 3 The dialog opens where you can view the history of certificate actions. Note that the **Revert** button is disabled. As noted at the beginning of this section, reverting a certificate action is not possible.
- 4 To view details for a particular audit entry, double-click it.

Permissions to manage certificates

Root and Power administrators always have rights to manage certificates. Custom administrators don't have them by default. To grant permissions to manage certificates to Power administrators, the **Certificates** global permission type is used.

If you are a Root or Power administrator, you can set certificate permissions as follows:

- 1 In the RAS Console, navigate to **Administration > Accounts**.
- 2 Select a Custom administrator account and click **Tasks > Properties**.
- 3 In the **Account Properties** dialog, click **Change Permissions**.

- 4 In the **Account Permissions** dialog, select a Site in the left pane and click **Change permissions** (or click the **Edit** link in the right pane).
- 5 In the left pane (Permission type), select **Certificates**.
- 6 In the right pane (Global permissions), select one or more permissions.
- 7 When done, close all dialogs.

A RAS administrator can also delegate his/her permissions to a custom administrator. To do so, navigate to **Farm > Site > Certificates** and click **Tasks > Delegate permissions**. In the dialog that opens, delegate permissions to a desired Custom administrator.

Upgrading from an older RAS version

When you upgrade Parallels RAS from a version prior to RAS 17.1 to a RAS 17.1 (or newer), every certificate that is used by RAS Secure Client Gateways and HALB is enumerated and only unique certificates are added to the **Certificates** subcategory. Gateways and HALB are then linked 1-to-1 to the certificates they were using before the upgrade.

Other actions related to an upgrade include the following:

- The **Inherit defaults** option in gateways is turned off after the upgrade.
- If a gateway is disabled during an upgrade, the Publishing Agent still has the information about the certificate that the gateway uses, so the gateway is configured properly when it comes back online.
- Site defaults settings are configured to use the default self-signed certificate.
- When a new gateway is added, it is configured to use the default self-signed certificate, provided the Site defaults are not changed afterwards.

Connection and Authentication Settings

A Parallels RAS administrator has the ability to customize how users connect to Parallels RAS. This chapter describes connection and authentication settings that can be configured according to your organization requirements. It then explains how to use two-factor authentication for higher level of security.

In This Chapter

RAS Publishing Agent connection settings	265
Remote session settings	266
Restricting access by Parallels Client type and build number	268
Multi-factor authentication	269
Allowing users to change domain password.....	286

RAS Publishing Agent connection settings

RAS Publishing Agent connection settings can be accessed from the **Connection** category.

Choosing authentication type

Select the **Authentication** tab. In the **Allowed authentication types** section, select one of the following options:

- **Credentials.** The user credentials are validated by the Windows system on which RAS is running. The credentials used for Windows authentication are also used to log in to an RDP session.
- **Smart Card.** Smart card authentication. Similar to Windows authentication, smart card credentials can be shared between both RAS and RDP. Hence, smart card credentials only need to be entered once. Unlike Windows authentication, the user only needs to know the smart card's PIN. The username is obtained automatically from the smart card, so the user doesn't need to provide it.
- **Web (SAML).** SAML SSO authentication.

Note that if smart card authentication is disabled, RAS Publishing Agent will not hook the Local Security Authority Subsystem Service (LSASS). Smart card authentication can be used in Parallels Client for Windows, Mac, and Linux. Please also note that smart cards cannot be used for authentication if Parallels Client is running inside an RDP session.

Smart card certificate

A valid certificate must be installed on a user device in order to use smart cards. To do so, you need to import the certificate authority root certificate into the device's keystore.

A certificate must meet the following criteria:

- The "Key Usage" field must contain digital signature.
- The "Subject Alternative Name" (SAN) field must contain a user principal name (UPN).
- The "Enhanced Key Usage" field must contain smart card logon and client authentication.

Authentication domain

To specify an authentication domain, select one of the following:

- **Specific:** Select this option and type a specific domain name.
- **All trusted domains:** If the information about users connecting to Parallels RAS is stored in different domains within a forest, select the **All Trusted Domains** option to authenticate against multiple domains.
- **Use client domain if specified:** Select this option to use the domain specified in the Parallels Client connection properties. If no domain name is specified on the client side, the authentication is performed according to the settings above.
- **Force clients to use NetBIOS credentials:** If this option is selected, the Parallels Client will replace the username with the NetBIOS username.

Note: If a certificate on your smart card does not contain a user principal name (UPN) in the "Subject Alternative Name" (SAN) field (or if it doesn't have the "Subject Alternative Name" field at all) you have to disable the **Force clients to use NETBIOS credentials** option.

Recommendation: After changing the domain names or some other authentication related changes, click the **Clear cached session IDs** button on the **Settings** tab.

Authenticating against non domain users

In order to authenticate users sessions against users specified on a standalone machine you must enter the [workgroup_name] / [machine_name] instead of the domain name. For example if you would like to authenticate users against a list of local users on a machine called SERVER1 that is a member of the workgroup WORKGROUP, enter the following in the domain field: WORKGROUP/SERVER1.

Remote session settings

The **Settings** tab in the **Connection** category allows you to configure the following remote session options.

Declare user session as idle after

This option affects reporting statistics, whereby a session is declared idle after the amount of time specified without any activity.

FIPS 140-2 encryption

The **FIPS 140-2 encryption** property allows you to specify whether FIPS-encrypted connections are allowed or even enforced on RAS Secure Client Gateways. When you allow (or enforce) the encryption, the Gateways will use the FIPS 140-2 encryption module. You can choose from the following options:

- **Disabled.** FIPS 140-2 encryption is disabled on RAS Secure Client Gateways.
- **Allowed.** RAS Secure Client Gateways accept both FIPS-encrypted and non-FIPS-encrypted connections.
- **Enforced.** RAS Secure Client Gateways accept FIPS-encrypted connections and will drop any non-FIPS-encrypted connection.

Note: For FIPS 140-2 encryption to work, a FIPS compliant certificate must be installed on each RAS Secure Client Gateway.

When you enable FIPS 140-2 encryption, the encryption status is displayed on the **Information > Site** tab in the RAS Console. Look for the **Encryption** property of a RAS Secure Client Gateway.

The following versions of Parallels Client support FIPS 140-2 encryption:

- Parallels Client for Windows 64-bit
- Parallels Client for Linux 64-bit

Please note that HALB is not supported when using a FIPS-encrypted connection.

Please also note that when FIPS 140-2 encryption is enforced, it is enforced all users in a given Farm. If there's a necessity to force FIPS for one user group and not forced for another, a new Farm must be deployed for this purpose.

Automatically log out client idle connection after

Specifies the time period after which an idle client connection should be logged out. Once the connection is logged out, the user is disconnected from Parallels RAS and is presented with the **Connections** dialog in Parallels Client as a way to notify them that they were logged out. They can use the dialog to log back on if desired. Parallels Client connection is considered idle after the last user session has been disconnected or logged off.

Cached authentication token timeout

Specify the amount of time that a session is cached for (higher amount of time reduces AD transactions).

Clear cached authentication tokens (a button)

Clears all cached session information.

Restricting access by Parallels Client type and build number

You can specify a minimum requirement for the Parallels Client type and version number in order for it to connect to the Parallels RAS Farm or to list published resources. In addition, you can set the Parallels Client security patch level (described later in this section).

To specify Parallels Client requirements:

- 1 In the RAS Console, select the **Connection** category and click the **Allowed Devices** tab.
- 2 The **Allow only clients with latest security patches** option specifies the Parallels Client security patch level. If the option is selected, only clients with latest security patches applied will be allowed to connect to Parallels RAS. This option must normally be selected to protect your environment from vulnerabilities. You should only clear it if you must use an older version of Parallels Client with no security patches installed. For more information, please see the following KB article: <https://kb.parallels.com/en/125112>.
- 3 In the **Mode** drop-down list, select from the following options:
 - **Allow all clients to connect to the system.** No restrictions. All Parallels Client types and versions are allowed full access.
 - **Allow only the selected clients to connect to the system.** Allows you to specify Parallels Client types and versions that are allowed to connect to the Parallels RAS Farm. Select the desired Parallels Client types in the **Clients** list. To set the **Minimum build** value, right-click the client type and choose **Edit**. Type the version number directly in the **Minimum build** column.
 - **Allow only the selected clients to list the published items.** Allows you to specify Parallels Client types and versions that can list published resources. Compared to the option above, this one does not restrict Parallels Clients connecting to Parallels RAS. Select this option and then select the desired Parallels Client types in the **Clients** list. To set the **Minimum build** value, right-click the client type and then click **Edit** in the context menu. Type the version number directly in the **Minimum build** column.

If a restriction is configured and a Parallels Client is excluded from the list, the user running it will receive a corresponding error message and will be advised to contact the system administrator.

Multi-factor authentication

Parallels RAS allows you to use multi-factor authentication for access control. When multi-factor authentication is used, users will have to authenticate through two successive stages to get the application list. While the first level will always use native authentication (Active Directory / LDAP), the second level can use one of the following solutions:

- Azure MFA (RADIUS) (p. 269)
- Duo (RADIUS) (p. 269)
- FortiAuthenticator (RADIUS) (p. 269)
- TekRADIUS (p. 269)
- Deepnet
- SafeNet (p. 282)
- Google Authenticator (p. 283)

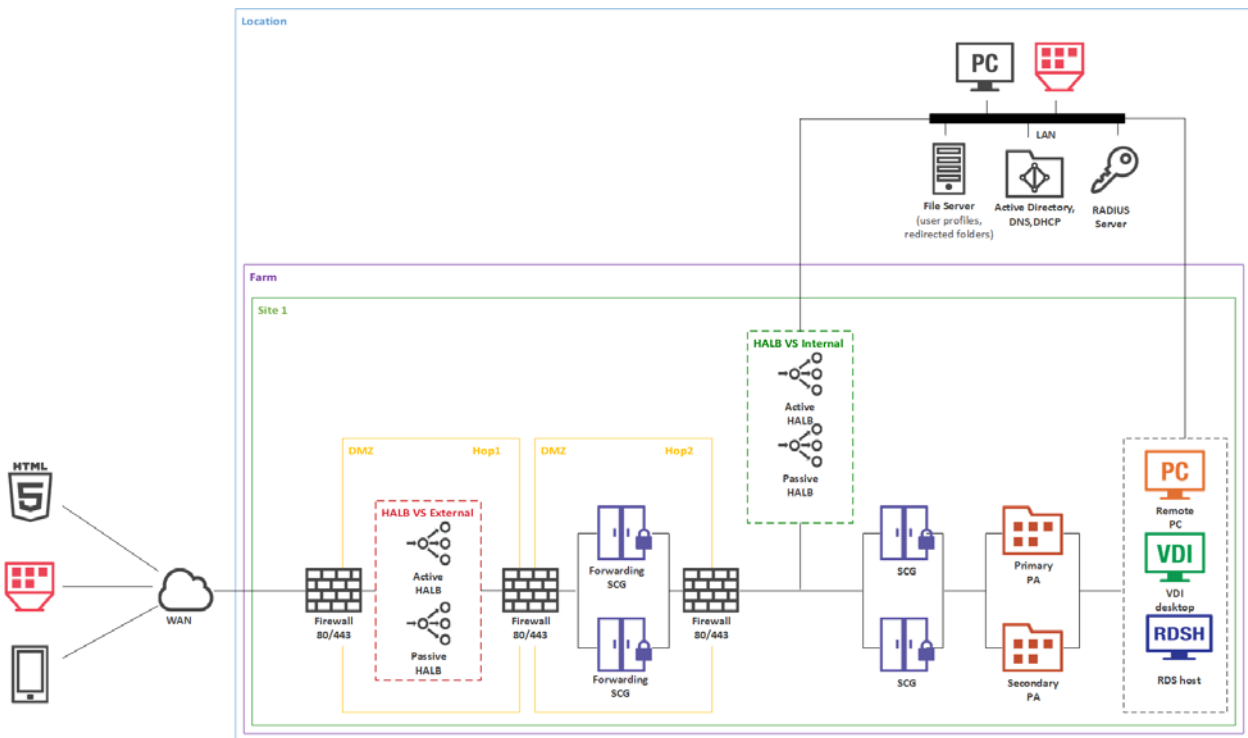
Multi-factor authentication is more secure because instead of using a standard user name and password, it uses a static user name and a one-time password generated by a token.

Multi-factor authentication can be configured in the Parallels RAS Console in **Connection > Multi-factor authentication**.

See also **Configuring MFA exceptions** (p. 285).

Using RADIUS

The below diagram shows the double hop perimeter network scenario with RAS Publishing Agent connected to a RADIUS server (RADIUS is located in Intranet but it can be placed in DMZ).



To configure RADIUS properties:

- 1 In the Parallels RAS Console, navigate to **Connection > Multi-factor authentication**.
- 2 In the **Provider** drop-down list, select a RADIUS MFA provider that you use in your organization. The following options are available:
 - Azure MFA server (RADIUS) (p. 273)
 - Duo (RADIUS) (p. 275)
 - FortiAuthenticator (RADIUS)
 - TekRADIUS
 - RADIUS
- 3 Click the **Settings** button. The **Properties** dialog opens where you can configure settings for the selected RADIUS provider.

Read on to learn how to configure RADIUS provider settings.

Connection

In the RADIUS **Properties** dialog, select the **Connection** tab and specify the following options:

- **Type Name:** Specify the name of the OTP connection type that will be displayed on the Logon screen on the client side. This should be the name that your users will clearly understand.

- **Primary server** and **Secondary server**: These two fields allow you to specify one or two RADIUS servers to include in the configuration. Specifying two servers gives you an option to configure high availability for RADIUS hosts (see below). Specify a server by entering its hostname or IP address or click the [...] button to select a server via Active Directory.
- When two RADIUS servers are specified, select one of the following high availability modes from the **HA mode** drop-down box: **Active-active (parallel)** means the command is sent to both servers simultaneously, the first to reply will be used; **Active-passive (failover)** means failover and timeout are doubled, Parallels RAS will wait for both hosts to reply.
- **HA mode**: See **Primary server** and **Secondary server** above. If only the **Primary server** is specified, this field is disabled.
- **Port**: Enter the port number for the RADIUS Server. Click the **Default** button to use the default value.
- **Timeout**: Specify the packet timeout in seconds.
- **Retries**: Specify the number of retries when attempting to establish a connection.
- **Secret key**: Type the secret key.
- **Password encoding**: Choose from PAP (Password Authentication Protocol) or CHAP (Challenge Handshake Authentication Protocol), according to the setting specified in your RADIUS server.

Click the **Check connection** button to validate the connection. If the connection is configured correctly, you will see a confirmation message.

Specify additional properties as required:

- **Forward username only to RADIUS server**: Select this option if needed.
- **Forward the first password to Windows authentication provider**: Select this option to avoid a prompt to enter the password twice (RADIUS and Windows AD). Note that for Azure MFA server, this option is always enabled and cannot be turned off.
- Please also read a note at the bottom of the dialog (if available) suggesting certain setting specifics for the selected RADIUS solution.

Attributes

If your RADIUS solution requires configuring attributes, click the **Attributes** tab and then click **Add**. In the dialog that opens, choose a desired preconfigured vendor and attribute:

- In the **Vendor** drop-down list, select a vendor.
- In the **Attribute** list, select a vendor attribute.
- In the **Value** field, enter a value for the selected attribute type (numeric, string, IP address, date, etc).

In certain scenarios you may need to add vendors and attributes that are not listed in this dialog. For the information about how to add vendors and attributes, please see the following KB article: <https://kb.parallels.com/en/125576>.

Click **OK** and then click **OK** again .to close all dialogs.

Automation

The **Automation** tab in the RADIUS **Properties** dialog allows you customize the OTP experience for Parallels Client users by configuring security verification methods and custom commands to be sent to a RADIUS server during the MFA login process. Different security verification methods can be assigned priority and configured to be automatically used.

With this functionality configured, users can choose their preferred security verification method from a predefined and configurable list including Push notification, Phone Callback, SMS, Email, and Custom. The methods appear as clickable icons on the OTP dialog in Parallels Client. When a user clicks an icon, a command is sent to the RADIUS server and the corresponding verification methods is used.

To configure a verification method (also called "actions" here and in the Parallels RAS Console), on the **Automation** tab, click **Tasks > Add**. In the **Add Action** dialog, specify the following properties:

- **Enable Action:** Enables or disables the action.
- **Title:** The text that will appear on the clickable icon in Parallels Client (e.g. "Push").
- **Command:** The OTP command to be used when the action icon is clicked in Parallels Client. Consult your MFA provider for command specifications.
- **Description:** A description that will appear on the user's screen as a balloon when the mouse pointer hovers over the action icon.
- **Action message:** A message to show to the user in the connection progress box.
- **Select an image:** Select an image from the provided gallery. The image is used as the action icon in the OTP dialog in Parallels Client.

When done, click **OK** to save the action. Repeat the steps above for other actions.

Note: You can create up to five actions. When all five are created, the **Tasks > Add** menu is disabled.

You can move the actions on the **Automation** tab up or down the list. This dictates in which order the action icons will be displayed in Parallels Client.

Autosend

There's one more option that you can configure for an action. It is called **Autosend**. The option can be enabled for one action only, making it a default action, which will be used automatically without user interaction.

To enable the **Autosend** option, select an action on the **Automation** tab and click **Tasks > Autosend**. To disable the option, click the same menu again. If you enable **Autosend** for a different action, it will be automatically disabled for the previous action.

There are two possible ways to make an action execute automatically in Parallels Client:

- Client is receiving the action icon configuration for the first time and one of the actions has **Autosend** enabled.
- Enabling the **Remember last method used** option in **Policies > Session > Connection > Multifactor authentication**. When the option is enabled, and Parallel Client receives the policy, the last method successfully used by the user will become the default automatic method.

Parallels Client

When the user logs in to Parallels RAS via MFA, the OTP dialog is shown in Parallels Client with the actions icons positioned above the OTP field. The user clicks an icon and the authentication is carried out according to the predefined action. For example, if the user clicks the "Push" icon, a push notification is sent to the user mobile device where they can simply tap "Approve". Or there could be a "Text me" icon, in which case a text is sent to the user mobile phone with a one-time password. If one of the actions has the **Autosend** option enabled, then this action is used automatically.

If a user always uses the same authentication method, they can make it the default one. To do so, the user enables the **Remember last method used** option in the **MFA authentication** section of the connection properties. Depending on the platform, the option can be found at the following locations:

- Parallels Client for Windows / Linux: **Connection Advanced Settings > MFA authentication**
- Parallels Client for Mac: **Advanced > MFA authentication**
- Parallels Client for Chrome: **Advanced Settings**
- HTML5 Client: **Settings**
- Parallels Client for iOS: **Connection Settings > MFA authentication**
- Parallels Client for Android: **Settings > MFA authentication**

As was already mentioned above, the **Remember last method used** can also be configured in Client Policies in the RAS Console. The option is enabled by default.

Configuring Azure MFA

Before reading this section, please read the following important note.

Note: As of July 1, 2019, Microsoft will no longer offer MFA Server for new deployments. New customers who would like to require multi-factor authentication from their users should use cloud-based Azure Multi-Factor Authentication. Existing customers who have activated MFA Server prior to July 1 will be able to download the latest version, future updates, and generate activation credentials as usual:
<https://docs.microsoft.com/en-us/azure/active-directory/authentication/howto-mfaserver-deploy>.

For new deployments, it is recommended to use Azure NPS Extension <https://docs.microsoft.com/en-us/azure/active-directory/authentication/howto-mfa-nps-extension> or Azure MFA Service along with SAML configuration in RAS.

Configure Azure MFA

Depending on the user location, there are four scenarios for the cloud MFA service:

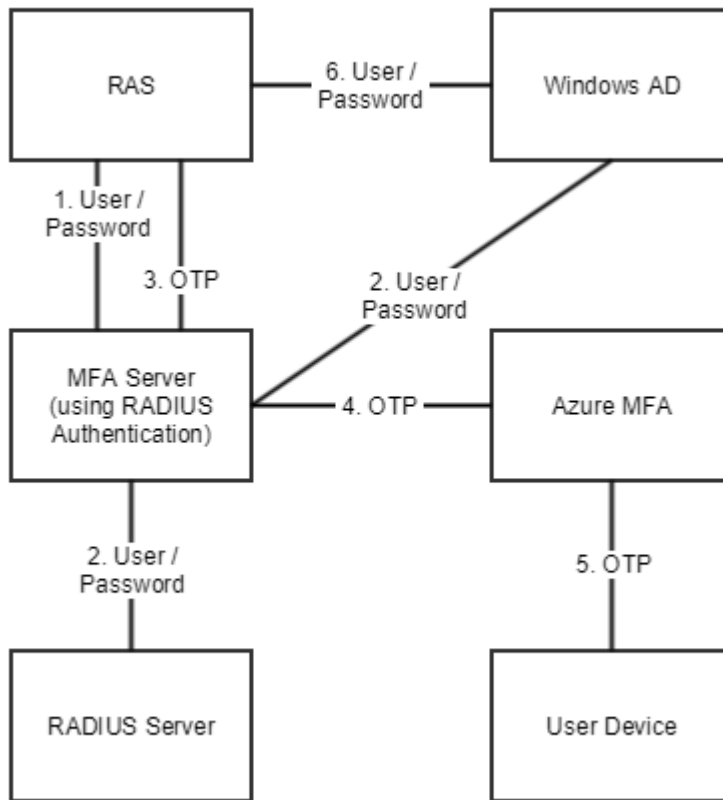
User location	MFA in the cloud	MFA Server
Azure Active Directory	Yes	
Azure AD and on-premises AD using federation with AD FS (is required for SSO)	Yes	Yes
Azure AD and on-premises AD using DirSync, Azure AD Sync, Azure AD Connect - no password sync	Yes	Yes
Azure AD and on-premises AD using DirSync, Azure AD Sync, Azure AD Connect - with password sync	Yes	
On-premises Active Directory		Yes

An Azure account with Global Administrator role is required to download and activate MFA Server. Syncing with Azure AD (via AD Connect) or a custom DNS domain aren't required to setup an MFA Server which runs exclusively on-premises.

Users need to be imported into MFA Server and be configured for MFA authentication.

Parallels RAS authenticates users with MFA Server using the RADIUS second level authentication provider. MFA Server thus needs to be configured to allow RADIUS client connections from the RAS server.

The authentication process goes through the following stages:



In stage 2 the user can be authenticated using either RADIUS or Windows AD. A prompt to enter the credentials twice (in stage 1 and 6) is avoided by enabling the option to forward the password.

Configuring Duo

For instructions on how to configure Parallels RAS with Duo RADIUS, please read the following Parallels KB article: <https://kb.parallels.com/124429>.

Using Deepnet DualShield

This section explains how to integrate Deepnet DualShield Authentication Platform 5.6 or higher with Parallels RAS.

In this section:

- **Supported tokens** (p. 276)
- **Configuring DualShield 5.6+ Authentication Platform** (p. 276)

- **Configuring Parallels RAS to use the DualShield Authentication Platform** (p. 279)
- **Connect to a RAS Farm** (p. 281)

You may also read the following documentation on DualShield Authentication Platform:

- DualShield Authentication Platform – Installation Guide
- DualShield Authentication Platform – Quick Start Guide
- DualShield Authentication Platform – Administration Guide

Supported tokens

The following is the list tokens supported by Parallels RAS:

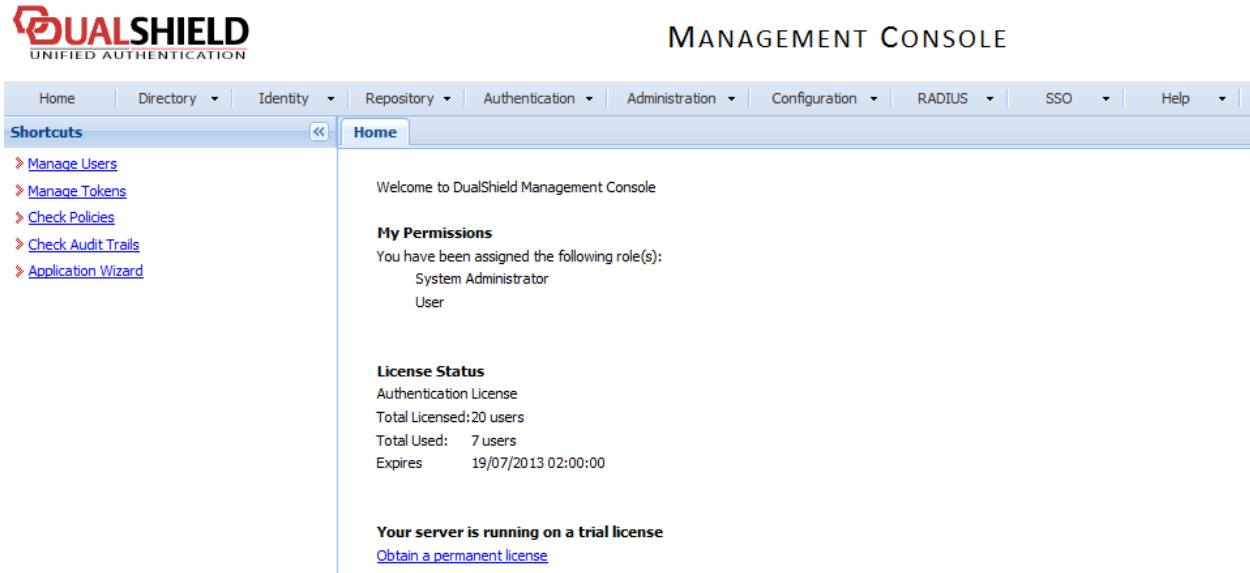
- MobileID (FlashID is not integrated with MobileID)
- QuickID
- GridID
- SafeID
- SecureID (RSA)
- DigiPass (Vasco)

If using hardware tokens such as SafeID the token information must first the XML file provided. Click on 'Import' and browse for the XML file provided. After the XML file has been imported each hardware token must be assigned to a user.

Configuring DualShield 5.6+ Authentication Platform

After following all the specified steps in "DualShield Authentication Platform – installation Guide" a URP is automatically opened in your internet browser ([http:// LOCALHOST:8073](http://LOCALHOST:8073)) which allows you to logon to the Management Console of DualShield.

Login in to the DualShield Management Console with the default credentials (User: sa, Password: sa). You will be prompted to change the default password.



Applications are set to provide a connection to realm, as the realm contains domains of users who will be allowed the access to the application.

Realm is set for multiple domain users to be able to access the same application.

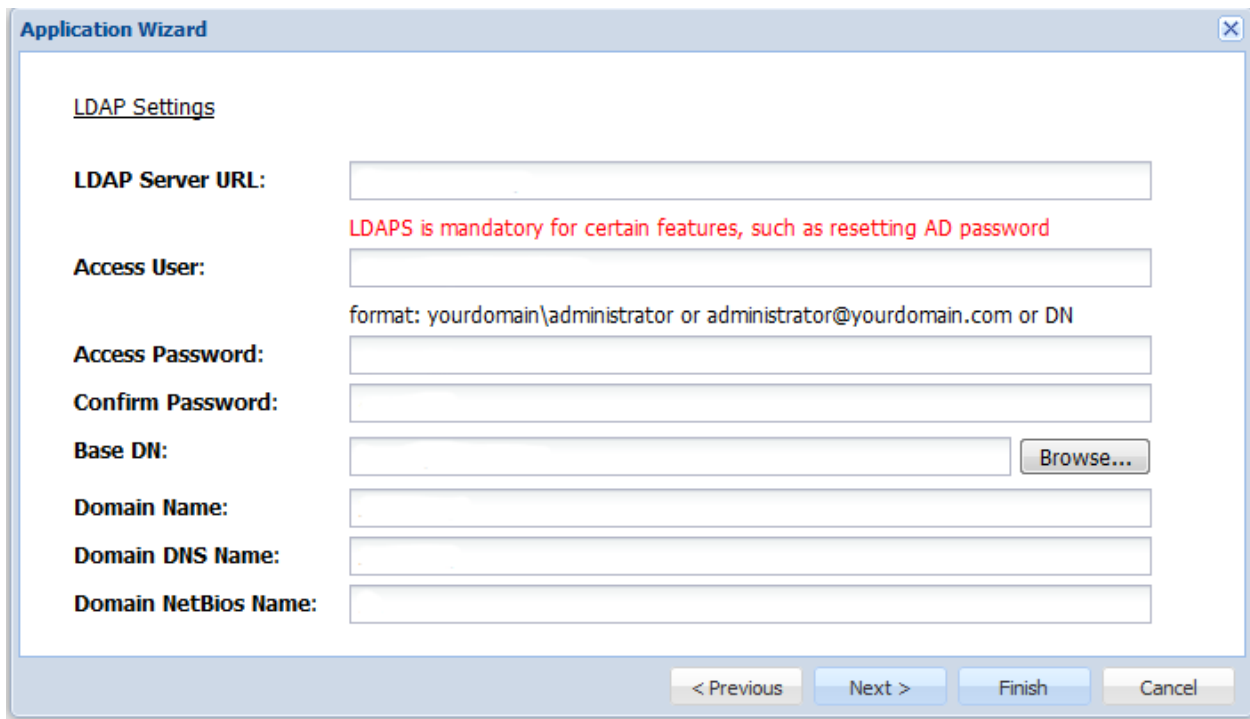
You need to create an Application which Parallels RAS will communicate with. Click on **Authentication > Application Wizard** and enter the information shown below and press **Next**.

The screenshot shows the 'Application Wizard' dialog box with the 'General Info' tab selected. It contains the following fields:

- Application Name:** A text input field.
- Application Type:** A dropdown menu.
- Identity Source Type:** A dropdown menu with 'LDAP' selected.
- Provider:** A dropdown menu with 'Active Directory' selected.

At the bottom of the dialog are three buttons: 'Next >', 'Finish', and 'Cancel'.

Specify the LDAP Server settings as shown below and press **Finish**.

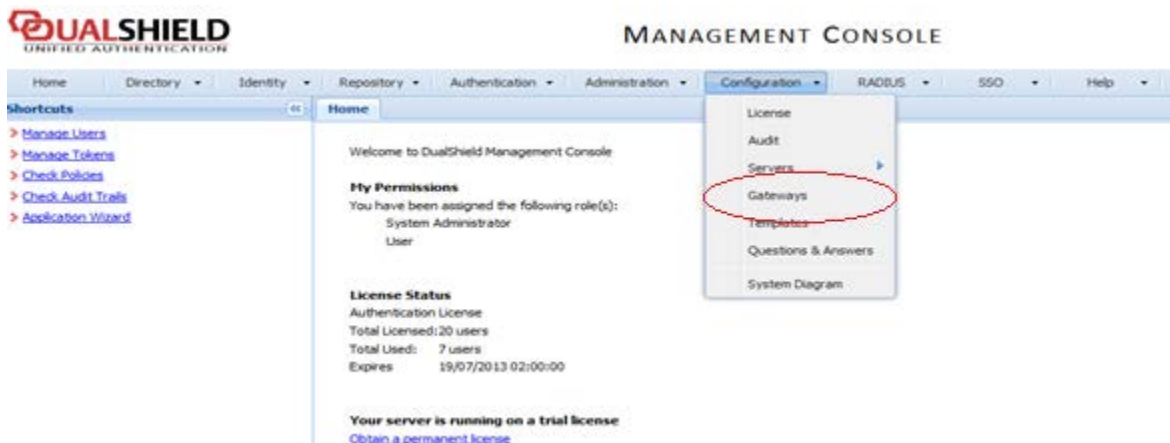


The screenshot shows the 'Application Wizard' dialog box with the 'LDAP Settings' tab selected. The fields are as follows:

- LDAP Server URL:** [Empty text box]
- Access User:** [Empty text box]
- Access Password:** [Empty text box]
- Confirm Password:** [Empty text box]
- Base DN:** [Empty text box] with a 'Browse...' button to its right.
- Domain Name:** [Empty text box]
- Domain DNS Name:** [Empty text box]
- Domain NetBios Name:** [Empty text box]

Below the fields are navigation buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'. A red text note reads: 'LDAPS is mandatory for certain features, such as resetting AD password'.

After you have configured the application you need to configure an Email or SMS gateway which are used by DualShield server to communicate with the end user. In this document we will be using an Email gateway. Select Gateways from the Configuration menu.



Configure your email gateway.

Message Gateway -- Edit

Type: EMAIL

Name:

Description:

Configuration:

Domains:

Enable:

Click **Edit** to enter your SMTP server information

SMTP Gateway

SMTP Server Address:

SMTP Server port: 465

Requires Authentication:

Transport Layer Security (TLS):

Secure Socket Layer(SSL):

User Name:

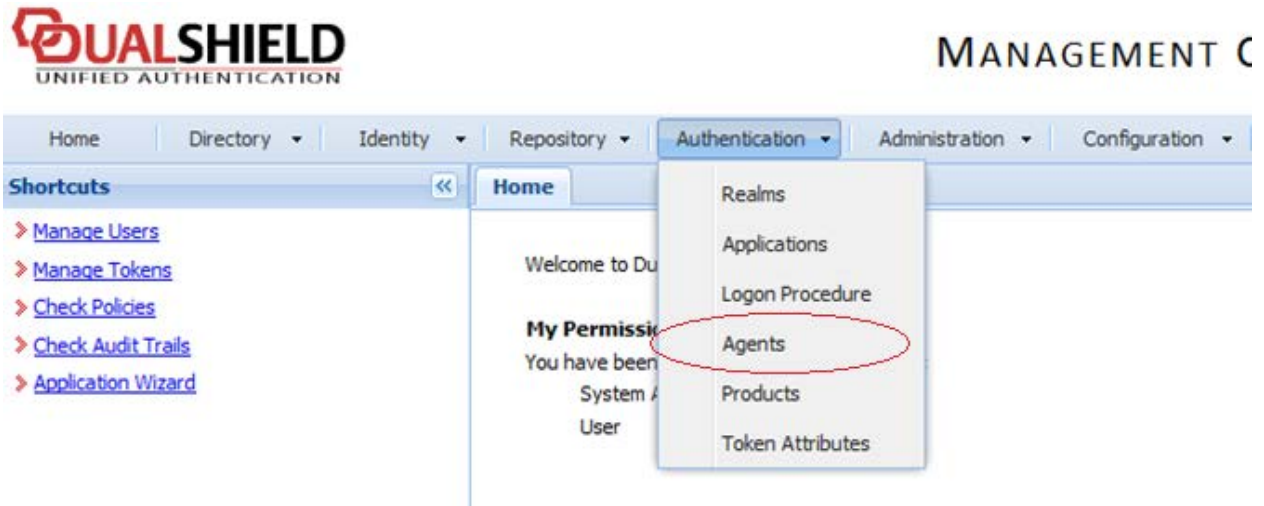
Password:

Configuring Parallels RAS to use the DualShield Authentication Platform

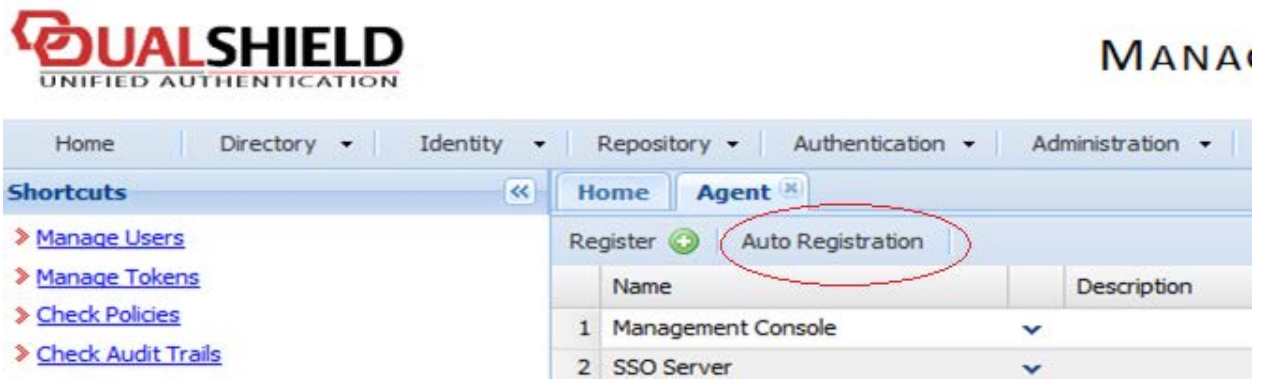
To begin:

- 1 In the RAS Console, navigate to the **Connection > Multi-factor authentication** tab.
- 2 In the **Provider** drop-down list, select **Deepnet**.
- 3 Click the **Settings** button.
- 4 Click the **Check Connection** button to test that the authentication server can be reached and to verify that the RAS Console is registered as a DualShield agent. If you see the "DeepNet server not valid" message, you have either specified an incorrect server information or you need to allow auto registration of the Parallels components as a DualShield agent.

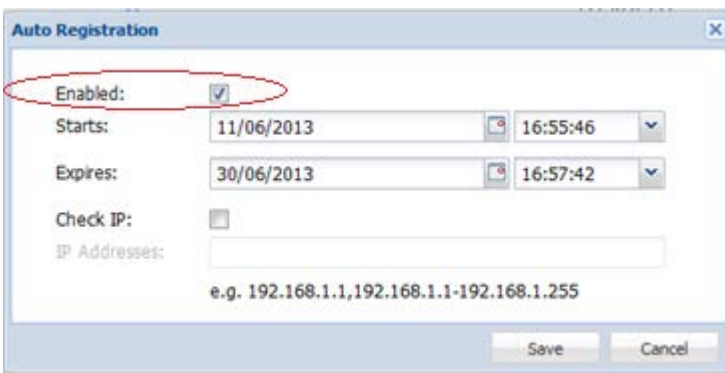
- 5 Go back to the DualShield Management Console and select **Agents** from the **Authentication** menu as shown below.



- 6 Select **Auto Registration**.



- 7 Select the **Enabled** option and set the date range.



- 8 Once the Agent Auto Registration is set, go back to the RAS Console and select **Yes**. You should see a message that the Dual Shield agent has been successfully registered.

Please note that all RAS Publishing Agents must be registered with Deepnet DualShield server. If you are using secondary Publishing Agents, you need to close all open windows until you can press **Apply** in the RAS Console. This will inform all the agents to self-register as DualShield agents.

- 9 In the Deepnet Properties dialog, click the **Applications** tab and browse for the Application name previously created from the DualShield Management Console.
- 10 Click the **Authentication** tab and select how you want your users to be authenticated:
 - **Mandatory for all users** means that every user using the system must log in using two-factor authentication.
 - **Create token for Domain Authenticated Users** will allow Parallels RAS to automatically create software tokens for Domain Authenticated Users. Choose a token type from the drop down list. Note that this option only works with software tokens, such as QuickID and MobileID
 - **Use only for users with a DualShield account** will allow users that do not have a DualShield account to use the system without have to login using two-factor authentication.
- 11 Go back to the **Connection > Multi-factor authentication** tab.
- 12 In the Exclusion section, specify the exclusion rules:
 - **User / Group exclude list** allows you to add users or groups within your active directory that will be excluded from using DualShield Authentication.
 - **Client IP exclude list** allows you to add IP addresses or a range of IP addresses that will be excluded from using DualShield Authentication. Please note that MFA exclusion by IP should not be used in scenarios when external and internal client subnets may overlap.
 - **Client MAC exclude list** allows you to add MAC addresses that will be excluded from using DualShield Authentication. You can also specify a MAC address range using double question marks as a wildcard in any part of the address. For example, 00-14-22-01-23-??, 00-14-22-01-??-??, or 00-14-22-??-??-??.
 - **Connection to the following Gateway IPs** allows you to set a Gateway where users connected to the Gateway will be excluded from using DualShield Authentication.

Connect to a RAS Farm

Parallels Client

Once DualShield has been enabled the users will have two-factor authentication. If using software tokens such as QuickID the administrator does not have to create a token for each user. RAS Publishing Agent will automatically create the token when the user tries to log in for the first time.

When a user tries to access a RAS Connection from Parallels Client, they are first prompted for the Windows username and password. If the credentials are accepted, RAS Publishing Agent will communicate with the DualShield server to create a unique token for that user.

If using MobileID or QuickID, an email about where to download the appropriate software will be sent to the user.

If using QuickID tokens, the application will ask for a One-Time Password which is sent by e-mail or SMS.

When asked for OTP, enter the One-Time Password to log in to the Parallels ApplicationServer XG Gateway.

Using SafeNet

SafeNet Token Management System provides a high-value of protection via secure tokens which makes it a perfect tool for second-level authentication in Parallels RAS.

In this section:

- Configuring SafeNet (p. 282)

Configuring SafeNet

To configure SafeNet:

- 1 In the Parallels RAS console, navigate to the **Connection > Multi-factor authentication** tab.
- 2 In the **Provider** drop-down list, select **SafeNet**.
- 3 Click the **Settings** button. The **SafeNet Properties** dialog opens.
- 4 On the **Connection** tab, enter the valid URL into the **OTP Service URL** field. To verify that the connection with the OTP Service can be established, click the **Check connection** button.

Note: RAS Publishing Agent communicates with the SafeNet Token Management System Server. It is highly recommended to have this behind a firewall for security reasons.

- 5 Click the **Authentication** tab.
- 6 In the **Mode** drop-down list, select how you want your users to be authenticated.
Mandatory for all users: every user using the system must login using two-factor authentication.

The available modes are:

- **Create token for Domain Authenticated Users:** Allows Parallels RAS to automatically create software tokens for Domain Authenticated Users. Choose a token type from the drop down list. Note that this option only works with software tokens.
 - **Use only for users with a SafeNet account:** Allows users that do not have a SafeNet account to use the system without having to login using two-factor authentication.
- 7 In the **TMS Web API URL** field, enter the location of the SafeNet API URL.
 - 8 In the **User Repository** field, enter the user repository destination.
 - 9 Click **OK** to save the values and close the **SafeNet Properties** dialog.

Parallels Client

In **Parallels Client – New Account Info** dialog:

- 1 Enter any four digits in the **OTP PIN** number field (these digits will be required further on in the process).
- 2 Enter your email address and then click on **OK**.
- 3 Log into your email account and retrieve the email containing the information you will need to activate your SafeNet authentication. An example of this email is shown below.

Activation Key: YZQHoczZWw3cBCNo

Token Serial: 4F214C507612A26A

Download MobilePASS client from:

<http://localhost:80/TMSService/ClientDownload/MobilePASSWin.exe>

**Login with domain credentials.*

**Place the attached seed file in the same folder as the MobilePASS client.*

Enter the One-Time Password to log into the RD Session Host Connection.

Application PIN: 4089

- 4 Download the MobilePASS client from the URL provided in the email.
- 5 Enter the Activation Key found in the SafeNet email.
- 6 Next, input the application PIN found in the email into the **MobilePASS PIN** field.
- 7 Click **Generate** to generate the eToken number and then click **Copy**.
- 8 Combine the OTP PIN and eToken in this order: OTP + eToken.
- 9 Enter this value into the Parallels Client and click **OK** to log in.

Using Google Authenticator

This section explains how to use Google Authenticator as a second-level authentication solution in Parallels RAS.

To configure Google Authenticator:

- 1 In the Parallels RAS Console, navigate to **Connection > Multi-factor authentication**.
- 2 In the **Provider** drop-down list, select **Google Authenticator**.
- 3 Click the **Settings** button.
- 4 In the **Google Authenticator Properties** dialog that opens, specify the following options:

- **Type Name:** The default name here is Google Authenticator. The name will appear on the registration dialog in Parallels Client in the following sentence, "Install Google Authenticator app on your iOS or Android device". If you change the name, the sentence will contain the name you specify, such as "Install <new-name> app on your iOS or Android device". Technically, you can use any authenticator app (hence the ability to change the name), but at the time of this writing only the Google Authenticator app is officially supported.
- The **User enrollment** section allows you to limit user enrollment via Google Authenticator if needed. You can allow all users to enroll without limitations (the **Allow** option), allow enrollment until the specified date and time (**Allow until**), or completely disable enrollment (the **Do not allow** option). If enrollment is disabled due to expired time frame or because the **Do not allow option** is selected, a user trying to log in will see an error message saying that enrollment is disabled and advising the user to contact the system administrator. When you restrict or disable enrollment, Google authenticator or other TOTP provider can still be used, but with added security which would not allow further user enrollment. This is a security measure to mitigate users with compromised credentials to enroll in MFA.
- The **Authentication** section allows you to configure TOTP tolerance. When using Time-based One-Time Password (TOTP), it is required to have the time synchronized between the RAS Publishing Agent and client devices. The synchronization must be performed against a global NTP server (e.g. time.google.com). Using the **TOTP tolerance** drop-down box, you can select a time difference that should be tolerated while performing authentication. Expand the drop-down box and select one of the predefined values (number of seconds). Note that changing time tolerance should be used with caution as it has security implications since the time validity of a security token can be increased, thus a wider time window for potential misuse.

Note: When using Time-based One-time Passwords (TOTP) providers, it is required to have both Publishing Agents and client devices time synchronized with a global NTP server (e.g. time.google.com). Adding TOTP tolerance increases the one-time password validity, which might have security implications.

- The **Reset User(s)** field in the **User management** section is used to reset the token that a user receives when they log in to Parallels RAS for the first time using Google Authenticator. If you reset a user, they'll have to go through the registration procedure again (see **Using Google Authenticator in Parallels Client** below). You can search for specific users, reset all users, or import the list of users from a CSV file.

5 Click **OK** when done.

Please also note that the TOTP available time is calculated as the default 30 seconds + x amount of seconds in the past + x amount of second in the future.

Using Google Authenticator in Parallels Client

Important: To use Google Authenticator or other TOTP provider, the time on a user device must be in sync with the time set on the RAS Publishing Agent server. Otherwise, Google authentication will fail.

Google Authenticator is supported in Parallels Client running on all supported platforms, including mobile, desktop, HTML5.

To use Google Authenticator, a user needs to install the Authenticator app on their iOS or Android device. Simply visit Google Play or App Store and install the app. Once the Authenticator app is installed, the user is ready to connect to Parallels RAS using two-factor authentication.

To connect to Parallels RAS:

- 1 The user opens Parallels Client or HTML5 Client and logs in using his/her credentials.
- 2 The multi-factor authentication dialog opens displaying a barcode (also known as QR code) and a secret key.
- 3 The user opens the Google Authenticator app on their mobile device:
 - If this is the first time they use it, they tap **Begin** and then tap **Scan a barcode**.
 - If a user already has another account in Google Authenticator, they tap the plus-sign icon and choose **Scan a barcode**.
- 4 The user then scans the barcode displayed in the Parallels Client login dialog.
If scanning doesn't work for any reason, the user goes back in the app, chooses **Enter a provided key** and then enters the account name and the key displayed in the Parallels Client login dialog.
- 5 The user then taps **Add account** in the app, which will create an account and display a one time password.
- 6 The user goes back to Parallels Client, clicks **Next** and enters the one time password in the **OTP** field.

On every subsequent logon, the user will only have to type their credentials (or nothing at all if the **Save password** options was selected) and enter a one time password obtained from the Google Authenticator app (the app will continually generate a new password). If the RAS administrator resets a user (see the **Reset Users(s)** field description at the beginning of this section), the user will have to repeat the registration procedure described above.

Configuring MFA exceptions

Multi-factor authentication (MFA) can be enabled or disabled for all user connections, but you can configure exceptions for specific connections.

To enable or disable MFA for all users:

- 1 First, select from the following two options:
 - **Enable MFA for all users except**
 - **Disable MFA for all users except**
- 2 In the **Restriction** section, specify exceptions from the common rule you selected above by selecting and configuring one of the options described below.

User or group list

Specify users or group names to be excluded from the common rule. Names must be entered in UPN format (username@domain.com).

Please note the following:

- For users to connect, the **Force clients to use NetBIOS credentials** option must be disabled (the option is located in **Connection > Authentication**). Users must log in using their names in UPN format (username@domain.com).
- The exclusion requires a domain environment and doesn't work in Workgroup.
- Group nesting is not supported when configuring an exclusion.

Client IP list

Specify a single IP address or a range of addresses to be excluded from the common rule.

Client MAC list

Specify MAC addresses. You can specify a MAC address range using a double question mark as a wildcard in any part of the address. For example, 00-14-22-01-23-??, 00-14-22-01-??-??, or 00-14-22-??-??-??.

Connection to the following Gateway IPs

Specify Gateway IP addresses. Users connecting through these addresses will be excluded from the common rule.

Allowing users to change domain password

Users can change their domain password directly from Parallels Client. In some cases, a user may be forced to change their domain password (e.g. when the password is about to expire). When changing a password, a username must be supplied in the UPN format (e.g. user@domain.com). Since users may not know the name of their domain, you can configure Parallels RAS to pass it to the client side automatically, so users don't have to enter it.

The domain name may be specified in the RAS Console in the following locations:

- The **Connection > Authentication** tab page. The tab page is described earlier in this section (p. 265). To force a domain name to the client side, select the **Specific** option and specify a domain name.

- In the **Theme Properties** dialog. Themes are described later in this guide in the **Configure Themes** (p. 346) section. Note that when you specify a domain name for a Theme, it overrides the domain name specified on the **Authentication** tab page (see above). To specify a domain name for a Theme, open the Theme properties dialog, select the **General** category, select the **Override authentication domain** option, and specify a domain name.

When Parallels Client connects to Parallels RAS, the domain name, specified as described above, is passed back to it. When the user opens a dialog in Parallels Client to change their domain password, the domain name is automatically added to the user name and the username field is grayed out. This way the user doesn't have to specify the domain name.

Load Balancing and HALB

This chapter describes load balancing options that you can use in Parallels RAS.

In This Chapter

Resource based & round robin load balancing	288
High availability load balancing (HALB).....	291

Resource based & round robin load balancing

Load Balancer in Parallels RAS is designed to balance RD Session Host connections from Parallels Clients.

The following types of load balancing are available:

- **Resource based.** Distributes sessions to servers depending on how busy the servers are. A new incoming session is always redirected to the least busy server.
- **Round robin.** Redirects sessions in sequential order. For example, let's say there are two RD Session Hosts in the Farm. The first session is redirected to server 1, the second session is redirected to server 2, and the third session is redirected to server 1 again.

Both methods are explained in this and the following subsections. Load Balancing options can be configured from the **Load Balancing** category in the RAS Console.

Selecting load balancing method

Load balancing is enabled by default when more than one server is available in a Site. The resource based load balancing is the default method. Load balancing method can be selected from the **Method** drop-down list.

Configuring resource counters

Resource-based load balancing uses the following counters to determine if a given server is busier than other servers and vice versa:

- **User sessions:** Redirect users to a server with the least number of sessions.
- **Memory:** Redirect users to the server with the best free/used RAM ratio.

- **CPU:** Redirect users to the server with the best free/used CPU time ratio.

When all of the counters are enabled, the Load Balancer adds the counter ratios together and redirects the session to the server with the most favorable combined ratio.

To remove a counter from the equation, clear the checkbox next to the counter name in the **Counters** section.

Session options

Reconnect to disconnected sessions: Enable this option to redirect incoming user sessions to a previously disconnected session owned by the same user.

Reconnect sessions using client's IP address only: When reconnecting to a disconnected session, the Parallels RAS will match the username requesting reconnection with the username of the disconnected session to match the sessions. With this option enabled, Parallels RAS will determine to which disconnected session to reconnect the session by matching the source IP address.

Limit each user to one session per desktop: Enable this option to ensure that the same user does not open multiple sessions. Please note that for this option to work, your session host must also be configured to restrict each user to a single session. In Windows Server 2008, you need to enable the "Restrict each user to a single session" option in Remote Desktop Session Host Configuration. In Windows Server 2012(R2), it's the "Restrict Remote Desktop Services users to a single Remote Desktop Services session" option in Local Group Policy \ Remote Desktop Services \ Remote Desktop Session Host \ Connections.

Disable Microsoft RD Connection Broker: If this option is enabled, the Microsoft RD Connection Broker will not interfere with the RAS brokering done by the RAS Publishing Agent if it is installed. Please note that this option will only work with Windows Server 2012 and above.

Agent timeout and refresh time

You can also change the default timeout and refresh time for RAS agents running on the servers. If you believe that it takes too long to wait for an agent to respond or if the timeout is not long enough, you can specify your own values.

To change default timeouts:

- 1 Click the **Configure** button.
- 2 In the dialog that opens, specify the time period in seconds in the **Declare agent dead if not responding for** field. If the agent is not responding within this time period, the server is excluded from the load balancer.
- 3 In the **Agent Refresh Time** field, specify the number of seconds needed to check if the agent is reachable.

Configure CPU optimization

The CPU optimization functionality allows you to optimize CPU load balancing according to your requirements. When configured, the CPU load balancer will lower the priority of a process when its CPU usage exceeds a specified value for a specified number of seconds. The load balancer will revert the priority to its original level when the process has been running below a certain percentage for a certain number of seconds.

To configure CPU optimization, select the **Enable CPU Optimization** option and then specify the values as described below.

Start

Specifies when the CPU optimization should be activated. The **Total CPU usage exceeds** field specifies the system wide CPU usage in percent.

CPU conditions

Specifies thresholds per process when a specific process exceeds or falls below the specified CPU percentage. Here you can specify **Critical** and **Idle** values. The CPU load balancer will adjust other priorities with respect to these values.

Please note that CPU usage values are attenuated and calculated based on the agent refresh time configured on the **Load Balancing** tab (p. 288).

Exclusions

Use the **Exclusions** list to specify processes that should be excluded from CPU optimization. Click **Tasks > Add** to select a process. To remove a process from the list, select it and click **Tasks > Delete**.

Irregular values for critical/idle may cause issues (processes set to idle due to incorrect configuration). If there are issues with getting the CPU usage counter, optimizations cannot be applied.

Log files can be found in %ProgramData%\Parallels\RASLogs\cpuloadbalancer.log. Use the log to confirm thresholds. You can check the CPU usage performance counter on Windows.

Note: Since the critical/idle thresholds are calculated based on the highest process CPU usage (not the absolute CPU usage), this value is not reflected in the logs when changing priorities.

Absolute CPU usage equals to total CPU usage. For example, if there are 2 processes taking 30% each, the total CPU usage is 60%. The usage threshold when CPU load balancer kicks in is 25% (default).

The highest process CPU usage is the CPU usage of the process taking the most CPU. For example, if you have three processes, two taking 10% and the third taking 40%, the highest CPU usage is 40%.

High availability load balancing (HALB)

High availability load balancing (HALB) in Parallels RAS is a functionality that load balances RAS Secure Client Gateways. The load balancer is built into a Parallels HALB appliance, which is a preconfigured virtual machine with the operating system installed and all relevant settings configured.

Parallels HALB appliance is available for the following hypervisors:

- Microsoft Hyper-V
- VMware

Please note that other hypervisors may also be used, but support is provided as best effort. The Parallels RAS HALB appliance uses the Open Virtualization Platform (OVA) format, which is natively supported by various hypervisor.

HALB is deployed in Parallels RAS on a Site level. You can have multiple HALB configurations per Site, which are called Virtual Servers. Each Virtual Server has its own IP address (called Virtual IP or VIP) and is assigned one or more HALB appliances (also called HALB devices in the Virtual Server context) that perform the actual load balancing. An HALB Virtual Server is a virtual representation of HALB devices. It provides traffic distribution to HALB devices when they are properly configured. Since the IP address of a specific Virtual Server is the single point of contact for the client software, it is recommended to have at least two HALB devices per Virtual Server for redundancy.

Multiple HALB devices assigned to a Virtual Server can run simultaneously, one acting as the primary and others as secondary. The more HALB devices a Virtual Server has, the lower the probability that end users will experience downtime. The Virtual Server is assigned the IP address of the primary HALB device, which is shared with secondary HALB devices. Should the primary HALB device fail, a secondary is promoted to primary and takes its place using the same IP address for client connections.

Note: Please note that when a secondary HALB device is promoted to primary, or when the primary HALB device goes back online, a user may experience up to two disconnects. The first disconnect will occur when an HALB device goes down. The second disconnect may happen when a device goes back online. The disconnects cannot be avoided because the virtual IP address has to be transferred from one HALB device to the other, which means that the first device has to stop communications over this address, while the other device will have start it. Note that disconnects don't affect user sessions. Users are able to reconnect to their sessions and no user data is lost.

Setting up High Availability Load Balancing consists of the following steps:

- 1 Deploying one or more Parallels HALB appliances (devices).
- 2 Configuring one or more Virtual Servers in the RAS Console.

Read on to learn how to download and deploy a Parallels HALB appliance.

Deploying a Parallels HALB appliance

To download a Parallels HALB appliance, visit <https://www.parallels.com/products/ras/download/links/>.

On the **Download Parallels Remote Application Server** web page, scroll down to the **Download Optional Server Components** table and find the **Parallels Remote Application Server HALB Appliances** row. The row contains the following download links:

- HALB Appliance OVA
- HALB Appliance VHD
- HALB Appliance VMDK

The appliance type that you need to download depends on the hypervisor that you are using. Please follow the instructions below for your hypervisor type.

VMware

For VMware, the appliance can be imported with either the OVA or zipped VMDK appliance file. If deployed via the OVA file, the VM is created already configured.

Alternatively, deployment via the VMDK file deploys the VM without preconfigured specifications. The minimum specifications for this VM are outlined below:

- One CPU
- 256 MB RAM
- One network card

Microsoft Hyper-V

For Microsoft Hyper-V, the appliance is imported with the VHD file.

Deploying a Parallels HALB appliance

After you download a Parallels HALB appliance, you need to import it to a hypervisor running on a separate machine connected to the same local network as Parallels RAS. For the information on how to import a virtual appliance, please consult your hypervisor documentation.

Adding a HALB virtual server

To add a HALB virtual server:

- 1 In the RAS console, navigate to **Farm** > <Site> > **HALB**.
- 2 On the **Virtual Servers** tab in the right pane, click **Tasks** > **Add**. The **HALB Configuration** wizard opens.
- 3 Make sure the **Enable HALB** option is selected.
- 4 Type a name for this virtual server and an optional description.
- 5 In the **Public address** field, type a public FQDN or IP addresses of this server. This is used by the Preferred routing functionality for redirecting client connections. Please see **Configuring preferred routing** (p. 250).
- 6 In the **Virtual IP** section, specify the virtual IP address properties which will be used for incoming client connections by a HALB device that you will assign to this Virtual Server later.
- 7 In the **Settings** section, select one or more of the following options. Note that at least one "LB" option must be selected. If you skip an option at this time, you can add it later in the virtual server properties dialog:
 - **LB Gateway Payload:** Enables load balancing of normal (unsecured) gateway connections.
 - **LB SSL Payload:** Enables load balancing of SSL connections.
 - **Client Management:** Enables management of Windows client devices connected through HALB.
- 8 Click **Next**.

From this point forward, depending on the payloads that you selected in the previous step, a wizard page will open where you can configure the payload properties. These pages are described below.

LB Gateway payload

Configure load balancing for normal connections:

- 1 Set the port number used by HALB devices to forward traffic to RAS Secure Client Gateways. The port is configured on a gateway. The default port is 80.
- 2 In the **Gateways** list, select a RAS Secure Client Gateway to be load balanced. Please note that only one IP address per gateway can be used. If you have more than one entry for the same gateway with different IP addresses, you can select just one.

LB SSL payload

Configure load balancing for SSL connections:

- 1 Set the port number used by HALB devices to forward SSL traffic to RAS Secure Client Gateways. The port is configured on a gateway. The default port is 443.
- 2 Select the SSL mode from **Passthrough** or **SSL Offloading**. By default, SSL connections are tunneled directly to gateways (referred to as Passthrough) where the SSL decryption process is performed.

The **SSL Offloading** mode requires an SSL certificate to be assigned to HALB. When you select it, click **Configure** and specify the following:

- **Accepted SSL Version:** Select an SSL version.
- **Cipher Strength:** Select the cipher strength of your choice. To specify a custom cipher, select **Custom** and then specify the cipher in the **Cipher** field.
- The **Use ciphers according to server preference** option is ON by default. You can use client preferences by disabling this option.
- **Certificates:** Select a desired certificate. For the information on how to create a new certificate and make it appear in this list, see the **SSL Certificate Management** (p. 258) chapter.

The **<All matching usage>** option will use any certificate configured to be used by HALB. When you create a certificate, you specify the "Usage" property where you can select "Gateway", "HALB", or both. If this property has the "HALB" option selected, it can be used with HALB. Please note that if you select this option, but not a single certificate matching it exists, you will see a warning and will have to create a certificate first.

- 3 Select a gateway to be load balanced. Note that only one IP address per gateway can be used.

Device Manager

Configure Windows client device management, select a gateway that will manage Windows client devices. Note that only one IP address per gateway can be used.

Devices

To assign HALB devices to the Virtual Server:

- 1 Click **Tasks > Add** and select or specify a HALB device. If you haven't deployed any HALB devices (appliances) yet, you can still save the Virtual Server configuration and assign HALB devices to it later. At least two HALB devices are recommended per Virtual Server. For more info, see **High Availability Load Balancing (HALB)** (p. 291). HALB device priority is set by positioning a device in the list. The device at the top is the primary HALB device. Devices under it are secondary HALB devices. To promote a device to primary, simply move it to the top of the list.
- 2 Finally, click **Finish** to save the Virtual Server settings and close the wizard.

The new virtual server will appear in the list in the RAS Console.

Modifying Virtual Server and configuring advanced options

To modify the Virtual Server settings, right-click it and choose **Properties**. The tabs in the **Properties** dialog have the same options as the wizard pages described above. The only exception is the **Advanced** tab, which is described below.

To view and configure advanced Virtual Server options, select the **Advanced** tab. The options that you see on this tab are applied to all HALB devices assigned to a Virtual Server. This list gives you a simple access to the HALB device options without logging in to the virtual machine directly. Please note that changing any of these values may potentially lead to undesired results. You should only change them according to specific network requirements.

The following advanced settings are available:

Option	Default value	Description
Enable RDP UDP tunneling	Enable	Enables RDP clients to transfer RDP over UDP traffic through HALB devices.
Minimum TCP connections	2000	Sets the maximum number of concurrent TCP connections.
Client inactivity timeout (s)	150	Maximum inactivity time on the client side in seconds.
Gateway connection timeout (s)	30	Maximum time to wait for a connection attempt to a gateway to succeed in seconds.
Client connection queue timeout (s)	30	When a device's Max TCP connections is reached, connections are left pending in a queue for the period of this timeout (seconds).
Gateway inactivity timeout (s)	150	Set the maximum inactivity time for gateways in seconds.
Amount of TCP connections per second	1000	Set a limit on the number of new connections accepted per second on an HALB device.
Gateway health check intervals (s)	5	Set the interval between two consecutive health checks in seconds.
VRRP virtual router ID	15	Used to differentiate multiple instances of VRRP running on the same network.
VRRP authentication password	-	Enable password authentication for VRRP communication between HALB devices used by for failover synchronization.
VRRP broadcast interval (m)	1	Minimum time interval in minutes for refreshing gratuitous ARPs while device is in active state.
VRRP health script check interval (s)	2	Set the interval between invocations of the script that ensures local HALB services are up and running (seconds).
VRRP health script check timeout (s)	10	Execution timeout for the script that ensures local HALB services are up and running (seconds).
VRRP advertisement interval (s)	1	The time interval between the advertisement packets that are being sent between HALB devices in the same VRRP group (seconds).
Enable OS updates	Disable	Allow HALB devices to automatically update OS packages.

Keep existing load balancing settings	Disable	Keep load balancing configuration currently present on the device and do not overwrite with new settings.
Keep existing VRRP/keepalived	Disable	Keep VRRP/keepalived configuration currently present on the device and do not overwrite with new settings.

HALB Device status and version number

HALB device status and version information can be verified in two places in the RAS Console, which are described below.

Site tab

You can view HALB devices and related information on the **Site** tab in the RAS Console. To see it, navigate to **Farm > Site**. Note the **Agent** and **Agent Version** columns. The two columns are described below.

The **Agent column** can have the following values:

- **Not verified** (red) - The agent is not verified and cannot communicate. If you see this, verify the agent.
- **Needs update** (yellow) - The agent is functioning normally but is an older version. If you see this, you should update the agent to the latest version.
- **Agent OK** (green) - The agent is OK. No actions are necessary.

The **Agent Version** column displays the actual agent version, including the Parallels RAS version and build numbers.

You can also ping a HALB device by right-clicking it and choosing **Tools > Ping host**. For additional information about using the **Tools** menu and the **Ping** tool specifically, please see **Computer management tools** (p. 431).

Devices tab

The HALB devices agent status and version can also be viewed in the main HALB subcategory. To see it, navigate to **Farm > Site > HALB** and select the **Devices** tab. The agent information displayed here is the same as on the **Site** tab described above.

HALB maintenance

When you need to replace or repair a HALB device (virtual machine), you can simply remove it from the Virtual Server configuration and then add the repaired or new device later. If you need to temporarily remove all HALB devices from a Virtual Server configuration, you can do that too.

You can also disable the Virtual Server during maintenance by clearing the **Enable HALB** option on the **General** tab in the Virtual Server properties dialog.

HALB connection and session information

To see the number of TCP connections per HALB device, navigate to **HALB > Devices** and examine the **TCP Connections** column in the device list. To refresh the list, click **Tasks > Refresh**.

To see the session information per Virtual Server, navigate to **Farm > Site**. The session count is displayed for each Virtual Server in the **Session** column.

Changing the HALB appliance password

To change the HALB appliance password:

- 1 Boot the appliance (virtual machine).
- 2 Press the <ALT> – <F1> key combination. A login prompt should be displayed.

```
Debian GNU/Linux 7 LB-00-0C-29-DA-92-7A tty1
LB-00-0C-29-DA-92-7A login: root
Password: _
```

- 3 Type in the following credentials:
 - **login:** root
 - **password:** Pa\$w0rd (note that "0" is zero, not the letter "O").

```
Debian GNU/Linux 7 LB-00-0C-29-DA-92-7A tty1
LB-00-0C-29-DA-92-7A login: root
Password:
Linux LB-00-0C-29-DA-92-7A 3.2.0-4-686-pae #1 SMP Debian 3.2.51-1 i686
Welcome to Lb-00-0c-29-da-92-7a, 2X HALB / Debian 7.2 Wheezy

System information (as of Fri Apr 17 09:47:25 2015)

System load: 0.03          Memory usage: 13%
Processes: 63             Swap usage: 0%
Usage of /: 71.5% of 494MB IP address for eth0: 10.124.4.119

root@LB-00-0C-29-DA-92-7A ~# passwd_
```

- 4 Once logged in, execute the password changing command and type a new password.

```
root@LB-00-0C-29-DA-92-7A ~# passwd  
Enter new UNIX password: _
```

Upon completion, you may log in to the HALB device with the new password.

RAS Multi-Tenant Architecture

In This Chapter

Overview	299
Architecture description.....	300
Deploying Tenant Broker and Tenants.....	303
Managing Tenants.....	311
Shared Gateways.....	312
Third-party network load balancers	313
HTML5 Client and Themes.....	314
Monitoring Tenants.....	315
Tenant Broker compatibility and updates.....	315
Upgrading from an older RAS version.....	316
Configuring notifications	316
Communication ports.....	317

Overview

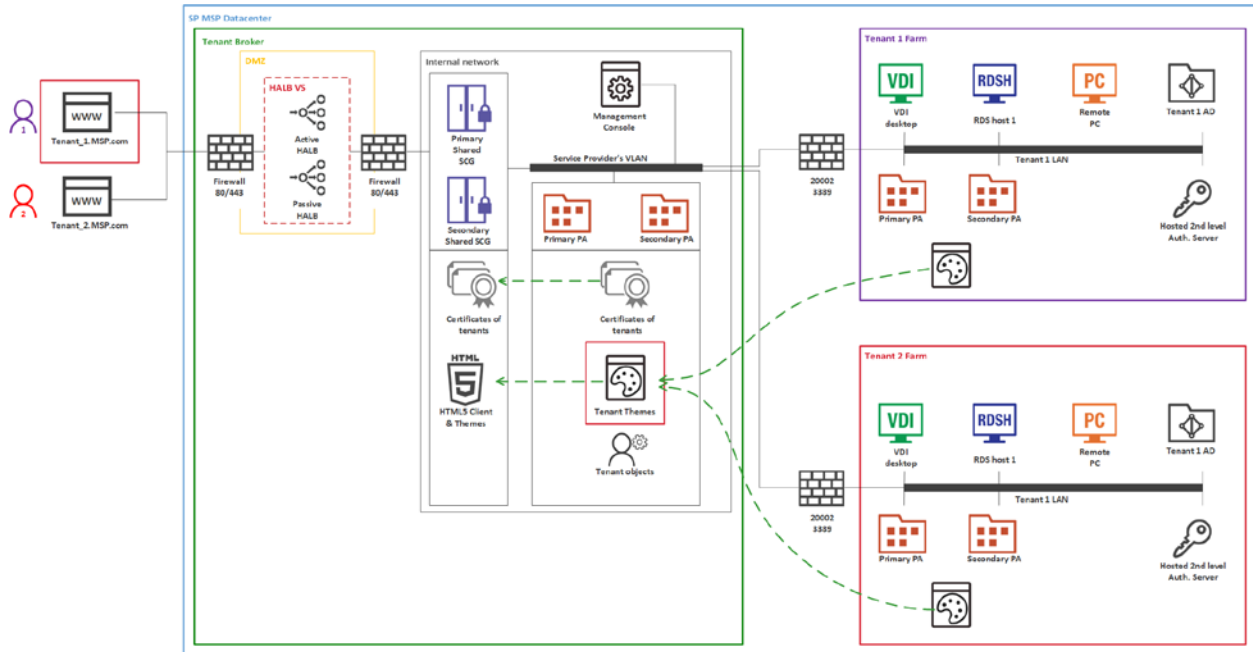
Beginning with RAS 17.1, Parallels introduces a new multi-tenant architecture, with the addition of Parallels RAS Tenant Broker, enabling organizations to share components from the same RAS infrastructure among different Tenants while keeping client data segregated and reducing costs.

The RAS multi-tenant architecture offers the following advantages to Service Providers and organizations:

- **Cost savings** due to reduction of number of RAS Secure Client Gateways and High Availability Load Balancers (HALBs) while maximizing resource usage and consolidation.
- **Faster onboarding** of new tenants/customers.
- **Simplified centralized management** of multi-tenant environments.
- **Extended market reach** through reduction of operational costs for organizations of any size by allowing cost scaling through shared infrastructure.

Architecture description

The following diagram illustrates a typical Parallels RAS deployment that uses the RAS multi-tenant architecture.



- Firewalls and HALB are installed in a DMZ and are shared by Tenants.
- Tenant Broker is a special RAS installation that hosts shared RAS Secure Client Gateways and HALB, and can also use RAS access layer. Tenant Broker is installed using the **Parallels RAS Tenant Broker** option in the Parallels RAS installer. Tenant Broker can be installed in its own domain or outside of a domain.
- Tenant farms are deployed just like traditional on-premises RAS environments and are joined to the Tenant Broker. Each Tenant Farm has its own RAS Publishing Agents and servers hosting published resources (VDI, RD Session hosts, or Remote PCs). No local RAS Secure Client Gateways and HALB (or third-party load balancers) are needed.
- Tenants are joined to the Tenant Broker and each Tenant is represented as a Tenant object in the Tenant Broker.
- Parallels Clients (both platform-specific and HTML5) connect to shared gateways in the Tenant Broker. When a client connects to an HTML5 gateway, a Theme from the corresponding Tenant is always used depending on which Tenant the client belongs to.

Implementation overview

The following is an implementation overview of the RAS multi-tenant architecture:

- Tenants are deployed as separate individual Farms or Sites. Tenants deployed as separate Farms are completely independent and never communicate with each other. If tenants are deployed as Sites, every Site must join the Tenant Broker separately.
- Shared resources include RAS Secure Client Gateways (including HTML5 gateways) and High Availability Load Balancers (HALB).
- A Tenant Farm doesn't need its own RAS Secure Client Gateways and HALB. However, deployments with Gateways and HALB are possible if you need them for internal connections. For example, if you have different policies for internal and external connections, you might want to install a Gateway and HALB to serve local users.
- The network configuration of a Tenant requires the Tenant Publishing Agent to Tenant Broker Publishing Agent connectivity. Additionally, shared RAS Secure Client Gateways need to communicate with servers hosting published resources and the Tenant's Publishing Agent. Depending on the implemented network architecture, it might require a VLAN to VLAN connectivity, VPN, etc. These communications require only a limited number of open ports. For the complete list, see **Communication ports** (p. 317).
- Communications with a Tenant domain are always performed from a local Tenant Publishing Agent and never from the Tenant Broker infrastructure.
- Every Tenant must have a unique public domain address, which can be assigned a number of different ways. For example, a service provider can register a subdomain (e.g. Tenant1.Service-Provider.com) and assign it to a Tenant. Another approach could be using a private domain address (e.g. RAS.Tenant1.com) and have it routed to RAS Secure Client Gateways in the Tenant Broker. Note that different public domain addresses can resolve to the same IP address if needed.
- When a Tenant is joined to the Tenant Broker, shared RAS Secure Client Gateways become aware of the Tenant and its configuration and can connect to the Tenant's RAS Publishing Agent(s). A route must be set for the incoming Tenant's traffic from the Internet to RAS Secure Client Gateways (or HALB) in the Tenant Broker.
- Tenant Broker comes with its own RAS Console allowing you to manage shared resources, Tenant objects and certificates, monitor Tenant performance, and carry out standard RAS administration tasks.
- All Tenant Themes are made available in the Tenant Broker. When user connects via a shared RAS Secure Client Gateway in the Tenant Broker, the corresponding Tenant Theme is presented to the user.
- Different SSL certificates can be used for different Tenants.

Licensing

Tenant Broker doesn't need a license. Licenses are managed on a Tenant level.

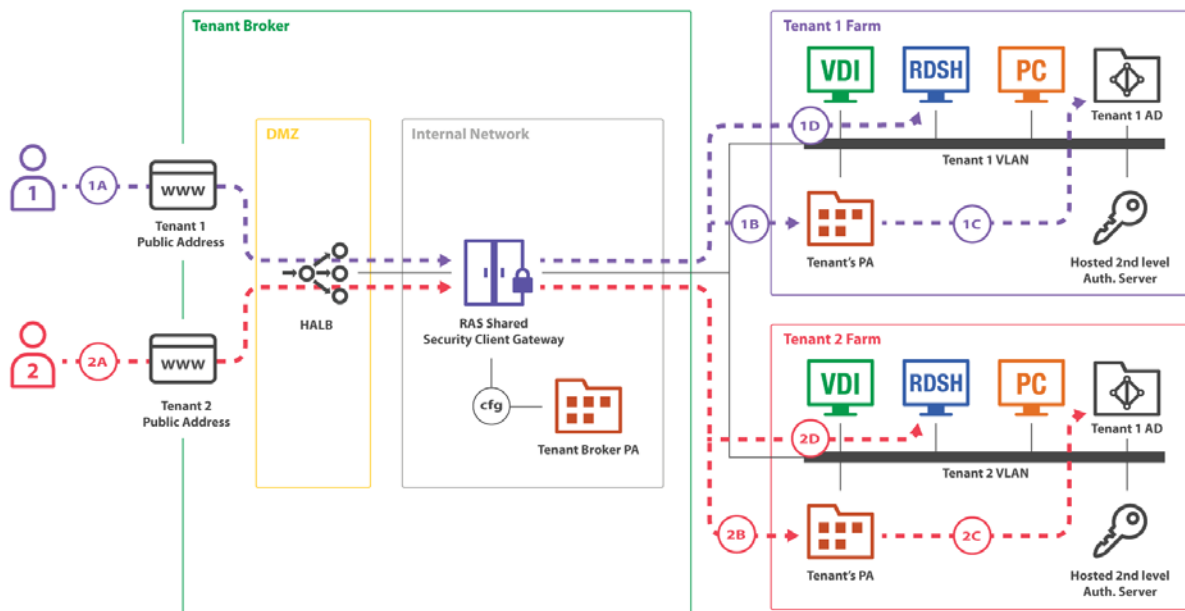
RAS version compatibility

Parallels RAS multi-tenant architecture is available in Parallels RAS 17.1 and newer. The following limitations apply when using older versions of Parallels RAS:

- Parallels Clients older than RAS 17.1 are incompatible with shared gateways and therefore cannot be used to connect to a Tenant Farm via the Tenant Broker.
- Parallels RAS installations older than RAS 17.1 are incompatible with Tenant Broker and cannot be joined as Tenants.

User connection flow

The following diagram illustrates the RAS user connection flow through Tenant Broker:



Shared RAS Secure Client Gateways installed in Tenant Broker are able to work with multiple concurrent user sessions in multiple Tenant farms. On the diagram above, you can see two users (1 and 2) connecting to different Tenant Farms (Tenant 1 Farm and Tenant 2 Farm). Both connections are tunneled through the same Gateway and then delivered to the correct Tenant Farm.

The connection flow consists of the following steps:

- 1 (1A), (2A) — A user initiates a RAS connection to a public address registered in the Tenant Broker. The (1A) connection goes to the Tenant 1 public address; the (2A) connection goes to the Tenant 2 public address.

- 2** (1B), (1C) — The shared Gateway makes a decision where to forward a user connection based on a hostname used in the initial connection (1A, 2A). After that each client establishes a RAS session with a Publishing Agent of their respective Tenant Farm. Tenant's Publishing Agent authenticates the user against Active Directory of the Tenant. After that, the user receives the list of published applications available to him or her.
- 3** (1D), (2D) — A user start a Remote User Session to a published application. The shared Gateway requests from Tenant's Publishing Agent an address of a server to forward the remote session to and forwards it.

The mapping of public addresses to Tenants is configured on shared Gateways by the Tenant Broker Publishing Agent.

Deploying Tenant Broker and Tenants

A typical scenario of deploying the multi-tenant architecture of Parallels RAS consists of the following steps:

- 1** Deploy Tenant Broker.
- 2** Deploy a traditional RAS Farm to operate as a Tenant.
- 3** Configure network between the Tenant Broker and the Tenant to allow the following connections:
 - Shared RAS Secure Client Gateways to Tenant RAS Publishing Agents.
 - Shared RAS Secure Client Gateways to resources hosts.
 - Tenant RAS Publishing Agents to Tenant Broker RAS Publishing Agent.

For the information about ports numbers, please see **Communication ports** (p. 317).

- 4** Create a Tenant object and a corresponding invitations hash in the Tenant Broker console, or create a secret key (more on this later in this chapter).
- 5** Join the Tenant to the Tenant Broker using the invitation hash or the secret key.
- 6** Assign a public domain address to the Tenant. This can be done at this point (after you join a Tenant) or it can be done in advance if you wish. Either way it has to be done or the clients will not be able to connect to the Tenant Farm.
- 7** Set up routing for incoming Tenant traffic from the Internet to shared RAS Secure Client Gateways and HALB.
- 8** Configure a certificate for the Tenant. By default, a self-signed certificate created during the installation will be used.
- 9** Test the client connectivity.

The subsequent sections describe the steps above in detail.

Deploying Tenant Broker

First you need to install Tenant Broker on a dedicated server. Please note that if you have Parallels RAS already installed on a computer where you are planning to install Tenant Broker, you need to uninstall it first. The two installation versions cannot coexist on the same machine.

To install Tenant Broker:

- 1 Run the standard Parallels RAS installer.
- 2 On the **Select Installation Type** page, select **Parallels RAS Tenant Broker**.
- 3 Click **Next** and follow the onscreen instructions.

Once the installation is finished, run the Parallels RAS Console.

When the console starts, you'll see that it has a different set of categories and managed objects compared to the standard RAS Console. The purpose of the Tenant Broker console is to manage shared resources and Tenants. It is not used to manage RD Sessions Hosts, VDI, or any other standard RAS resources because they are deployed and managed in individual Tenant Farms.

The Tenant Broker console

You can manage the following categories and object in the Tenant Broker console:

- **Farm.** This category allows you to manage Tenants, Gateways, Publishing Agents, HALB, and Certificates. The **Settings** subcategory allows you to manage global logging and the Tenant Broker itself.
- **Administration.** Allows you to perform management tasks similar to the standard RAS Console: Accounts, Settings, Mailbox, Reporting, Settings Audit.
- **Information.** Lists services and components running in the Tenant Broker and their status.

As with the standard RAS Console, every time you modify any of the objects, you need to click the **Apply** button for the changes to be saved in the configuration database.

Install RAS Secure Client Gateways

By default, Tenant Broker does not have any RAS Secure Client Gateways installed. To add a Gateway, log in to the Tenant Broker console, navigate to **Farm > Gateways** and click **Tasks > Add**. If you already have one or more RAS Secure Client Gateways, which are not used in any other RAS Farm, you can also add such a Gateway to the Tenant Broker. Please note that existing RAS Secure Client Gateway installations must be RAS version 17.1 or newer. Gateways from older RAS versions cannot operate as shared gateways.

To install a new gateway, run the Parallels RAS installer on a desired server, choose **Custom** and select the **RAS Secure Client Gateway** component. After the installation is finished, go back to the Tenant Broker console and add the gateway to the Tenant Broker.

Deploying a Tenant

A Tenant Farm is deployed just like a traditional Parallels RAS Farm. The only difference is, when installing the Farm, you don't need to install RAS Secure Client Gateways in it.

Note: If you decide to install a local (private) RAS Secure Client Gateway in a Tenant Farm (e.g. for local connections), you can do that, but please keep in mind that you cannot mix HALB and Gateways from the Tenant Broker and a Tenant Farm. The HALB appliance installed in the Tenant Broker will not support this scenario.

To set up a Parallels RAS Farm to be used as a Tenant:

- 1 Run the Parallels RAS installer.
- 2 On the **Select Installation Type** page, select **Custom**.
- 3 Click **Next**.
- 4 Make sure that the following components are selected for installation:
 - RAS Publishing Agent
 - Parallels RAS Console (optional; you can have the RAS Console installed on a different machine)

Other components are optional. You can install them now or you can install them later if needed.

- 5 Click **Next** and follow the onscreen instructions to complete the installation.

Join a Tenant to Tenant Broker

Once the Tenant Farm is operational, you can join one or more sites in it to the Tenant Broker.

Note: A Tenant is a Site in a separately deployed Parallels RAS Farm. When you join a Tenant to Tenant Broker, you join a Site. When you want to join the whole Farm, you do it one Site at a time. Of course, if you have just one Site in a Farm (and have no plans to create more sites), you are essentially joining the whole Farm.

There are two ways you can join a Tenant: (1) Using an invitation hash or (2) Using a shared secret key. The difference between the two is as follows:

- **Invitation hash.** An invitation hash is an automatically generated encrypted string that can be used to join a single Tenant to Tenant Broker. Invitation hash is a property of a Tenant object, which is created in the Tenant Broker console. You email the hash to the Tenant Farm administrator, so they can use it to join the Tenant Broker. Once used, an invitation hash cannot be used again by any other Tenant.

- **Shared secret key.** A shared secret key is similar to an invitation hash, with one important difference. It can be used to join an unlimited number of Tenants. A Tenant object is not pre-created for a secret key in the Tenant Broker. Instead, the object is created when the key is used to join a Tenant. Because of its unlimited usage capability, only the Tenant Broker admins should have access to a shared secret key. This scenario is useful when there are multiple Tenants, all managed by the same Tenant Broker administrator.

The invitation hash scenario is described below. For the secret key scenario see **Joining with a secret key** (p. 307).

First, you need to generate an invitation hash and create a Tenant object on the Tenant Broker side:

- 1 Log in to the Tenant Broker.
- 2 In the RAS Console, navigate to **Farm > Tenants**.
- 3 Click **Tasks > Add**.
- 4 In the **Tenant properties** dialog, specify the following:
 - **Name:** Type a Tenant name (this can be any name that you like).
 - **Public domain address:** If you've already assigned a public domain address to the Tenant, specify it here. If not, you can leave it blank. The address is not required for the Tenant to join the Tenant Broker. However, without the address specified here, end users will not be able to connect to the Tenant, so you will need to come back and fill it in later. For details, see **Assign a public domain address** (p. 309).
 - **Description:** Type an optional description.
 - **Publishing Agents:** This field is disabled and will be populated automatically when the Tenant joins the Tenant Broker. See more in **Tenant configuration** (p. 311).
 - **Tenant invitation hash:** This is the hash that the admin of the Tenant Farm will need to use to join the Tenant Broker. A hash is generated automatically when you open this dialog. To generate a new hash, click **Create new hash**.
 - **Send via email.** You can give the invitation hash to the Tenant admin directly or you can use this button to send it via email. When you click the button, you'll see a dialog where you can enter the recipients and where you can review and modify the email message. By default, the message contains instructions on how to join the Tenant Broker. Please note that SMTP settings must be configured in the RAS Console before you can use the email option. You can configure SMTP first and then return to this screen to complete this step.
- 5 Click **OK** to close the **Tenant properties** dialog. The new Tenant will appear in the **Tenants** list in the console. At this time, the Tenant is not joined yet. Read on to learn how to join it.

To join the Tenant to the Tenant Broker:

- 1 Log in to the Tenant Farm.
- 2 In the RAS console, navigate to **Farm > Site**. Note that you are joining a Site to the Tenant Broker, not the whole Farm, so if you have more than one Site, you need to join them one by one.

- 3 Click **Tasks > Join Tenant Broker**.
- 4 In the **Join Tenant Broker** dialog, enter the invitation hash that you obtained from the Tenant Broker in the previous steps (or, if you are an admin of a Tenant Farm, the one your received in the invitation email).
- 5 Click **Join**.

On successful join, you will see a message welcoming you to the Tenant Broker. If the primary Publishing Agent in your Tenant Farm can't reach the Tenant Broker, you will see a corresponding error message. Make sure that the Tenant Broker computer is reachable from the machine where you have the Tenant's RAS Publishing Agent running.

Overriding Tenant Broker IP address

The Tenant Broker IP address is detected automatically when you generate an invitation hash (or a secret key) and is embedded into the hash. If a Tenant can't reach the Tenant Broker using this address, you have the ability to override it as follows:

- 1 Log in to the Tenant Broker.
- 2 In the RAS Console, navigate to **Farm > Settings** and click the **Tenant broker** tab.
- 3 Select the **Override Tenant Broker address in tenant invitations and secret keys** option.
- 4 Enter the desired IP address in the field provided.

When done, the specified IP address will be used instead of the auto-detected address when generating an invitation hash or secret key. When the hash is used on the Tenant side to join the Tenant Broker, the Tenant will use this address to connect to the Tenant Broker.

Once used on the Tenant side, an invitation hash binds the Tenant Farm to the corresponding Tenant object in the Tenant Broker and the tenancy becomes effective.

Joining with a secret key

In addition to an invitation hash, you can join a Tenant to the Tenant Broker using a secret key. As described earlier (p. 305), a secret key can be used to join an unlimited number of Tenants to the same Tenant Broker.

To create a secret key:

- 1 Log in to the Tenant Broker console.
- 2 In the RAS Console, navigate to **Farm > Settings**.
- 3 Select the **Tenant broker** tab.
- 4 Select the **Allow RAS Farms to register in Tenant Broker using a secret key**.
- 5 The secret key is generated automatically. To generate a different key, click **Generate**.

Once you have the key, you can use it to join one or more Tenants to the Tenant Broker.

Note: Due to its unlimited usage capability, only the Tenant Broker administrator should have access to a shared secret key. Secret keys can be practical when the Tenant Broker administrator manages Tenant Farms, so instead of generating a hash for every Tenant, he/she can use a single secret key to join all of them to the Tenant Broker.

To join a Tenant using a secret key:

- 1 Log in to the Tenant Farm.
- 2 In the RAS Console, navigate to **Farm > Site**.
- 3 Click **Tasks > Join Tenant Broker**.
- 4 In the **Join Tenant Broker** dialog, enter the secret key and then enter a Tenant name of your choosing. The name you enter will be used in the Tenant Broker to name the corresponding Tenant object.
- 5 Click **Join**.

On successful join, you will see a message welcoming you to the Tenant Broker. If the primary Publishing Agent in your Tenant Farm can't reach the Tenant Broker, you will see a corresponding error message. Make sure that the Tenant Broker computer is reachable from the machine where you have the primary Publishing Agent running.

Overriding the Tenant Broker IP address

The Tenant Broker IP address is detected automatically when you generate a secret key and is embedded into it. If a Tenant can't reach the Tenant Broker using this address, you have the ability to override it as follows:

- 1 Log in to the Tenant Broker.
- 2 In the RAS Console, navigate to **Farm > Settings** and click the **Tenant broker** tab.
- 3 Select the **Override Tenant Broker address in tenant invitations and secret keys** option.
- 4 Enter the desired IP address in the field provided.

Verify join status

After you join a Tenant to the Tenant Broker, you should verify that the procedure was successful.

First, verify the Tenant Broker status in the Tenant console:

- 1 Log in to the Tenant Farm.
- 2 In the RAS Console, navigate to **Farm > Site** and select the **Site** tab in the right pane.
- 3 You should see the **Tenant Broker** section with the **Status** column, which should say **OK**. If the status is **Not verified**, make sure that the Tenant Broker is operational (or contact the Tenant Broker admin if you are not him or her).

You can also see additional Tenant Broker information by right-clicking it and choosing **Properties**. The information includes the following:

- **Name:** The Tenant Broker name.
- **Primary address:** The primary RAS Publishing Agent address.
- **Secondary address:** The secondary RAS Publishing Agent address (if available).

You should then verify the Tenant status in the Tenant Broker console:

- 1 Log in to the Tenant Broker.
- 2 In the RAS Console, navigate to **Farm > Tenants**.
- 3 In the **Tenants** tab, find the Tenant of interest and examine the **Status** column, which should say **OK** if the Tenant is joined properly. For other possible **Status** column values, see **Tenant configuration** (p. 311).

Configure network

After deploying a Tenant, you need to configure networking between Tenant Broker and Tenant in order to allow the following communications:

- Tenant Publishing Agent > Tenant Broker Publishing Agent: port 20003
- Tenant Broker Gateway > Tenant Broker Publishing Agent: port 20002
- Tenant Broker Gateway > Tenant Publishing Agent: port 20002
- Tenant Broker Gateway > Servers hosting published resources: port 3389

These are standard RAS ports, which are also described in the **Port reference** section.

Assign a public domain address

Every Tenant must have a unique public domain address for end users to connect to it through Tenant Broker. Although every Tenant must have a unique public domain address, it is not required for every Tenant to have a unique IP address. Different public domain address can be configured to resolve to the same IP address to reach the Tenant Broker shared Gateways. This way the Tenant Broker is still able to forward traffic to the right tenant based on the hostname requested by an end user.

A public domain address can be chosen a number of different ways. For example, a service provider can register a subdomain (e.g. Tenant1.Service-Provider.com) and assign it to a Tenant. Another approach could be using a private domain address (e.g. RAS.Tenant1.com) and have it routed to RAS Secure Client Gateways in the Tenant Broker. For testing purposes, you can even use an IP address.

The **Public domain address** is also a property of a Tenant object in the Tenant Broker console. After joining a Tenant to the Tenant Broker, you must ensure that this property contains the correct address. Otherwise end users will not be able to connect to the Tenant through the Tenant Broker.

To verify (and set if necessary) the Tenant's public domain address:

- 1 Log in to the Tenant Broker.
- 2 In the RAS Console, navigate to **Farm > Tenants**.
- 3 Right-click a Tenant and choose **Properties**.
- 4 In the **Properties** dialog, verify that the **Public domain address** field contains the correct address.

Configure an SSL certificate

The public domain address assigned to a Tenant must have a matching certificate. The Tenant Broker admin must create a certificate for every Tenant in the Tenant Broker console. Shared RAS Secure Client Gateways must then be configured to use these certificates. Tenant certificates are created and managed in Parallels RAS the same way as other certificates using the **Farm > Site > Certificates** subcategory. For the complete information about how to create certificates and how to assign them to RAS Secure Client Gateways and HALB, please see the **SSL Certificate Management** chapter (p. 258).

When a user connects to the Tenant's public domain address, a certificate with the common name matching the requested public domain address is selected automatically for every connection. The first available certificate is used which might not be the self-signed (say it was deleted)

If no matching certificate is found, the default self-signed certificate will be used, but the user will see a certificate warning in the web browser.

Set up routing for incoming traffic

One other thing that you have to do after you join a Tenant to the Tenant Broker, is set up routing for the incoming traffic from the Internet to shared RAS Secure Client Gateways or HALB.

User authentication

User authentication in the RAS multi-tenant architecture is performed by the RAS Publishing Agent running in the Tenant Farm. The Publishing Agent is selected randomly by a shared RAS Secure Client Gateway. If the Publishing Agent is unavailable, then it's marked accordingly and no communication is conducted with it from the same shared gateway for a period of time. The gateway checks the Publishing Agent status periodically and resumes communications as soon as the agent becomes available.

Unjoining from Tenant Broker

To unjoin a Tenant from the Tenant Broker, do the following:

- 1 Log in to the Tenant Farm.
- 2 In the RAS Console, navigate to **Farm > Site**.

3 Click **Tasks > Unjoin from Tenant Broker**.

The Tenant will be unjoined from the Tenant Broker. As a result, the Tenant users will no longer be able to connect to the Tenant Farm through the Tenant Broker.

Managing Tenants

In this section:

- Tenant configuration (p. 311)
- Deleting a Tenant object (p. 312)
- Opening a Tenant console (p. 312)

Tenant configuration

To see the list of existing Tenants in the Tenant Broker console, select **Farm > Tenants**.

The **Status** column indicates the Tenant status, which can be one of the following:

- **OK** — The Tenant has joined and has been verified.
- **Not Joined** — The Tenant object was created for the Tenant and the invitation hash was generated, but the Tenant has not joined the Tenant Broker yet.
- **Not Verified** — The Tenant has joined, but no connection to the Tenant's RAS Publishing Agent has been established yet. This status is usually displayed for a minute or so immediately after the Tenant joins the Tenant Broker. Once the connection is established, the status changes to **OK**.

This status can also appear when the Tenant Broker loses a connect with the Tenant's primary Publishing Agent. Shared gateways will be able to process connections only if they are still able to communicate with the Tenant's Publishing Agent on their own. They are independent from the Tenant Broker's Publishing Agent, but Tenant's Publishing Agent is still required to authenticate users.

- **Disabled** — The Tenant is disabled in the Tenant Broker configuration. You can enable and disable Tenant objects as described below.

To see and modify Tenant properties, click **Tasks > Properties** (or right-click > **Properties**). The **Properties** dialog opens where you can view and modify the following properties:

- **Enable Tenant:** Enable or disable the Tenant object in the Tenant Broker.
- **Name:** The Tenant name (must be unique).
- **Public domain address:** The unique address that end users connect to from the outside (e.g. RAS.tenant.com, tenant1.MSP-FARM.com, etc.). See more in **Assign a public domain address** (p. 309).

- **Clients in gateway mode connect to published tenant resources by server IP:** When selected, clients will use the Tenant IP address instead of the DNS name. You can use this option when a Tenant farm does not share the same DNS provider as the Tenant Broker farm.
- **Forward tenant sessions tunneled through gateway using server IP:** When a client session is forwarded to a server hosting published resources, either the server name (FQDN, hostname) or IP address can be used. When this option is selected (default) the IP address is used to forward the session internally. When the option is cleared, the configured host name is used.
- **Description:** An optional Tenant description. The Tenant description is a property that exists and can be viewed only in the Tenant Broker console.
- **Publishing Agents:** An IP address of one or more RAS Publishing Agents installed in the Tenant Farm. This is a read-only field.
- **Tenant invitation hash:** The hash that was used to join the Tenant to the Tenant Broker. This is a read-only field.

Deleting a Tenant object

A Tenant object can be deleted any time. To delete an object, click **Tasks > Delete** (or right-click > **Delete**). This deletes the Tenant configuration from shared RAS Secure Client Gateways, so no RDP sessions can be established from the gateway to the deleted Tenant anymore. The Tenant's RAS Console will show the Tenant Broker status as "Join Broken" after this. To completely remove any references to the Tenant Broker, the Tenant admin needs to unjoin the Tenant from the Tenant Broker (p. 310).

Opening a Tenant console

As a Tenant Broker admin, you can open the Tenant console right from the Tenant Broker console. To do so, navigate to **Farm > Tenants**, right-click a Tenant and choose **Open tenant console**. This will open a new instance of the RAS Console and will prompt you to log in to the Tenant Farm. Please note that the Tenant Farm must be configured to allow remote console connections, which means that the corresponding port must be open on the Tenant Publishing Agent and you need to know the credentials of the Tenant Farm administrator.

When you log in to a Tenant from the Tenant Broker console, the Tenant Farm is automatically added to the **Location** drop-down list (in the upper left-hand corner of the RAS Console window), so you can connect to the Tenant again by simply selecting it in the **Location** list.

Shared Gateways

All RAS Secure Client Gateways that exist in the Tenant Broker are shared among Tenants. For the most part, shared gateways operate similarly to standard RAS Secure Client Gateways but there are differences, which are described below.

Tunneling policies

Tunneling policies are allowed. Tunneled connections are sent to a Tenant Farm mapped to the public address used. The policies however are limited to "None" and "All servers in Site".

WYSE

WYSE is not supported.

Session counters

For each shared gateway, a session counter is displayed in the Tenant Broker console. To see how many sessions a gateway is running, navigate to **Farm > Site** and examine the **Sessions** column in the **Gateways** section.

Client connection routing

Each shared gateway is aware of a configuration of each existing Tenant and is able to route client connections to a correct RAS Publishing Agent running in a Tenant Farm. The routing works as follows:

- 1 A new client connection is established.
- 2 A shared gateway determines which Tenant the client belongs to based on the Tenant configuration.
- 3 The correct RAS Publishing Agent in the Tenant Farm is selected for this connection.
- 4 Two-factor authentication and application listing requests are forwarded to the selected RAS Publishing Agent. All subsequent client operations are also carried out using that Publishing Agent. See also **User authentication** (p. 310).

Shared gateway maintenance

When you need to take a shared RAS Secure Client Gateway offline for maintenance, you can do it the same way it's done in a traditional Parallels RAS Farm. You disable the gateway and wait for active sessions to drain. To see the number of active sessions for a gateway, navigate to **Farm > Site**. The session count is displayed in the **Sessions** column.

You can safely take shared Gateways offline. Parallels Clients will reconnect to the same sessions automatically.

Third-party network load balancers

Third-party network load balancers are possible to use with shared RAS Secure Client Gateways the same way they are used with traditional (not shared) RAS Secure Client Gateways.

HTML5 Client and Themes

One of the important features of the RAS multi-tenant architecture is the ability to use a shared HTML5 gateway (which is a part of the RAS Secure Client Gateway) for all browser-based client connections, while at the same time using tenant-specific HTML5 client Themes defined on the Tenant side. This allows Service Providers to implement white-labeling by creating unique custom Themes for individual Tenants.

An HTML5 client Theme is created in a Tenant Farm. The user interface and the functionality remain the same as with a traditional Parallels RAS Farm. When Tenants join the Tenant Broker, Themes are pulled from the Tenant's RAS Publishing Agent and added to the configuration of every shared RAS Secure Client Gateway.

When connecting to a Tenant Farm via the HTML5 gateway, a user must enter the Tenant public domain address (not the gateway address). The correct Theme is then used by the shared gateway as follows:

- The default Tenant Theme is used when the user enters the default URL: `https://<public-tenant-address>`.
- A specific Theme is used when the user adds the Theme name after the Tenant address: `https://<public-tenant-address>/<Theme-name>`

HTML5 configuration

The HTML5 Client is normally configured on the RAS Secure Client Gateway level (the **HTML5** tab in the gateway **Properties** dialog). When configuring a Theme, you have the ability to override the gateway settings by specifying them for a specific Theme in a Tenant Farm. To do so, in the Tenant RAS Console, select a Theme, open its properties and then select the **Gateway** category where you can specify your own settings. For more information, see **HTML5 Client Theme settings > Gateway** (p. 349).

Viewing Tenant Themes in Tenant Broker

If you are a Tenant Broker administrator, you can view Tenant Themes right in the Tenant Broker console:

- 1** In the Tenant Broker console, select **Farm > Tenants**.
- 2** Select a Tenant and click **Tasks > View tenant themes**.
- 3** The dialog opens where you can view Themes that were pulled from the Tenant and added to the configuration of every RAS Secure Client Gateway in the Tenant Broker.

Use this functionality to ensure that all Tenant Themes are properly synchronized on the Tenant Broker side, so when users connect to a Tenant through Tenant Broker, the appropriate Theme is used.

Monitoring Tenants

Parallels RAS Performance Monitor is a RAS component used to analyze Parallels RAS deployment bottlenecks and resource usage. RAS Performance Monitor can be used to monitor Tenants and view their performance metrics right from the Tenant Broker console.

To configure RAS Performance Monitor to collect information about Tenants:

- 1 Install RAS Performance Monitor as described in **Parallels RAS Performance Monitor** chapter (p. 422).
- 2 Log in to the Tenant Broker console.
- 3 In the console, navigate to **Administration > Reporting**.
- 4 Select the **Enable RAS Performance Monitor** option (the **RAS Performance Monitor configuration** section).
- 5 In the **Server** and **Port** fields, specify the name or IP address of the server where you have RAS Performance Monitor installed.
- 6 Click **Apply**.
- 7 Now open a Tenant console and repeat steps 3 to 6 above, so both Tenant Broker and the Tenant are configured to use the same RAS Performance Monitor. This way, when Tenant(s) report their performance data to the RAS Performance Monitor, it can be viewed on the Tenant Broker side.

Tenants will report statistics to RAS Performance Monitor and you can view these statistics in the Tenant Broker console. When viewing the data in the RAS Performance Monitor dashboard, you can switch between Farms and sites, so you can select a specific Tenant and view its performance metrics.

Tenant Broker compatibility and updates

When updating Parallels RAS to a newer version, the following rules apply to the RAS Multi-Tenant architecture:

- Parallels RAS Tenant Broker supports Tenants up to two major RAS versions older. For example, if you upgrade RAS Tenant Broker from RAS 17 to RAS 18 (or the next major version when it becomes available), it will support Tenants running Parallels RAS 17.
- When doing updates, you should first update the Tenant Broker, so your RAS Multi-Tenant installation remains fully operational. You can update Tenants later during their own maintenance window.

Upgrading from an older RAS version

If you have an existing Farm running RAS v16.x and would like to join it as a Tenant to Tenant Broker, follow these steps:

- 1 Upgrade the Farm to RAS 17.1 (or newer).
- 2 To join your Farm as a Tenant to Tenant Broker, follow the instructions in the **Deploying Tenant Broker and Tenants** section (p. 303).
- 3 Once the Farm is joined, you can remove local RAS Secure Client Gateways if you are not planning on using them for local connections. See **Implementation overview** (p. 300) for additional info.

Configuring notifications

System event notifications are used to alert RAS administrators about system events via email. You can configure system event notifications in **Farm > Site > Settings > Notifications**. For the complete description of this functionality, please see **System event notifications** (p. 440). The rest of this section describes notifications, which are specifics to Tenant Broker and Tenants.

Tenant event notifications

As a Tenant Broker administrator, you can receive notifications about the following Tenant events:

- **New Tenant enrollment.** Triggers when a new Tenant joins the Tenant Broker.
- **Tenant unjoins the broker.** Triggers when a registered Tenant unjoins the Tenant Broker.
- **Tenant status alert.** Triggers when the RAS Publishing Agent in a Tenant Farm goes offline.

When a Tenant event occurs, the Tenant Broker administrator receives an email containing the following information (depending on the event type):

- Tenant name.
- Tenant Broker name.
- Tenant enrollment method (invitation hash or secret key).
- Tenant status.
- Date.

To enable Tenant notifications, do the following:

- 1 Log in to the Tenant Broker.
- 2 In the RAS Console, navigate to **Farm > Site > Settings > Notifications**.
- 3 In the **Notification handlers** section, click **Tasks > New > Tenant events**.

- 4 In the **Tenant Events Notification Handler Properties** dialog, specify the following:
 - On the **General** tab, select the **Send email to RAS administrators** option and specify one or more email addresses separated by a semicolon.
 - On the **Settings** tab, either select the **Use the default settings** option (to use Site defaults) or clear it and specify your own settings.
- 5 Click **OK** to save your settings and close the dialog.

Tenant Broker event notifications

A Tenant Farm administrator can receive notifications when the Tenant Broker becomes unavailable. This usually happens when the RAS Publishing Agent in the Tenant Broker goes offline. The notification handler is configured the same way as described above, but this one is configured in the Tenant Farm (not the Tenant Broker).

Common event notifications

In addition to the **Tenant events** handler, you can configure notifications for common events, such as CPU utilization, Memory utilization, RAS Agent events, etc. The only limitation here when it comes to Tenant Broker is the Tenant Broker has a limited set of system events for which notification handlers can be configured (see the list of available handlers below). This is due to the fact that the Tenant Broker doesn't have RD Sessions Hosts, VDI provider, licensing limits, published resources, etc. A Tenant Farm has the complete set of notification handlers, so the Tenant admin can configure any of them.

The following notification handlers are available in the Tenant Broker:

- CPU utilization
- Memory utilization
- Number of gateway tunneled sessions
- Failed gateway tunneled sessions
- RAS Agent events

For additional information, please see **System event notifications** (p. 440).

Communication ports

Tenant Broker and Tenants communicate with each other using the following ports:

- Tenant Publishing Agent > Tenant Broker Publishing Agent: port 20003
- Tenant Broker Gateway > Tenant Broker Publishing Agent: port 20002
- Tenant Broker Gateway > Tenant Publishing Agent: port 20002
- Tenant Broker Gateway > Servers hosting published resources: port 3389

These are standard RAS ports, which are also described in the **Port reference** section.

SAML SSO Authentication

Parallels RAS 17.1 and newer support the Security Assertion Markup Language (SAML) authentication mechanism. SAML is an XML-based authentication that provides single sign-on (SSO) capability between different organizations by allowing user authentication without sharing the local identity database.

As part of the SAML SSO process, the new RAS Enrollment Server communicates with Microsoft Certificate Authority (CA) to request, enroll, and manage digital certificates on behalf of the user to complete authentication without requiring the users to put in their Active Directory credentials. Service providers and enterprises with multiple subsidiaries don't have to maintain their own internal Identity Management solutions or complex domains/forest trusts. Integrating with third-party Identity Providers allows customers and partners to provide end users with a true SSO experience.

In This Chapter

Introduction	319
System requirements.....	321
SAML basics	322
SAML configuration	323
Parallels Client configuration	341
Parallels client policy configuration.....	342
Test the SAML SSO deployment	342
Error messages	343

Introduction

Security Assertion Markup Language (SAML) is an XML-based authentication that provides single sign-on (SSO) capability between different organizations by allowing user authentication without sharing the local identity database. Parallels RAS 17.1 and newer support the SAML authentication mechanism.

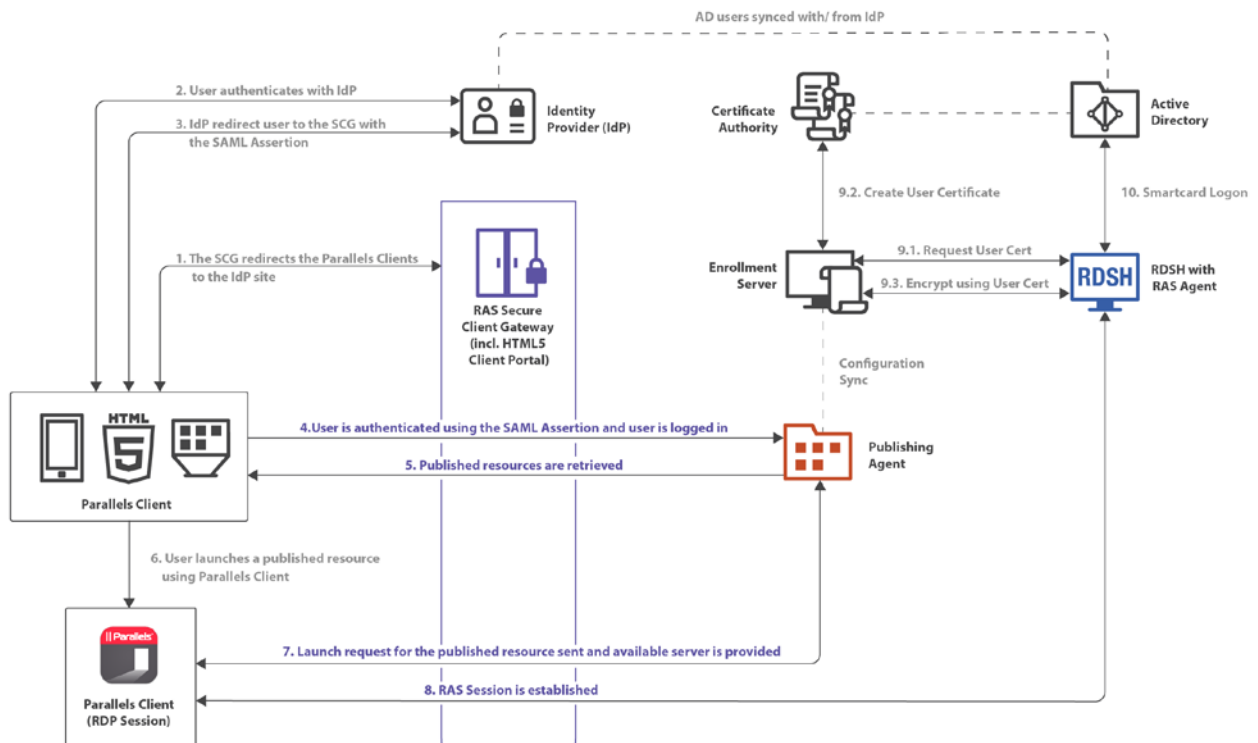
SAML (2.0) SSO was introduced in Parallels RAS 17.1 supporting HTML5 initiated authentication using HTML5 Client or Parallels Client for Windows. Parallels RAS 18 extends the client support for initiating SAML authentication using the default OS browser or the browser embedded in Parallels Client.

As part of the SAML SSO process, the new RAS Enrollment Server communicates with Microsoft Certificate Authority (CA) to request, enroll, and manage digital certificates on behalf of the user to complete authentication without requiring the users to put in their Active Directory credentials. Service providers and enterprises with multiple subsidiaries don't have to maintain their own internal Identity Management solutions or complex domains/forest trusts. Integrating with third-party Identity Providers allows customers and partners to provide end users with a true SSO experience.

Supported delivery options are:

- HTML5 Client
- HTML5 Client portal initiated SAML for Windows
- HTML5 Client Portal initiated SAML for Mac and Linux
- HTML5 Client Portal initiated SAML for Android, iOS and iPadOS
- Parallels Client for Windows initiated SAML authentication
- Parallels Client for Mac initiated SAML authentication
- Parallels Client for iOS/iPadOS initiated SAML authentication
- Parallels Client for Android initiated SAML authentication

The below high-level logical diagram depicts SAML authentication and login process within a Parallels RAS environment:



The SAML authentication and login steps on the diagram above are:

- 1 RAS Secure Client Gateway redirects the Parallels Client login request to the IdP site.
- 2 The user authenticates with IdP.
- 3 IdP redirects the user to the RAS Secure Client Gateway with the SAML Assertion.
- 4 The user is authenticated using the SAML Assertion and the user is logged in.
- 5 The list of the available RAS published resources is retrieved.
- 6 The user chooses a published resource and launches it from Parallels Client.
- 7 The launch request from the user is sent to the server side and the resource is started on the available server.
- 8 A Parallels RAS session is established.
- 9 User certificate is processed:
 - Certificate is requested.
 - Certificate is created.
 - Encryption is preformed using the certificate.
- 10 Smartcard logon.

System requirements

RAS Enrollment Server

- Windows Server 2008 R2 up to Windows Server 2022

RD Session Hosts

- Windows Server 2008 (x64 bit versions) up to Windows Server 2022

Desktop operating systems (guest VMs and Remote PCs)

- Windows 7 up to Windows 11

Please note that 32-bit operating systems are not supported.

Parallels Client

- Parallels Client version 18 is required.
- Supported platforms include Windows, Mac, Linux, iOS, Android.

SAML basics

Security Assertion Markup Language (SAML) is a standard for exchanging authentication information between identity and service providers. SAML authentication is a single sign-on mechanism where a centralized identity provider (IdP) performs user authentication, while the service provider (SP) only makes access control decisions based on the results of authentication.

The main benefits of using SAML authentication are as follows:

- Service providers don't need to maintain their own user databases. User information is stored in a centralized database on the identity provider side. If a user has to be added or removed, it only needs to be done in a single database.
- Service providers don't need to validate users themselves, so there's no need for a secure authentication and authorization implementation on the provider's side.
- Single sign-on means that a user has to log in once. All subsequent sign-ons (when a user launches a different application) are automatic.
- Users don't have to type in credentials when signing in.
- Users don't have to remember and renew passwords.
- No weak passwords.

The single sign-on process

SAML single sign-on can be initiated on the service provider side or on the identity provider side. The two scenarios are outlined below.

The SAML single sign-on process initiated on the service provider side consists of the following steps:

- 1** A user opens Parallels Client (one of the supported versions) (p. 321) and connects to the service provider.
- 2** The service provider sends a message to the identity provider, asking to authenticate the user.
- 3** The identity provider asks the user for a username and password.
- 4** If the user credentials are correct, an authentication response (assertion) is sent to the client and then passed to the service provider. The response contains a message that the user has logged in successfully; the identity provider signs the assertion.
- 5** The user is presented with the published applications list. When the user launches an application, there's no prompt for credentials.

Single sign-on can also be initiated on the identity provider side, in which case the basic steps are the following:

- 1 A user logs in to identity provider via a web browser and is presented with a list of enterprise applications, including Parallels RAS.
- 2 Once Parallels RAS is selected, the assertion is sent to the client, then passed to the service provider configured for Parallels RAS.
- 3 Users are presented with the RAS published applications list.
- 4 When the user launches an application, there is no prompt for credentials.

SAML configuration

In this section:

- Prerequisites (p. 323)
- IdP side configuration (p. 324)
- SP side configuration (RAS side) (p. 325)
- Active Directory user account configurations (p. 327)
- Configure certificate authority templates (p. 329)
- RAS Enrollment Server configuration (p. 338)
- RAS Enrollment Server high availability (p. 340)
- SAML integration examples and tips (p. 340)

Prerequisites

To configure SAML in Parallels RAS, you need the following:

- 1 Microsoft Active Directory with the following two user accounts present:
 - **Enrollment agent user:** used to enroll certificates through RAS Enrollment Server (ES) on behalf of the authenticated user.
 - **NLA User:** used to initiate the NLA connection with RD Session Hosts and/or VDI guests.

See **Active Directory user account configuration** (p. 327) for required permissions and delegations. Note that Azure Active Directory Domain Services (AADDs) are not supported to be used with SAML SSO.

- 2 Microsoft Enterprise Certification Authority (CA) including the following templates:
 - Enrollment Agent Certificate Template
 - Smartcard Logon Certificate Template

- 3 Third-party Identity Provider (IdP) such as Azure, Okta, Ping Identity, Gemalto SafeNet, and others. This is where the user accounts will reside. User accounts in IdP must be synchronized with the Microsoft Active Directory environment. Please consult with the provider on how to properly synchronize users.
- 4 Domain Controllers must have Domain Controller certificates. The certificates on the Domain Controllers must support smart card authentication. Certificates are created using the Microsoft CA certificate template named Domain Controller Authentication. Manually created Domain Controller certificates might not work. If you get an error "Request Not Supported", you may need to recreate Domain Controller certificates. Make sure RD Session Hosts and VDIs have the root certificate issued by the CA in the Trusted Root Certification Authorities store.
- 5 A Parallels RAS Farm with RD Session Host and/or VDI workloads (running on 64-bit OS).
- 6 For security reasons, the RAS Enrollment Server is recommended to be installed on a dedicated host. The host should be a standalone server that does not have any other components and roles installed.
- 7 Both SAML and RAS Enrollment Server configurations are Site-specific settings within the RAS environment. RAS administrators must have "Allow viewing of site information" and "Allow site changes" permissions delegated.

Note: Prerequisite knowledge of Microsoft Active Directory and Group Policy configuration is required for some of the above tasks.

Azure Active Directory Domain Services (AADDS) and Azure Virtual Desktop access are not currently supported with Parallels RAS SAML SSO.

IdP side configuration

On the identity provider side, you need to do the following:

- 1 Log in to your preferred IdP platform and create a generic or RAS specific SAML-based application to be used with the Parallels RAS environment.
- 2 Configure the application and take note of the following configuration properties to be added in Parallels RAS later:
 - **Entity ID**
 - **Logon URL**
 - **Logout URL**
 - **Certificate (base64)**
- 3 Alternatively, you can export a metadata file to be imported in Parallels RAS. For additional help, see **IdP example and tips**.

SP side configuration (RAS side)

On the service provider side (the Parallels RAS side), you need to enable Web (SAML) authentication and add the identity provider to the RAS Farm.

Enable Web (SAML) authentication

- 1 In the RAS Console, navigate to **Connection > Authentication**.
- 2 In the **Allowed authentication types** section, select the **Web (SAML)** option.

Adding an IdP to the RAS Farm

To add an IdP:

- 1 In the RAS Console, navigate to **Connection > SAML**. If the tab page is disabled, make sure you enabled Web (SAML). See above.
- 2 Click **Tasks > Add**.
- 3 In the **Add Identity Provider** wizard, specify a provider name.
- 4 In the **Use with Theme** drop-down box, select a Theme (p. 346) to which the IdP will be assigned. If you don't have a specific Theme yet, you can use the default Theme or you can select "<not used>" and assign a Theme later. Note that there can be multiple IdPs configured in the same RAS Farm. However, at this time, one IdP can be assigned to one Theme.
- 5 Select one of the following methods that the wizard will use to obtain the IdP information:
 - **Import published IdP metadata:** Import from an XML document published on the Internet. Specify the document URL taken from the IdP side configuration.
 - **Import IdP metadata from file:** Import from a local XML file downloaded from the IdP application. Specify the file name and path in the field provided.
 - **Manually enter the IdP information:** Select this option and then enter the information manually on the next wizard page.
- 6 Click **Next**.
- 7 If the configuration was imported in the previous step, the next page will be populated with data obtained from the XML file. If you've selected to enter the IdP data manually, you'll have to enter the values yourself:
 - **IdP entity ID:** Identity provider entity ID.
 - **IdP certificate:** Identity provider certificate data. To populate this field, you need to download the certificate from the IdP side, then open the downloaded file, copy its contents and paste it into this field.
 - **Logon URL:** Logon URL.
 - **Logout URL:** Logout URL.

Select the **Allow unencrypted assertion** option if needed.

Note: By default, the **Allow unencrypted assertion** option is disabled. Ensure that the IdP configuration is set to encrypt assertion or change the default setting within the RAS configuration.

- 8 At this point, you can configure service provider (SP) settings to be imported on the IdP side (IdP portal). You can do it now or you can do it later. To do it now, follow the steps below. To do it later, click **Finish** and then, when needed, open the identify provider object properties, select the **SP** tab and do the same steps as described below.
- 9 To configure SP settings, click the **Service provider information** button.
- 10 In the dialog that opens, enter the host address. The IdP will redirect to this address, which should be accessible from the end user browser.
- 11 The other fields including **SP Entity ID**, **Reply URL**, **Logon URL** and **Logout URL** are prepopulated based on the host address. The SP Certificate is autogenerated.
- 12 Next step is to complete the IdP configuration based on the values above. These values can be manually copied or exported as a metadata file (XML). Click the **Export SP metadata to file** link. Save the metadata as an XML file. Import the XML file into your IdP.
- 13 Close the dialog and click **Finish**.

Configuring user account attributes

When user authentication is performed by the IdP, user account attributes in Active Directory are compared with the matching attributes in the IdP user database. You can configure which attributes should be used for comparison as described below.

The following table lists available attributes:

RAS name	SAML name *	AD name	Description
UserPrincipalName	NameID	userPrincipalName	User Principal Name (UPN) is the name of a system user in an email address format.
Immutable ID	ImmutableID	objectGUID	A Universally Unique Identifier.
SID	SID	objectSid	An ObjectSID includes a domain prefix identifier that uniquely identifies the domain and a Relative Identifier (RID) that uniquely identifies the security principal within the domain.
sAMAccountName	sAMAccountName	sAMAccountName	The sAMAccountName attribute is a logon name used to support clients and servers from previous version of Windows, such as Windows NT 4.0 and others.
Custom	Email	Mail	A custom attribute that can be used to allow any SAML attribute name to match any AD attribute value. By default, it is the email address.

* The attributes in the **SAML name** column are editable and can be customized based on the IdP that you are using.

To configure attributes:

- 1 In the RAS Console, right-click an IdP that you've added in previous steps.
- 2 In the IdP **Properties** dialog, select the **Attributes** tab. On this tab, you can select or clear the attributes to be used for comparison or create custom ones:
 - Attributes that are selected will be compared for a match.
 - The names of all of the preconfigured SAML attributes (the IdP side) can be modified to match the AD attributes as required.
 - The custom attribute can be used to allow any SAML attribute name to match any AD attribute value. By default, it is the email address.
- 3 Configure and enable the desired attributes as needed based on the attributes configured on the IdP side.
- 4 Click **OK** to close the dialog.

Note 1: Multiple attributes are used in the presented order. If an attribute fails, the next configured attribute is used. Only one attribute is used at a time (in either/or fashion).

Note 2: If multiple AD users are configured with the same AD attribute value, user matching will fail. For example, if the email attribute is chosen and different AD users have the same email address, attribute matching between IdP account and AD User account will not be successful.

Attributes configuration tips

- When possible, use automation for user synchronization (such as Microsoft Azure AD Connect for Azure IdP configuration) between your Active Directory and the IdP to minimize user identity management overhead.
- Choose a user identification attribute that is unique to your environment, such as the User Principal Name (UPN) or Immutable ID (ObjectGuid) when possible. Alternatively, you can use other unique identifiers such as email address. In this case make sure that the **Email address** field in the user object in the AD is configured. If you use Microsoft Exchange Server, use the **Exchange Addresses** tab and Exchange policies.
- If using UPN as an attribute, you can also configure alternative UPN suffixes. This can be done from Active Directory Domains and Trusts (select root > right-click to open the **Properties** dialog). Once a new alternative UPN suffix is created, you can change the UPN on the user object properties from Active Directory Users and Computers.

Active Directory user account configuration

The enrollment agent user and NLA user must be created in Microsoft Active Directory. The following describes how to create these users.

Enrollment agent user account

The enrollment agent user account is required to enroll certificates through RAS Enrollment Server on behalf of the authenticated user. Please note that the enrollment agent user requires logon privileges on the machine where RAS Enrollment Server Agent is installed.

NLA user account

The NLA User is needed to initiate the NLA connection with RD Session Hosts and/or VDI guests. Please note that the NLA user requires log on privileges to the session host.

The NLA User must be a member of the Remote Desktop Users group and be granted the **Allow log on through Remote Desktop Services** permission. At the same time the NLA User must be prohibited to logon via Remote Desktop Services.

To exclude the NLA User account, it must be assigned the **Deny log on through Remote Desktop Services** user right.

To achieve both goals, you can use local or domain GPOs (linked to OU or domain wide).

A restart of the device is not required for this policy setting to be effective. Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy settings are applied through GPOs in the following order, which will overwrite settings on the local computer at the next Group Policy update:

- 1 Local policy settings
- 2 Site policy settings
- 3 Domain policy settings
- 4 OU policy settings

Create a new GPO or use Default Domain Policy GPO as follows:

- 1 Open the Group Policy Management Console (GPMC).
- 2 Open or create a GPO linked with the OU where the RDSH or VDI objects reside.
- 3 Navigate to **Computer Configuration > Windows Settings > Security Settings > Local Policies > User Rights Assignment** and open "Allow log on through Remote Desktop Services" option.
- 4 Choose to add User or Group..., add the NLA user and click **OK**.

Note: The option will override default settings (on workstation and servers: Administrators, Remote Desktop User; on domain controllers: Administrators) therefore do not forget to add the groups like local administrators group or domain admins group.

- 5 Navigate to **Computer Configuration > Windows Settings > Security Settings > Local Policies > User Rights Assignment** and open the **Deny log on through Remote Desktop Services** option.
- 6 Choose to add User or Group..., add the NLA user and click **OK**.

Configure certificate authority templates

In this section:

- Create an Enrollment Agent template (p. 329)
- Create a smartcard logon certificate template (p. 333)

Create an Enrollment Agent template

To create the Enrollment Agent template:

- 1 From the Certificate Authority server, launch the Certificate Authority management console (MMC) from Administrative Tools.
- 2 Expand the CA, right -click on the "Certificate Templates" folder and select **Manage**.
- 3 Right-click the Enrollment Agent template and choose **Duplicate Template**. The new template properties window opens. On the **General** tab, configure the following properties:
 - **Template display name:** PrIsEnrollmentAgent
 - **Template name:** PrIsEnrollmentAgent
 - **Validity period:** 2 years
 - **Renewal period:** 6 weeks
 - **Publish certificate in Active Directory:** ON
 - **Do not automatically re-enroll if a duplicate certificate exists in Active Directory:** OFF

Note: The display name can be any name you choose, however the template name must match the template name highlighted above.

The screenshot shows the 'Properties of New Template' dialog box with the following details:

- Subject Name:** PrisEnrollment Agent (highlighted with a red box)
- Server:** (empty)
- Issuance Requirements:** (empty)
- Superseded Templates:** (empty)
- Extensions:** (empty)
- Security:** (empty)
- Compatibility:** (empty)
- General:** (selected tab)
- Request Handling:** (empty)
- Cryptography:** (empty)
- Key Attestation:** (empty)
- Template name:** PrisEnrollmentAgent
- Validity period:** 2 years
- Renewal period:** 6 weeks
- Publish certificate in Active Directory (highlighted with a red box)
- Do not automatically reenroll if a duplicate certificate exists in Active Directory
- Buttons:** OK, Cancel, Apply, Help

4 Select the **Cryptography** tab and set the following values:

- **Provider category:** Legacy Cryptographic Service Provider (read-only).
- **Algorithm name:** Determined by CSP
- **Minimum key size:** 2048

In the section **Choose which cryptographic providers can be used for requests**, choose **Requests must use one of the following providers**. In the following list of providers, clear all options except **Microsoft Strong Cryptographic Provider** and set priority as the preferred provider:


[X] Microsoft Strong Cryptographic Provider

- [] Microsoft Enhanced Cryptographic Provider v 1.0
- [] Microsoft Base Cryptographic Provider v 1.0
- [] Microsoft Enhanced RSA and AES Cryptographic Provider

The screenshot shows the 'Properties of New Template' dialog box with the 'Security' tab selected. The 'Cryptography' sub-tab is active. The 'Provider Category' is set to 'Legacy Cryptographic Service Provider' and the 'Algorithm name' is 'Determined by CSP'. The 'Minimum key size' is set to '2048'. Under 'Choose which cryptographic providers can be used for requests', the radio button 'Requests must use one of the following providers:' is selected. The 'Providers' list includes 'Microsoft Strong Cryptographic Provider' (checked), 'Microsoft Base Smart Card Crypto Provider', 'Microsoft Enhanced Cryptographic Provider v1.0', 'Microsoft Base Cryptographic Provider v1.0', and 'Microsoft Enhanced RSA and AES Cryptographic Provider'. The 'Request hash' is set to 'Determined by CSP' and 'Use alternate signature format' is unchecked. Buttons for 'OK', 'Cancel', 'Apply', and 'Help' are at the bottom.

- 5 Select the **Security** tab and do the following:
- Click **Add**.
 - Add the enrollment agent user account.

- Allow (select) the "Read" and "Enroll" permission. Click **Apply** and **OK**.



Issue the certificate template

To issue the certificate template that you've created:

- 1 Run Certificate Authority again and right click on **Certificate Templates**, select new and click on **Certificate Template to Issue**.
- 2 Select the certificate template you've created in the previous steps (i.e. Prls Enrollment Agent) and click **OK**.
- 3 The certificate template should appear in the **Certificate Templates** list.

Note: After creating the Enrollment Agent template and the Smartcard Logon template (described later), you should restart the **Active Directory Certificate Services** service in Windows.

Create a smartcard logon certificate template

To create a smartcard logon certificate template:

- 1** From the Certificate Authority server, launch the Certificate Authority management console (MMC) from Administrative Tools.
- 2** Expand the CA, right-click on the "Certificate Templates" folder and select **Manage**.
- 3** Right click on the "Smartcard Logon" certificate template and then select **Duplicate**.
- 4** The new template properties open in the **General** tab. Type a template name in the text box. Note that the real name automatically appears in the second text box with no spaces. Remember this name. You will need it later to configure of SAML feature. The options on this tab should be configured as follows:
 - **Template display name:** PrlsSmartcardLogon
 - **Template name:** PrlsSmartcardLogon
 - **Validity period:** 1 years
 - **Renewal period:** 6 weeks
 - **Publish certificate in Active Directory:** OFF
 - **Do not automatically re-enroll if a duplicate certificate exists in Active Directory:** OFF

Note: The display name can be any name you choose, however the template name must match the template name highlighted above.

The screenshot shows a dialog box titled "Properties of New Template" with a close button (X) in the top right corner. The dialog has several tabs: "Subject Name", "Server", "Issuance Requirements", "Superseded Templates", "Extensions", "Security", "Compatibility", "General", "Request Handling", "Cryptography", and "Key Attestation". The "General" tab is currently selected. Inside the dialog, there are several input fields and controls:

- Template display name:** A text box containing "Pris Smartcard Logon", which is highlighted with a red rectangular border.
- Template name:** A text box containing "PrisSmartcardLogon".
- Validity period:** A dropdown menu showing "1" and "years".
- Renewal period:** A dropdown menu showing "6" and "weeks".
- Publish certificate in Active Directory
- Do not automatically reenroll if a duplicate certificate exists in Active Directory

At the bottom of the dialog, there are four buttons: "OK", "Cancel", "Apply", and "Help".

5 Select the **Cryptography** tab and set the following:

- **Provider category:** Legacy Cryptographic Service Provider (read-only).
- **Algorithm name:** Determined by CSP
- **Minimum key size:** 2048

In the section **Choose which cryptographic providers can be used for requests**, choose **Requests must use one of the following providers**. In the following list of providers, clear all options except **Microsoft Strong Cryptographic Provider** and set priority as the preferred provider:

Microsoft Strong Cryptographic Provider

- [] Microsoft Enhanced Cryptographic Provider v 1.0
- [] Microsoft Base Cryptographic Provider v 1.0
- [] Microsoft Enhanced RSA and AES Cryptographic Provider

Properties of New Template

Subject Name	Server	Issuance Requirements
Superseded Templates		Extensions
Security		Key Attestation
Compatibility	General	Request Handling
Cryptography		

Provider Category: Legacy Cryptographic Service Provider

Algorithm name: Determined by CSP

Minimum key size: 2048

Choose which cryptographic providers can be used for requests

Requests can use any provider available on the subject's computer

Requests must use one of the following providers:

Providers:

- Microsoft Strong Cryptographic Provider
- Microsoft Base Smart Card Crypto Provider
- Microsoft DH SChannel Cryptographic Provider
- Microsoft Enhanced Cryptographic Provider v1.0
- Microsoft Enhanced DSS and Diffie-Hellman Cryptographic Pr

Request hash: Determined by CSP

Use alternate signature format

OK Cancel Apply Help

- 6 Select the **Issuance Requirements** tab and set the following:
- **CA certificate manager approval:** OFF
 - **This number of authorized signatures:** 1
 - **Policy type required in signature:** Application policy
 - **Application policy:** Certificate Request Agent

- Same criteria as for enrollment: ON

PrIs Smartcard Logon Properties

Superseded Templates Extensions Security Server

General Compatibility Request Handling Cryptography Key Attestation

Subject Name Issuance Requirements

Require the following for enrollment:

CA certificate manager approval

This number of authorized signatures: 1

If you require more than one signature, autoenrollment is not allowed.

Policy type required in signature:

Application policy

Application policy:
Certificate Request Agent

Issuance policies:

Add...
Remove

Require the following for reenrollment:

Same criteria as for enrollment

Valid existing certificate

Allow key based renewal (*)


Requires subject information to be provided within the certificate request.

* Control is disabled due to [compatibility settings](#).

OK Cancel Apply Help

- 7 Select the **Security** tab and do the following:
- Click **Add**.
 - Add the enrollment agent user account.

- Allow (select) the "Read" and "Enroll" permissions. Click **Apply** and **OK**.



Issue the certificate template

To issue the certificate template that you've created:

- 1 Run Certificate Authority again and right click on **Certificate Templates**, select new and click on **Certificate Template to Issue**.
- 2 Select the certificate template you've created in the previous steps (i.e. Prls Smarcard Logon) and click **OK**.
- 3 The certificate template should appear in the **Certificate Templates** list.

Note: After creating the Smartcard Logon template and the Enrollment Agent template (described earlier), you should restart the **Active Directory Certificate Services** service in Windows.

RAS Enrollment Server configuration

RAS Enrollment Server communicates with Microsoft Certificate Authority (CA) to request, enroll, and manage digital certificates on behalf of a user for SSO authentication in the Parallels RAS environment.

Note: For security reasons, RAS Enrollment Server should be installed on a secure, dedicated server similar to an Active Directory Domain Controller or Certificate Authority with no other Parallels RAS components installed.

Setup and configure RAS Enrollment Server

You can remotely install the RAS Enrollment Server Agent on a specified server from the RAS Console. You can also install the Agent by running the standard RAS installer on the desired server.

To remotely install the RAS Enrollment Server:

- 1 In the RAS Console, navigate to **Farm > Site > Enrollment servers**.
- 2 Click **Tasks > Add**.
- 3 Specify the FQDN or IP address of the server where you want the RAS Enrollment Server Agent to be installed.
- 4 Click **Next**.
- 5 In the **Enrollment Server Agent Information** dialog, click **Install** and follow the onscreen instructions.

To install the RAS Enrollment Server using the Parallels RAS installer:

- 1 Run the Parallels RAS installer on the server where you want the RAS Enrollment Server Agent installed.
- 2 On the **Select Installation Type** page, select **Custom** and click **Next**.
- 3 Clear all other components and select the Parallels RAS Enrollment Server component.
- 4 Click **Next** and follow the onscreen instructions.
- 5 Once the RAS Enrollment Server is installed, open the RAS Console and navigate to **Farm > Site > Enrollment servers**.
- 6 Click **Tasks > Add**.
- 7 Enter the Enrollment Server FQDN or IP address and click **Next**.
- 8 Follow the onscreen instructions to add the server to the Farm.

Obtain and copy the registration key

If you perform a manual installation using the RAS installer, it is necessary to place a registration key file on the Enrollment Server host. This step is not required if the RAS Enrollment Server Agent was remotely deployed from the RAS Console.

First, you need to obtain the registration key file as follows:

- 1 Open the RAS Console and navigate to **Farm > Site > Enrollment servers**.
- 2 Click **Tasks > Export registration key**.
- 3 Save the key to a file named *registration.crt*.

Once you have the registration.crt file, copy it to the following folder on the server where you have the RAS Enrollment Server installed, by default in the following path:

```
C:\Program Files (x86)\Parallels\ApplicationServer\x64
```

Note: It is mandatory for the registration key file to be named "registration.crt".

Configure AD integration

After you added the RAS Enrollment Server in the RAS Console, you need to configure AD integration for it as follows:

- 1 In the RAS Console, navigate to **Farm / Site > Enrollment Servers**.
- 2 Select the **AD Integration** tab.
- 3 In the **Certificate authority (CA)** section, specify the configuration string of your Enterprise CA where the new certificate templates, (Pris Enrollment Agent and Pris Smartcard Logon) were created. This should be done in the following format:
`C\hostname.domain\issuing CA name`
Alternatively, you can click the [...] button to select a CA. For configuration details, see **Configure certificate authority templates** (p. 329).
- 4 In the **Enrollment Agent** section, specify the Enrollment Agent username and password. For configuration details, see **Active Directory user account configuration** (p. 327).
- 5 In the **NLA user** section, specify the NLA username and password. For configuration details, see **Active Directory user account configuration** (p. 327).
- 6 Click the **Validate AD integration settings** button to make sure that the information you've entered is valid.

Using computer management tools

You can perform standard computer management tasks on a RAS Enrollment Server host right from the RAS Console. These include Remote Desktop Connection, PowerShell, Computer Management, Service Management, Event Viewer, IPconfig, Reboot, and others. To access the **Tools** menu, click **Tasks > Tools** and choose a desired tool. For requirements and usage information, see **Computer management tools** (p. 431).

RAS Enrollment Server high availability

For high availability, multiple Enrollment Servers (ESs) can be added to each Site. All enabled and verified ESs will be used in an active/active fashion. Upon user login, requests from workload VMs such as RD Session Hosts or VDIs are equally distributed among the available ESs. In case of failures on a particular ES, the next available ES is selected and the SAML SSO authentication process continues. Specifically required for manual deployment of multiple ESs, it is important to note that all ESs in the same site share the same registration key which is required to be deployed in the specified path as mentioned in the **RAS Enrollment Server configuration** (p. 338) section.

Note: Multiple ESs do not share a common certificate repository store and all certificates are segregated on each ES. This means that in case of multiple ESs, same user might have different certificates available on different ESs.

SAML integration examples and tips

For examples of how to integrate various Identity Providers with Parallels RAS, please read the **SAML SSO Authentication Examples** guide, which is available on the Parallels website at <https://www.parallels.com/products/ras/resources/>.

User account attributes

When user authentication is performed by the IdP, user account attributes in Active Directory and the IdP are compared with each other for a match. Attributes to be compared are configured on the IdP and in the RAS Console. For details, see **SP side configuration (RAS side)** (p. 325).

Security tip

For security reasons, it is advisable to configure enrollment agent restrictions for a CA to allow only the newly created Enrollment Agent User permissions to enroll certificates on behalf of the users. To do so, follow the steps below.

- 1 Open the Certification Authority snap-in, right-click the name of the CA, and then click **Properties**.
- 2 Click the **Enrollment Agents** tab, click **Restrict enrollment agents**, and click **OK** on the message that appears.

- 3 Under **Enrollment agents**, click **Add**, type the name of the Enrollment agent user created in the previous steps and then click **OK**. Click **Everyone**, and then click **Remove**.
- 4 Under **Certificate Templates**, click **Add**, select the templates that were created (Pris Enrollment Agent and Pris Smartcard Logon) and then click **OK**. When you have finished adding the names of certificate templates, click **<All>**, and then click **Remove**.
- 5 Under **Permissions**, click **Add**, type the names or groups, which are the users or group expected to login to the RAS environment using SAML, and then click **OK**. Click **Everyone**, and then click **Remove**.
- 6 If you want to block the enrollment agent from managing certificates for other users, computers, or groups, under **Permissions**, select this user, computer, or group, and then click **Deny**.
- 7 When you are finished configuring enrollment agent restrictions, click **OK** or **Apply**.

Note: The user or group that you applied enrollment agent restrictions to must have a valid enrollment agent certificate for the CA before they can act as an enrollment agent, whether restricted enrollment agent permissions have or have not been configured.

Parallels Client configuration

Parallel HTML5 Client

No additional configuration is needed.

Parallels Client for Windows

To configure Parallels Client for Windows for SAML SSO authentication:

- 1 Select a connection (or create a new one) and open its properties.
- 2 In the **Login** section, select **Web** as authentication type.
- 3 Select the **Advanced Settings** tab and click the **Connection Advanced Settings** button.
- 4 In the **Web authentication** section, select or clear the **Use default OS browser** option. If the option is selected, the SAML SSO login dialog will open in the default browser. If the option is cleared, the browser built into the Parallels Client will be used.
- 5 Close all dialogs to save the connection properties.

When connecting to Parallels RAS, a dialog will open in a web browser asking the user to enter credentials to be verified by the identity provider. If the credentials are valid, the list of published applications will appear in the Parallels Client.

Parallels Client for Mac

To configure Parallels Client for Mac for SAML SSO authentication:

- 1 Select a connection (or create a new one) and open its properties.
- 2 On the **Connection** tab, in the **Login** section, select **Web** as authentication type.
- 3 Select the **Advanced** tab and select or clear the **Use default OS browser** option (**Web authentication** section). If the option is selected, the SAML SSO login dialog will open in the default browser. If the option is cleared, the browser built into the Parallels Client will be used.
- 4 Close the dialog to save the connection properties.

When connecting to Parallels RAS, a dialog will open in a web browser asking the user to enter credentials to be verified by the identity provider. If the credentials are valid, the list of published applications will appear in the Parallels Client.

Parallels client policy configuration

In addition to specifying SAML SSO options directly on the client side, you can also set them via the Parallels Client Policy in the RAS Console. To do so:

- 1 In the RAS Console, select the **Policies** category.
- 2 Open the policy properties and navigate to **Session > Connection > Primary Connection**.
- 3 In the **Authentication type** drop-down list, select **Web**.
- 4 Navigate to **Session > Connection > Web authentication**.
- 5 Select or clear the **Use default OS browser** option. If the option is selected, the SAML SSO login dialog on the client side will open in the default browser. If the option is cleared, the browser built into the Parallels Client will be used.

See also **Parallels Client configuration** (p. 341).

Test the SAML SSO deployment

When you have the SAML SSO authentication configured, you can test it as described below.

Service provider initiated authentication

- 1 Use a web browser to open the RAS HTML5 Client (or use a platform-specific Parallels Client), specifying the Theme to which you assigned the identity provider.
- 2 Start a published application. Check that the application session was started successfully.

- 3 Create one more Theme, add one more IdP providers and then connect specifying the new Theme. Launch another application.

Identity provider initiated authentication

- 1 Use the web browser to connect to the identity provider portal.
- 2 Start a published application. Check that the application session was started successfully.

Error messages

Error messages appear in the web browser when something goes wrong with SAML SSO authentication.

Pre HTML5 loading

Error message	Notes
Unable to parse SAML Assertion	<p>There was an error while parsing and validating the SAML Assertion. Further details can be found in HTML5 Logs.</p> <p>Most common causes:</p> <p>SAML Response is not valid for this audience: The most probable cause for this issue is having wrong configuration on the IDP, especially the Entity ID URL. The entity ID URL in the assertion will not match with the Entity ID provided in the SP SAML settings.</p> <p>Expected 1 Assertion or 1 EncryptedAssertion; found 0: The Assertion / EncryptedAssertion tag was not found in the response. The HTML5 Client will be expecting an encrypted assertion while the IDP is sending a non encrypted one. This can either be fixed by changing the IDP settings to send an encrypted assertion or tick the checkbox found in 'RAS Console > Connection > SAML > IDP Settings > Allow unencrypted assertion'</p> <p>SAML Response is not yet valid: This might happen if the time of the server where RAS Gateway is installed is incorrect, for instance 4 seconds behind. In this case the assert will be created before actually trying to parse it.</p> <p>SAML Response is no longer valid: This might happen if the time of the server where RAS Gateway is installed is incorrect. In case it's manually set in the future, assert might be seen as not valid anymore while trying to validate it.</p>
SAML Assertion body is empty	SAML Assertion was not found in the response. Further details can be found in HTML5 Logs
Unable to create SAML logout request	There was an error while creating SAML logout request. Further details can be found in HTML5 Logs.
Unable to create SAML logout response	There was an error while creating logout response. Further details can be found in HTML5 Logs.

Post HTML5 loading

Error code	Error message	Notes
0x00000029	SAML IdP settings not found. IdP Id:'xxx'	Check the Identity Provider settings. Check if the IdP metadata are correctly imported.
0x0000002A	SAML IdP info keys loading failed. IdP Id:'xxx'	Check if the IdP certificate is present in the IdP settings.
0x0000002B	SAML Theme mismatch	Check if the theme is correctly set in the IdP settings.
0x0000002C	Logon using SAML failed. Error: 0x00001	See errors below
0x00000029	No Enrollment Sever available	Check Enrollment server(s) status
0x0000002A	Missing NLA User Configuration	Enter NLA User details
0x00000003	Logon using SAML failed. Error: Failed to match AD User. 0x00000006	Check if the Attributes settings are correct in the IdP properties.
0x00000003	Logon using SAML failed. Error: Failed to validate and decrypt the response. 0x00000009	Check if the IdP certificate is present in the IdP settings.
0x00000003	Logon using SAML failed. Error: Assertion not encrypted. 0x0000001C	Check if the IdP settings for the logon request are correct.
0x00000003	Logon using SAML failed. Error: Failed to decrypt the assertion. 0x0000001D	Check the SP certificate is correctly set in the IdP settings.
0x00000003	Logon using SAML failed. Error: Failed to verify assertion. 0x0000001F	Check if the IdP certificate is present in the IdP settings.

Once an application or desktop is launched

Error message	Description and reference
Invalid username or password	The user certificate is valid, but the domain controller did not accept it. Check the Kerberos logs on the domain controller.
The system could not log you on. Your credentials could not be verified.	Check connectivity with the domain controller and check that the appropriate certificates installed.
The request is not supported	The "Domain Controller" and "Domain Controller Authentication" certificates on Domain Controller require enrolling, even if they are already available.
The system could not log you on. The smartcard certificate used for authentication was not trusted.	The intermediate and root certificates are not installed on the machine where the error is shown. The CA root certificate and any intermediate certificates must be added to the "Trusted root certificates" in the local computer account.
You cannot logon because smart card logon is not supported for your account.	The user account has not been fully configured for smart card logon.
No valid smart card certificate could be found.	Check the configuration of the PrIsSmartcardCertificate. The extensions might not be set correctly, or the RSA key is less than 2048 bits.
Bad Request	Check the configuration of the PrIsSmartcardCertificate. The extensions might not be set correctly, or the RSA key is less than 2048 bits.

Parallels HTML5 Client

Parallels HTML5 Client is a client application that runs in a web browser. Users can use Parallels HTML5 Client to view and launch remote applications and desktop from a web browser.

Compared to platform-specific Parallels Clients (Parallels Client for Windows, Parallels Client for iOS, etc.), Parallels HTML5 Client does not require end users to install additional software on their computers or mobile devices. Feature-wise, platform-specific Parallels Clients give users more options than Parallels HTML5 Client. Nonetheless, Parallels HTML5 Client is a fully-featured platform-independent client providing end users with an alternative method of working with remote resources published via Parallels RAS.

Please note that the RAS HTML5 Gateway (the server side) requires Windows Server 2008 R2 or higher (it will not work on Windows Server 2008). The only requirement for the client side is an HTML5-enabled web browser.

In This Chapter

Configure HTML5 Client	345
Configure Themes	346
Open Parallels HTML5 Client	353
Main menu options.....	355
Running remote applications and desktops	357
Auto login.....	359
Direct App access.....	360
Using the toolbar	361

Configure HTML5 Client

The HTML5 Client is a part of RAS Secure Client Gateway. To be used by end users, the HTML5 Gateway must be enabled and configured in the RAS Console as described in **Configure HTML5 Gateway** (p. 74).

Session persistence based on a cookie

RAS HTML5 Client session persistence is normally set by user's IP address (source addressing). If you can't use source addressing in your environment (e.g. your security policy doesn't allow it), you can use the Session Cookie to maintain persistence between a user and a server. To do so, you'll need to set up a load balancer that can use a session cookie for persistence. The cookie that you should use is ASP.NET_SessionId. If you are using a load balancer that doesn't use ASP.NET, you can specify a different cookie on the Web Requests tab of the RAS Secure Client Gateway Properties dialog. For more information, see **Web request load balancing** (p. 78).

Configure Themes

Themes in Parallels RAS is a functionality that allows you to do the following:

- Allow access to a Theme to specified groups of users while configuring certain Theme properties that will apply to these groups. This functionality is supported by Parallels Client on all available platforms.
- Customize the appearance of Parallels Client, which enables you to implement custom branding of Parallels Client for different groups of users. Note that this functionality is only available for RAS HTML5 Client and Parallels Client for Windows.

To manage Themes, in the Parallels RAS Console, navigate to **Farm** > <Site> > **Themes**. The **Themes** view in the right pane displays the available Themes. The list contains at least one default Theme. This Theme cannot be removed but you can customize it as needed. In addition to the default Theme, you can create your own Themes.

To create a new or modify an existing Theme:

- Click **Tasks** > **New Theme** (or click the **[+]** icon) to create a new Theme.
- Double-click an existing Theme (or right-click it and choose **Properties**).

The **Theme Properties** dialog opens. Use the dialog to create a new or modify an existing Theme. The instructions in the subsequent sections apply to both scenarios.

General Theme settings

The Theme settings described below apply to Parallels Client for all available platforms.

General

Select **General** in the left pane and specify the following Theme properties:

- **Enable Theme:** Enable or disable the Theme (the default Theme cannot be disabled).
- **Name:** Specify a Theme name.

- **Description:** Specify an optional Theme description.
- **Override authentication domain:** Allows you to specify a domain name that will be passed to Parallels Client, so users don't have to enter it manually. This setting overrides the domain name setting in **Connection > Authentication**. For more information, see **Allowing users to change domain password** (p. 286).
- **Limit access to this Theme to members of these Active Directory groups:** If this option is cleared, any Parallels RAS user can access the Theme if they know its URL. To limit access to a particular group (or groups), select this option and then click **Tasks > Add** (or click the **[+]** icon) and select the desired group(s).

Message

Select **Messages** in the left pane and specify a post-logon message (up to 500 characters). The post-logon message appears as a message box immediately after the user successfully logs in. The message can be overridden for HTML5 client and Windows client individually (see **Messages** for each client in the subsequent sections).

HTML5 Client Theme settings

The **HTML5 client** category allows you to configure Theme settings for Parallels HTML5 client. These settings affect how the HTML5 client looks and behaves in a web browser.

Note: To see how your HTML5 client Theme looks, click the **Preview HTML5 Theme** button in the lower left-hand corner of the dialog at any time.

URLs

The **URLs** category is used to specify the Theme login page URL and add additional URLs to the HTML5 Client page:

- **Theme login page:** Specifies a postfix for the Theme login page URL. This field is populated automatically with the Theme name when you save it, but you can specify a name of your choice. The complete URL of the Theme login page is comprised of "https://<host-name>/" followed by the name specified in this field. For the explanation of what the <host-name> should be, please see **Web request load balancing** (p. 78).

For example, if you name the Theme "Theme-S1", the complete URL is https://<host-name>/Theme-S1. When you save the Theme, the URL is displayed on the **Themes** tab in the RAS Console (the **HTML5 URL** column).

Please note that the URL described above is the short version, which is easier to remember and use. The full version is:

https://<host-name>/RASHTML5Gateway/?theme=<team-name>

Both the short and the long versions are equally valid.

- **Show Parallels Client download URL.** If selected, users will see the **Download Client** link on the HTML5 client page, which can be used to download, install, and configure Parallels Client on users' computers.
- **Override download URL for branded Parallels Client (Windows):** Specifies a location from which your Windows users will download Parallels Client for Windows. By default, Parallels Client is downloaded from the Parallels web site. If you use a branded version of Parallels Client, you can specify its location in this field.
- **Footer URLs.** This option allows you to specify custom URLs that will be placed in the HTML5 client footer. To add a URL, click **Tasks** > **Add** and specify a URL, a text that will appear on the page footer, and a tooltip text. When entering similar URLs, you can duplicate an existing one by right-clicking it and choosing **Duplicate** (or select an entry and click the "duplicate" icon next to the **[-]** icon). If you've added multiple URLs, you can reorder them by clicking the up or down arrow icons or selecting **Up** or **Down** items in the **Tasks** menu. The URLs will appear in the footer in the order listed (you can click the **Preview HTML4 Theme** button to see how it looks).

Branding

The **Branding** category allows you to customize the appearance of HTML5 Client pages.

The following properties can be customized:

- **Webpage title:** Specifies the title that appears on the webpage. You can type any title you like.
- **Login to:** Specifies a name that will appear in the HTML5 Client login dialog. For example, if you type "ABC" here, the login page will say, "Log in to ABC". There are two predefined variables that you can use here: %FARM% (the actual Farm name; this is the default value) and %SITE% (the Licensing Site name).
- **Company logo:** Displays the image which is displayed on the HTML5 client page header. To change the image, select browse and then specify the image file. Note that changing the logo image also removes the default **Remote Application Server** part from the page header.
- **Favicon icon:** Displays the currently set favicon icon. To change the icon, click **Browse** and select an icon file.

Colors

Specify the desired colors for various HTML5 Client elements, such as header, footer, work area, buttons, etc.

To change a color, do one of the following:

- Click the **Change color** button and choose a color.
- Click the "pencil" icon in the HEX column and type a HEX value of the desired color.

Language bar

Select languages that will appear in the language selector on the Parallels HTML5 Client page. The selector appears as a language flag icon on the page header to the right of the user name.

Messages

On this pane, you can specify pre-logout and post-logout messages:

- A pre-logout message will appear on the **Log in** page.
- To override the default post-logout message (see **Messages** in the beginning of this topic), select the **Override post-logout message** option and enter a message.

The messages must be 500 characters or less.

Input prompt

Input prompts specified here will appear on the login page to help users enter their username and password correctly in the fields provided. For example, the default **user@domain** login prompt will appear as a light gray text in the login field, hinting the user that they should enter their name in the UPN format. Predefined input prompts are provided for every supported language. You can specify your own prompts if needed.

Gateway

The **Gateway** category can be used to override the default HTML5 Client settings, which are configured for the RAS Secure Client Gateway. Normally, you shouldn't override the gateway settings if you are running a traditional Parallels RAS Farm and using a single Theme in a Site. Scenarios when overriding the settings may be needed include the following:

- You have multiple Themes for different groups of users and would like different Themes to behave differently in terms of application launching methods and restrictions.
- You are using RAS multi-tenant architecture where RAS Secure Client Gateways are running in Tenant Broker and are shared by Tenants, which are separate Farms. Themes in this kind of deployment are defined on the Tenant level, so each Tenant can have its own HTML5 Client look and feel. Since gateways are shared by Tenants, it is logical to configure these settings on a Theme level, which is exactly what the **Gateway** category allows you to do. For the complete description of what Tenant Broker and Tenants are, please read the **RAS Multi-Tenant Architecture** chapter (p. 299).

To override the RAS Secure Client Gateway settings, select the **Override gateway settings for the Theme** option and then specify your own settings. For the description on how to configure these settings, see **Configure HTML5 Client (p. 74)**.

Legal policies

Cookie consent

Select the **Enable cookie consent** option to show a notification about Parallels HTML5 Client cookie policy to users on first time use. This provides users with information regarding the use of cookies and the option to accept.

End User License Agreement

Select the **Enable EULA** option to show the Parallels End User License Agreement (EULA) to users on first time use. User intervention is required to read and accept the agreement to complete the login process.

Parallels Client for Windows Theme settings

Panes under the **Windows client** heading allow you to configure Theme settings for Parallels Client for Windows. By configuring a Windows client Theme, you can make the client appear to end users as your organization requires.

Branding

On the **Branding** pane, specify the following:

- **Company name:** Used to create the Start menu hierarchy: Start \ Company Name \ App Name.
- **Application name:** Displayed in the app caption and the Start menu entry name.
- **Connection banner:** Displayed when a connection is being established.
- **Application icon:** The application icon used for the Start menu and by the main app window.

Messages

To override the default post-logon message, select the **Override post-logon message** option and enter a message.

Custom menu

The **Custom Menu** pane allows you add a menu item to the **Help** menu in white-labeled Parallels Client for Windows. For example, if you enter "&Notepad" in the **Menu item** field and "notepad.exe" in the **Command** field, a new menu item will appear under the **Help** menu in every white-labeled Parallels Client for Windows connecting to this Farm. The item will be named **Notepad** (with the "N" being the shortcut) and it will open the Notepad.exe application when clicked. The **Command** field can contain an executable name, a URL, or any other command that can be properly executed on a Windows machine. For instance, you can add a menu item specifying a URL of your Helpdesk solution, so your users can easily reach it when needed.

Create Windows client package for mass distribution

After configuring a Windows Client Theme, you can create a Windows client package for mass distribution as follows:

- 1 While still in the Windows client section of the **Theme Properties** dialog, click the **Generate Windows client package** button.
- 2 In the dialog that opens, specify the following options:
 - Specify the target folder on your local computer where the package will be created.
 - Select or clear the "Open the folder in Windows Explorer " option as needed.
- 3 Click **Generate**. This will create the ClientDownloader.exe file. When you run the file, it will download the latest version of Parallels Client for Windows installer (MSI) and will apply your custom Theme to it.

You can now distribute the installer to end users. When they run the installer, it will install Parallels Client for Windows with all customizations (start menu shortcuts, desktop shortcut, images and icons) specified in the Theme. In the future, if you need to upgrade an installed copy of Parallels Client for Windows to a newer version, you don't need to repeat the instructions described above. Simply upgrade the older version and the branding features will remain intact.

General Theme tasks

When you are done customizing a Theme, click **OK** to save it and return to the Parallels RAS console.

You can also perform the following actions on the **Themes** tab in the Parallels RAS Console:

- **Duplicate a Theme** — right-click a Theme and choose **Duplicate** (or select a Theme and click **Tasks > Duplicate**).
- **Preview HTML5 Theme** — right-click a Theme and choose **Preview HTML5 Theme** (or **Tasks > Preview...**).
- **Delete a Theme** — right-click a Theme and choose **Delete** (or **Tasks > Delete**).

When done creating or modifying Themes, click **Apply** in the Parallels RAS Console to commit the changes to Parallels RAS. You can now test the Theme by opening its URL in an HTML5-enabled web browser.

Delegating session management permissions

If your organization has multiple user groups, all sharing centralized Parallels RAS resources, you have the ability to delegate session management permissions to an administrator of a particular group. When you do, the administrator can see and manage Parallels RAS sessions only for users who belong to that group.

Here's how this functionality works:

- 1 A separate Theme is created for each group. Session management permissions for the Theme are delegated to a custom administrator (see **Managing Administrator accounts** (p. 50)).
- 2 When a custom administrator logs in to the Parallels RAS Console, they are presented with a limited user interface displaying sessions that belong to the Theme (or multiple Themes) that the administrator is allowed to manage.

The rest of this section describes how to configure and use this functionality.

Create a Theme and delegate session management permissions

If you don't have a Theme for a user group, you need to create it. Follow the instructions provided earlier in this chapter (p. 346). To delegate session management permissions, you specifically need to do the following:

- 1 When specifying settings on the **General** page, select the **Limit access to this Theme to members of these Active Directory groups** option and add one or more groups.
- 2 After creating or configuring the Theme, close the **Theme Properties** dialog, then right-click anywhere in the list and choose **Delegate Permissions**.
- 3 If you already have a custom administrator account that you would like to use, it will appear in the list. If you don't have an account, create one as follows:
 - a Click **Tasks > Add**.
 - b In the **Account Properties** dialog, click the [...] button next to **Name** and select an account.
 - c The **Permissions** field is read-only and set to **Custom administrator** (the type that must be used here).
 - d Populate the rest of the fields (email, mobile, etc.) as needed.
 - e Click **OK**.
- 4 Back in the **Delegate Permission** dialog, select the administrator in the left pane.

- 5 In the lower portion of the right pane, select permissions (view, modify, manage sessions) for the desired Theme. You can also set permissions in the upper portion of the right pane, but they will apply to all existing Themes, and this is probably not what we are trying to do here.
- 6 Click **OK**.

Manage sessions

Once the above is complete, the custom administrator can manage sessions that belong to the specified Theme(s). To manage sessions:

- 1 Run the Parallels RAS Console and log in using the credentials of a custom administrator.
- 2 The right pane will contain sessions that belong to the members of the group(s) assigned to the Theme.
- 3 To manage a session, select it, then click the **Tasks** drop-down menu and choose a desired option (Disconnect, Log off, Send message, etc.).

Settings audit

Any changes to administrator permissions are recorded in the settings audit. Possible actions are create, update, and delete. You can view the changes by going to **Administration > Settings audit** or **Farm > Themes > Settings Audit**.

Using Themes in Parallels Client for Windows

In order for a user to use a corresponding Theme, the connection properties must be properly set. To do so:

- 1 In Parallels Client for Windows, right-click a connection and choose **Connection Properties**.
- 2 On the **Connection** tab, the server name must be followed by the Theme name after a forward slash, as in `Server-name/Theme-name`.

When the administrator views sessions in the RAS Console, a client using a Theme can be identified by the Theme name in the **Theme** column.

Open Parallels HTML5 Client

To open Parallels HTML5 Client in a web browser, enter one of the following in a web browser, depending on your setup:

- The DNS name of an HALB device or HALB Virtual Server (if in use). For example, `https://ras.msp.com`.
- The FQDN or IP address of a specific RAS Secure Client Gateway. For example, `https://ras-gw1.company.dom`.

For more information about the HTML5 Client URL, please see **Web request load balancing** (p. 78).

When you open the HTML5 Client in a web browser, the login page is displayed.

Note: By default, when a user opens HTML5 Client in a web browser for the first time, the cookie consent message is displayed at the top of the page in accordance with the GDPR regulation. To read the Parallels cookie policy, the user clicks the provided link. To agree with the policy, the user clicks **Got it** to close the message and continue. The RAS administrator can disable the cookie consent message in the Theme settings dialog (p. 350).

To log in to Parallels RAS, specify your user name in the UPN format (username@domain.com) and password and click **Log in**.

Note: If Parallels RAS is configured to use Google Authenticator as a second-level authentication provider, an additional dialog opens where the user can either scan a QR code or use a secret key to generate a one time password (OTP). For details, please see **Using Google Authenticator** (p. 283).

Once the user is logged in, one of the scenarios described below takes place depending on how the HTML5 Client is configured on the server side. For details, please see **Configure HTML5 Client** (p. 74).

Launch apps in Parallels Client and fallback to HTML5

With this option configured on the server side, you will see a dialog box in the web browser with the following options:

- **Install Parallels Client.** Opens the Parallels Client download and installation page. Follow the instructions and install Parallels Client.

Note: If you don't have administrative permissions on this computer, a dialog will open saying so. The dialog has two buttons: **Install Full Client** and **Install Basic Client**. If you know credentials of an administrative account on this computer, click **Install Full Client** and enter the credentials when asked. The installation will continue using these credentials and the full version of Parallels Client will be installed. If you don't know the credentials, click **Install Basic Client**. The basic version of Parallels Client will still work but some of the functionality will be missing.

After the installation, you should see Parallels HTML5 Client displaying published resources that you can use. Please also note a link in the lower left corner of the screen displaying the Parallels Client version and build number.

You can now run remote applications and desktop in Parallels Client or in a browser (HTML5). The default method for running applications and desktops is Parallels Client. To run a remote application or desktop in a browser, right-click it (or tap and hold on a mobile device) and then choose Parallels HTML5 Client.

- **Open in Parallels HTML5 Client.** Closes this dialog box and opens the main Parallels HTML5 Client screen. Remote applications or desktops will be launched in the web browser. When you open Parallels HTML5 Client the next time, you will again see the same dialog box with the same options.

- **Always open in Parallels HTML5 Client.** This option works similarly to the option above but your selection is remembered the next time you open Parallels HTML5 Client.

Launch apps in Parallels Client

When this option is configured on the server side, you will see a dialog box prompting you to install Parallels Client. Click the link provided to open the Parallels Client download and installation page and follow the instructions. After you install Parallels Client, the main Parallels HTML5 Client screen opens displaying published resources that you can use. If you now double-click or tap a resource, it will be launched in Parallels Client.

Launch apps in browser only (HTML5 only)

With this option configured, the main Parallels HTML5 Client screen opens with no additional prompts. Remote applications and desktops will be launched in the web browser.

Main menu options

To open the Parallels HTML5 Client main menu, click or tap the "person" icon in the upper-right. You can select from the menu options described below.

Settings

Allows you to configure the following settings:

- **Clipboard redirection:** Enable or disable the clipboard in a remote session. Select from the following options: **Bidirectional** (copying is allowed in both directions), **Server to client only**, **Client to server only**, **Disabled** (copying in either directions is not allowed).
- **Sound:** To play the sound on the local computer, select the **Bring to this computer** option. If sound is not supported by your browser, the menu will be disabled and you'll see a corresponding text message below it.
- **Remote audio recording:** Enable or disable the sound input redirection from the local computer to the remote application. For example, if you would like to use a microphone in Skype or a similar app for teleconferencing, you need to enable audio recording in Parallels HTML5 client. Select **Record from this computer** to enable recording or select **Do not record** to disable it.

Note: Audio input is supported in Chrome, Firefox, Edge and Safari 11. If your browser doesn't support audio input, this setting will be disabled and you will see a text message instead.

- **Redirect links:** Select a desired redirection option from the following: **Do no redirect**, **Redirect URLs**, **Redirect email**, **Redirect all**. When redirection is enabled, a link will be opened on the local computer.
- **Windows touch input:** Enable or disable pen input redirection with pressure sensitivity support. Please note that the eraser button is not supported.

Note: Pen input redirection is supported with the following software: Chromium-based browsers running on Windows 8.1 or later, Google Chrome running on Chrome OS. This functionality was tested on Chrome OS 97.X and 98.X.

- **Redirect printers:** Select a printer redirection option: **RAS Universal Printer** (uses the RAS Universal Printing technology) or **Do not redirect** (printers will not be redirected).
- **Keyboard mode:** Select **Universal Keyboard** or **PC Keyboard**. If you have problems typing certain characters, try selecting **PC Keyboard** and then selecting a proper layout in the **Keyboard Layout** drop-down list (see below).
- **Keyboard layout:** Select a keyboard layout (e.g. English (US), English (UK), Japanese). To enable this drop-down list, the Keyboard Mode option must be set to PC Keyboard.
- **Auto login:** Enable or disable auto login in HTML5 Client. If this option is on, and the user credentials have been saved before, the user will not have to enter them again. This option may not be available if a Client Policy was applied where this option is turned off. Note that the auto login option is supported on the latest Chromium-based browsers, such as Google Chrome and Microsoft Edge. For more information, please see **Auto login** (p. 359).
- **Connection timeout, seconds:** Specify the connection timeout.
- **MFA: Remember last method used:** If using multi-factor authentication, enable this option so the last method used is remembered and used by default.
- **Always ask for credentials when starting application:** If this option is enabled, a user will be asked to enter credentials when starting an application even if the session is still active. You can use this option as added security to prevent unauthorized users to access applications. For example, if a user disconnects from a session, no one else will be able to take over the session and run remote applications. As another example, if a user leaves a device with an open HTML5 Client displaying the app listing (with or without running RDP sessions) then any user who tries to open a new application or another instance of a running application will be prompted for credentials. Please note that the **Auto login** option (described above) must be disabled for this functionality to work; otherwise saved credentials will be used automatically.

Change Password

Allows the user to remotely change their domain password. When the password is being changed, the password requirements are displayed on the screen, so the user can follow them for the new password to be accepted. This option can be disabled through Client Policies (**Control settings > Password > Prohibit changing password**).

Detect Client

Determines if Parallels Client is installed on the local computer. If Parallels Client is not installed, gives user an option to install it or skip the automatic Parallels Client detection on subsequent logons.

Download Client

Opens a web page with instruction on how to download and install Parallels Client.

Logout

Ends user session with Parallels RAS and logs the user out.

Running remote applications and desktops

To launch a remote application or desktop in Parallels HTML5 Client, do one of the following:

- Double-click (or tap on a mobile device) an application or a desktop icon. The resource will open inside a web browser or in Parallels Client depending on the server-side HTML5 configuration (RAS Secure Client Gateway Properties > HTML5 > Launch sessions using option).
- Right-click (or tap and hold on a mobile device) an application or a desktop to display a context menu. The menu will appear if the **Allow user to select launch method** or **Allow opening applications in a new tab** (or both) options are selected on the **RAS Secure Client Gateway Properties > HTML5** tab in the RAS console. The menu allows you to choose whether to open the resource in Parallels Client or Parallels HTML5 Client (depending on the setting mentioned above) and it also allows you to choose whether to open an application in the same or new tab in the web browser.
- If a resource cannot be opened in Parallels Client due to an error, a message will be displayed with an option to open it in the web browser instead.

Please note that to open a resource in Parallels Client from the HTML5 page, a URL with a custom scheme is used. When you double-click on a resource on the HTML5 page, the URL is executed and is then passed to Parallels Client which uses the instructions that it contains to open the resource. For more information see **RAS HTML5 Gateway API and Parallels Client URL scheme** (p. 467).

Using drag and drop functionality

Parallels HTML5 Client supports drag and drop functionality when running remote applications and desktops.

Note: The **Allow file transfer command** option must be enabled on the Gateway for the drag and drop functionality to work. See **Configure HTML5 Client** (p. 74).

Here's how to use drag and drop when working with a remote application:

- 1 Select a file on your local computer.
- 2 Drag and drop the selected file to an app. The 'Save as' window will pop up.
- 3 Enter a name for the file and save it. The file will be saved on the server hosting the app.

You can also drag and drop files between two remote apps running on different hosts.

Here's how to use drag and drop with a remote desktop:

- 1 Select a file on your local computer.
- 2 Drag and drop the selected file to a remote desktop. The 'save as' window will pop up.
- 3 Enter a name for the file and save it. The file will be saved on the desktop on the server that hosts it.

Native clipboard experience

In Parallels RAS 18.2 and newer, you can copy and paste plain text between a local device and remote session in both directions. Simply use Ctrl+C/Cmd+C to copy (Ctrl+X/Cmd+X to cut) and Ctrl+V/Cmd+V to paste. This functionality is fully supported in Chromium based browsers (Chrome, Edge Chromium, Opera) and Internet Explorer. On Firefox, only server to client copy/paste is supported. Other web browsers do not support this functionality.

Other useful features

Other useful functionality on the main Parallels HTML5 Client screen includes the following:

- **Favorites list.** You can add a remote application or a desktop to the Favorites list, so you can easily find them. To do so, point to or tap an application or a desktop and then click or tap the "star" icon. To view the list, click or tap the "star" icon on the footer toolbar (in the lower left). To remove a resource from the list, point to it and click the "X" icon (or point to or tap the resource icon and then click or tap the star icon).
- **Search.** To search for a resource, begin typing its name in the **Search** box (upper right). The list will be filtered as you type to contain only the resources with matching names.
- **View a description.** To view a resource description, position the mouse pointer over it. The description will appear as a tooltip. This could be helpful if one or more resources are published using the same name. By reading the description, you can distinguish between them.
- **Taskbar.** When you launch a remote application or a desktop, its icon is added to the taskbar at the bottom of the screen. When the taskbar is full, items of the same type are grouped to save space. You can click or tap on a group to see the list of all running instances and to switch to or close a particular instance.

Auto login

Auto login for RAS HTML5 Client portal facilitates the frequent use of the portal by providing auto login option for the user without requiring user intervention to input their credentials. With Auto login enabled, as soon as user opens the Parallels HTML5 Client portal, user will be automatically logged in, able to see the list of resources that were made available by the administrator and to launch resources accordingly. This experience can be configured from the HTML5 Client settings or centrally controlled from RAS policies (p. 389). It is set to accelerate user login and increase user experience by reducing the number of times users are prompted to login to the Parallels HTML5 Client Portal.

To use this feature, the following requirements must be met:

- RAS HTML5 Gateway must have a valid and trusted (by end point devices) certificate.
- There is only one set of valid credentials already saved for a given domain (including subdomains). Auto login is not applicable in case of using shared devices due to multiple user accounts accessing the same Parallels RAS environment.
- Not using incognito — when using incognito, and only one set of credentials is available, sign-in pop up still shows, so auto login cannot happen without the user intervention.
- Auto login is supported on Chromium-based browsers only, such as Google Chrome and Microsoft Edge.

Configuration

The following settings control the Auto login functionality:

- When a user logs in to HTML5 Client for the first time, they need to click **Save** when asked to save the password. For this to happen, the **Offer to save password** and **Auto sign-in** options must be enabled in the browser (these are the default settings in a Chromium-based browser).
- The user needs to confirm to reduce the number of times they asked for credentials. This will enable the Auto login option.
- The RAS administrator may also use RAS policies to enforce auto login (enabled/disabled). This can be done from **Policy > Session > Connection > Primary Connection > Auto Login**.
- An expiry time of 60 days is set when the **Auto login** option is enabled due to security.

Using Auto login

The following describes how Auto login works:

- 1** The user opens the HTML5 Client web page in a browser and logs in. Note that direct app access is also supported (p. 360).
- 2** On the first login, HTML5 Client will suggest the user to enable Auto login.

- When the user opens the HTML5 Client (or uses a direct app link) the prompt to enter credentials will not appear.

To see the Auto login setting in the HTML5 Client, click on the user icon in the top right and click **Settings**. Examine the **Auto Login** setting.

Direct App access

Specific published resources in Parallels RAS 18 may be directly accessed through RAS HTML5 Client. This can be achieved with the introduction of a new parameter, *appid*, which allows administrator to use links to access the published resource directly. This allows a more flexible and easier way for users to access Parallels RAS published resources such as using browser shortcuts or bookmarks or third-party portals such as Azure My Apps Portal to access independent SAAS applications and Parallels RAS virtual apps and desktops.

To launch a published resource directly, you need to specify a URL using one of the following formats:

URL format	Description
https://FQDN?appid=	This format omits the Theme name and uses the default HTML5 Client Theme. The "appid" parameter specifies the published resource ID as seen in the Publishing category in the RAS Console. The ID is automatically generated when a resource is published. To see it, select a published resource and examine the Application field on the Information tab. For example, #5: Microsoft Office Word — the application ID of the Microsoft World here is 5.
https://FQDN/<Theme-name>?appid=<app-ID>	This format is similar to the one above, but specifies a Theme name.
https://FQDN/RASHTML5Gateway?theme=&appid=	This format is the same as the one above, but uses the full URL specification. It is listed here just for reference.

Supported parameters:

Parameter	Description
appid	The published item (application or desktop) to be launched.
overrideparams	[Optional]. URL Encoded Override arguments that needs to be passed to the published application.

Example:

https://FQDN?appid=14&overrideparams=C%3A%2Ftest.txt

When opening a published resource using a direct link, the **Auto login** option (p. 359) is also used depending on the settings.

Using the toolbar

Parallels HTML5 Client includes a special toolbar that becomes available when you launch a remote application or desktop. The toolbar appears differently for remote desktops and remote applications. The toolbar has also slightly different functions for desktop computers and mobile devices. The differences are explained in the subsequent topics.

In this section:

- Using the toolbar on desktop computers (p. 361)
- Using the toolbar on mobile devices (p. 363)
- Using the remote clipboard (p. 364)
- Hiding toolbar items (p. 365)

Using the toolbar on desktop computers

Remote desktop toolbar

When you launch a remote desktop in a web browser on a desktop or laptop computer, the toolbar appears as follows:



The top area of the toolbar is used to drag the toolbar up or down. Click and hold it and then drag the toolbar to the desired position. The arrow icon is used to show or hide the toolbar items.

The main toolbar items are (from top to bottom):

- **Full screen.** Display the remote desktop in full screen on the local computer.
- **Upload a file.** Upload a file from the local computer to the remote server. After clicking this item, you are presented with two dialogs, one after another. In the first dialog, select a file on the local computer you wish to upload. In the second dialog, select a location on the remote server where you want to save the file.
- **Download a file.** Download a file from the remote server to the local computer. After clicking this item, select a file on the remote server you wish to download. Depending on your web browser configuration, the download will start automatically or you will be asked to select a destination folder on your local computer.
- **Shortcuts.** Display the **Shortcuts** menu (see below for the menu description).
- **Clipboard.** Display the remote clipboard. Please see **Using the remote clipboard** (p. 364) for more information.

The **Shortcuts** menu allows you to send keystrokes and key sequences to the remote desktop:

- **Escape.** Sends the "Escape" keystroke to the remote desktop.
- **Tab.** Sends the "Tab" keystroke.
- **Backspace.** Sends the "Backspace" keystroke.
- **Print screen.** Sends the "Print Screen" keystroke. The screen will be printed to the clipboard of the remote desktop from where you can paste it into an application (e.g. Paint) running on the same remote computer.
- **Windows Key.** Sends the "Windows logo key" keystroke.
- **Control+Alt+Delete.** Sends the "Ctrl+Alt+Delete" key sequence.

Remote application toolbar

When you launch a remote application, the toolbar is embedded into the page footer and it's collapsed by default. To expand the toolbar, click the "arrow-up" icon in the lower right-hand corner.




The toolbar items are (from top to bottom):

- **Download.** Download a file from the remote server to the local computer. After clicking this item, select a file on the remote server you wish to download. Depending on your web browser configuration, the download will start automatically or you will be asked to select a destination folder on your local computer.
- **Upload.** Upload a file from the local computer to the remote server. After clicking this item, you are presented with two dialogs, one after another. In the first dialog, select a file on the local computer you wish to upload. In the second dialog, select a location on the remote server where you want to save the file.
- **Clipboard.** Display the remote clipboard. Please see **Using the remote clipboard** (p. 364) for more information.

Using the Toolbar on Mobile Devices

Remote Desktop Toolbar

When you launch a remote desktop in a web browser on a mobile device, the toolbar appears as follows:



The small arrow icon at the top is used to show or hide the toolbar items.

The main toolbar items are (from top to bottom):

- **Upload a file.** Upload a file from the local device to the remote server. Note that in iOS, you can upload from the Photos folder only.
- **Download a file.** Download a file from the remote server to the local device (not available in iOS).
- **Shortcuts.** Display the **Shortcuts** menu (see below for the menu description).
- **Clipboard.** Display the remote clipboard. Please see **Using the remote clipboard** (p. 364) for more information.

- **Keyboard.** Display the native keyboard. This opens your mobile device native keyboard so you can type in an application on the remote desktop.
- **Arrow.** The arrow icon is used to switch between the two available mouse input modes:
 - Mode 1:** The first mode (the arrow icon is white) follows the movement of your finger on the screen and performs a click on a remote desktop where you tap.
 - Mode 2:** The second mode (the arrow icon is red) displays a virtual mouse pointer on the remote desktop and allows you to move that pointer to a precise position with your finger. When you tap anywhere on the screen, the click on the remote desktop is performed at the precise position of the virtual mouse pointer.

The **Shortcuts** menu allows you to send keystrokes and key sequences to the remote desktop:

- **Escape.** Sends the "Escape" keystroke to the remote desktop.
- **Tab.** Sends the "Tab" keystroke.
- **Backspace.** Sends the "Backspace" keystroke.
- **Print screen.** Sends the "Print Screen" keystroke. The screen will be printed to the clipboard of the remote desktop from where you can paste it into an application (e.g. Paint) running on the same remote computer.
- **Windows Key.** Sends the "Windows logo key" keystroke.
- **Control+Alt+Delete.** Sends the "Ctrl+Alt+Delete" key sequence.

Remote application toolbar

When you launch a remote application, the toolbar is embedded into the page footer and it's collapsed by default. To expand the toolbar, click the "arrow-up" icon in the lower right-hand corner.

The toolbar items are (from top to bottom):

- **Download.** Download a file from the remote server to the local device (not available in iOS).
- **Upload.** Upload a file from the local device to the remote server. Note that in iOS, you can upload from the Photos folder only.
- **Clipboard.** Display the remote clipboard. Please see **Using the remote clipboard** (p. 364) for more information.
- **Keyboard.** Display the native keyboard. This opens your mobile device native keyboard so you can type in an application on the remote desktop.

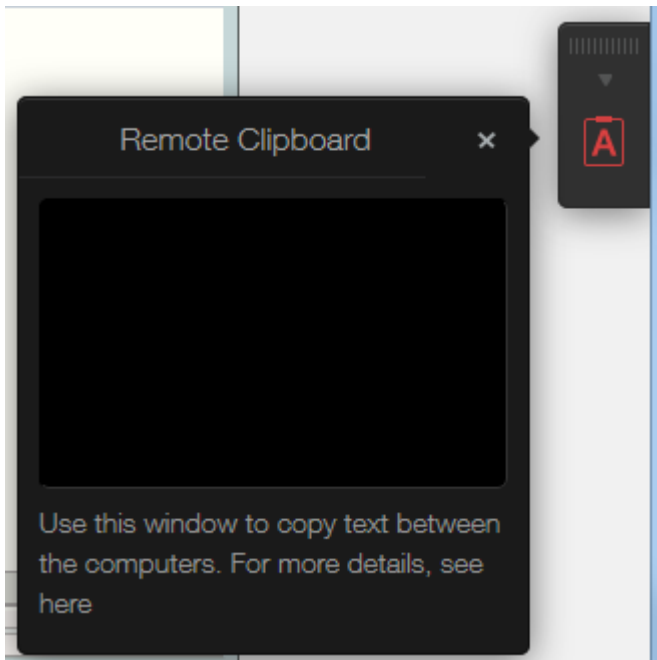
Using the remote clipboard

The Remote Clipboard allows you to copy and paste text between your local client system and a remote application or desktop. The clipboard is accessed from the toolbar.

Note: In Parallels RAS 18.2 and newer, you can copy and paste plain text between a local device and remote session without using the Remote Clipboard. For more information, see **Native clipboard experience** (p. 358).

To use the clipboard:

- 1 Expand the toolbar click the **[A]** icon.
- 2 This opens the **Remote Clipboard** window. On the screenshot below, a remote desktop toolbar is shown. A remote application toolbar looks differently, but it functions exactly the same.



- 3 To copy text from the local computer to a remote application, type (or paste) it in the **Remote Clipboard** window. The text is automatically saved on the remote computer clipboard, so you can use a standard paste command (e.g. Ctrl+V) to paste it into a remote application.
- 4 To copy text from a remote application to the **Remote Clipboard** window, highlight it and use the standard copy command (e.g. Ctrl+C). The text will appear in the **Remote Clipboard** window from where you can copy it to a local application.

Hiding toolbar items

You can hide the clipboard and file transfer items on the toolbar if you believe that it's a security risk. The clipboard can be disabled on a RAS Secure Client Gateway or Client Policy level.

To disable the clipboard for a Gateway:

- 1 In the Parallels RAS Console, navigate to **Farms** > <Site> > **Gateways**.
- 2 Right-click a desired RAS Secure Client Gateway and choose **Properties**.

- 3 Select the **HTML5** tab and clear the **Allow clipboard command** option in the **Restrictions** section.

You can also disable the clipboard on the Client Policy level, which will disable it for a given user or user group on any Gateway they connect to:

- 1 In the Parallels RAS Console, select the **Policies** category.
- 2 Right-click a policy and choose **Properties**.
- 3 Select the **Connection Properties** item in the left pane and then select the **Local Resources** tab in the right pane.
- 4 In the **Local devices and resources** section, clear the **Clipboard** option.

Note: Please note that when enabling or disabling the clipboard on a client policy level, this will also affect the clipboard functionality on desktop and mobile versions of Parallels Client. This means that if you disable the clipboard, the desktop and mobile device users will not be able to use their local clipboard when working with a remote application.

You can also disable the file upload and file download items on the toolbar. For instructions, please read the **Configuring remote file transfer** section (p. 408).

Universal Printing

Printer redirection enables users to redirect a print job from a remote application or desktop to their local printer, which can be connected to the user's computer or be a local network printer attached via an IP address. RAS Universal Printing simplifies the printing process and solves most printer driver issues by eliminating the need for a remote server to have a printer driver for a specific local printer on the client side. Therefore, a user can print regardless of which printer they have installed locally, and the RAS administrator doesn't have to install a printer driver for each printer connected to the local network.

In This Chapter

Managing Universal Printing settings	367
Universal Printing drivers	368
Font management	369

Managing Universal Printing settings

To configure RAS Universal Printing, select the **Universal Printing** category in the RAS Console.

By default, the Universal Printing driver is automatically installed together with an RD Session Host Agent, VDI Guest VM Agent, or a Remote PC Agent. Therefore, upon adding a server to the Farm, the Universal Printing is already enabled. The Universal Printing driver is available as a 32 bit and 64 bit version.

Enabling and disabling Universal Printing support

To enable or disable the Universal Printing support for a server, right-click the server in the **Servers in Site** list and click **Enable** or **Disable** in the context menu.

Configuring a printer renaming pattern

By default, Parallels RAS renames printers using the following pattern: `%PRINTERNAME% for %USERNAME% by Parallels`. For example, let's say a user named Alice has a local printer named Printer1. When Alice launches a remote application or desktop, her printer is named `Printer1 for Alice by Parallels`.

To change the default printer renaming pattern, select the Universal printing category. On the Universal printing tab, specify a pattern in the **Printer rename pattern** field. To see the predefined variables that you can use, click the [...] button next to the input field. The variables are:

- `%CLIENTNAME%` – the name of the client computer.
- `%PRINTERNAME%` – the name of a printer on the client side.
- `%SESSIONID%` – RAS session ID.
- `%USERNAME%` – the name of the user connected to RAS.
- `<2X Universal Printer>` – This is a legacy mode where only one printer object will be created in the RDP session.

You can also use certain other characters in a printer renaming pattern. For example, you can define the following commonly used pattern: `Client/%CLIENTNAME%#/%PRINTERNAME%`. Using this pattern (and the user named Alice from the example above), a local printer will be named `Client/Alice's Computer#/Printer1`

You can specify a different printer renaming pattern for each server in the **Servers in Site** list.

Note: Redirected printers are only accessible by the administrator and the user who redirected the printer.

Printer retention

When client-defined printers are redirected to a remote session, it takes time and impacts overall session establishing time. To improve user experience, you can reuse previously created user's printers. To do so, on the **Universal printing** tab, set the **Printer retention** option to **On**.

Universal Printing drivers

A system administrator can control the list of client-side printer drivers which should be allowed or denied the Universal Printing redirection privileges.

Using this functionality you can:

- Avoid server resource overloading by non-useful printer redirection. Since the majority of users choose to redirect all local printers (this is default setting), a large number of redirected devices is created on the server which are not really used. It's mostly related to various paperless printers like PDFCreator, Microsoft XPS Writer, or various FAX devices.
- Avoid server instability with certain printers. There are some printers that might create server instability (spooler service component) and as the result deny printing services as a whole for all connected users. It is very important that the administrator has the ability to include such drivers to the "deny" list to continue running printing services.

To specify printer drivers:

- 1 In the Parallels RAS Console, navigate to **Universal Printing > Printer drivers**.
- 2 In the **Mode** drop-down list, select which printers should be allowed redirection from the following options:
 - **Allow redirection of printers using any driver** — (default) This option places no limitation on the type of driver a printer is using to use redirection privileges.
 - **Allow redirection of printers using one of the following drivers** — Only the printers using the drivers listed in the box below the **Mode** field are allowed redirection. To add a printer driver to the list, click the **Tasks > Add** (or click the + icon) and type the printer driver name in the edit field provided.
 - **Don't allow redirection of printers that use one of the following drivers** — This is probably the most useful option in the context of this feature. The printers that use the drivers specified in the list will be denied redirection privileges. All other printers will be allowed to use redirection. To add a printer driver to the list, click the **Tasks > Add** (or click the + icon) and type the printer driver name in the edit field provided.
- 3 To delete a printer driver from the list, click **Tasks > Delete** or click the minus-sign icon.
- 4 When done making changes, click the **Apply** button to save the changes.

Please make a note of the following:

- When adding a printer driver to the list, type the printer *driver* name, not the printer name.
- The driver names comparison is case insensitive and requires full match (no partial names, no wildcards).
- The settings that you specify on this tab affect the entire Site (not an individual server).

Font management

Fonts need to be embedded so that when printing a document using Universal Printing the document is copied to the local spooler of the client machine to be printed. If the fonts are not present on the client machine the print out would not be correct.

To control the embedding of fonts within a print job use the **Fonts Management** tab page and check/uncheck the option **Embed Fonts**.

Excluding fonts from embedding

To exclude a specific font type from being embedded, click **Tasks > Add** in the **Exclude the following Fonts from embedding** section and select a font from the list.

Automatically Install fonts on servers and clients

To automatically install a specific font type on servers and clients, click **Tasks > Add** in the **Auto install fonts** section and select the fonts from the list.

Note: By default, fonts added to the auto install list will be excluded from the embedding list because the fonts would be installed on the Windows clients, therefore there is no need for them to be embedded. Clear the option **Automatically exclude font from embedding** in the select font dialog so the font is not excluded from the embedding list.

Resetting excluded fonts to default

To reset the list of excluded fonts to default, click **Tasks > Reset to default**.

You can also specify a universal printing compression policy. For more info see **Client Policies > Experience** (p. 398).

Universal Scanning

Scanner redirection enables users who are connected to a remote desktop or accessing a published application to make a scan using the scanner that is connected to the client machine. This chapter describes how to configure and use RAS Universal Scanning services.

In This Chapter

Managing Universal Scanning.....	371
Adding scanning applications	372

Managing Universal Scanning

Universal Scanning uses TWAIN and WIA redirection to let any application using either technology hardware connected to the client device for scanning. With Universal Scanning there is no need to install a specific scanner driver on the server.

Note: The server feature **Desktop Experience** is required in order to enable both WIA and TWAIN scanning on RD Session Hosts.

To configure Universal Scanning, select the **Universal Scanning** category in the RAS Console.

By default, the Universal Scanning driver is automatically installed with RD Session Host, Guest VM, and Remote PC agents. Therefore, upon adding a server to the Farm the Universal Scanning is installed.

Configuring a Scanning Rename Pattern

By default, Parallels RAS renames scanners using the following pattern: %SCANNERNAME% for %USERNAME% by RAS. For example, if a user named Lois, who has SCANNER1 installed locally, connects to a remote desktop or published application, her scanner is renamed to "SCANNER1 for Lois by RAS".

To change the pattern used to rename scanners, specify a new pattern in the **Scanner rename pattern** input field. The variables that you can use for renaming are:

- %SCANNERNAME% — client side scanner name.
- %USERNAME% — username of the user connected to the server.

- %SESSIONID% — ID of the active session.

You can configure a different renaming pattern specifically for each server in the list.

Note: Redirected scanners are only accessible by administrator and the user who redirected the scanner.

Enabling and Disabling Universal Scanning Support

To enable or disable the WIA or Twain Universal Scanning support for a particular server, click the **WIA** tab or the **TWAIN** tab, then right-click a server and click **Enable** or **Disable** in the context menu.

Adding scanning applications

Adding a scanning application

TWAIN applications that will use the Universal Scanning feature have to be added in the TWAIN tab by selecting the **TWAIN Applications** button so they can use the Twain driver, hence making it easier for the administrator to set them up.

To add an application to the list of scanning applications:

- 1 With the **Universal Scanning** category selected in the RAS Console, click the **TWAIN** tab.
- 2 Click the **Twain Applications** button (below the **Servers in Site** list) and then click **Add**.
- 3 In the **TWAIN Applications** dialog, click **Tasks > Add** and browse for the application executable. Select the executable and click **Open**.

Note: Some applications might use different or multiple executables. Make sure that all required executables are added to the list of scanning applications.

Deleting a scanning application

To delete a scanning application from the list, highlight it and click **Tasks > Delete**.

Note: If you delete an application from the list, the installation of the application will not be affected.

You can also specify a universal scanning compression policy. For more info see **Client Policies > Experience** (p. 398).

User Device Management and Client Policies

This chapter describes tasks that a Parallels RAS administrator can perform to manage user devices, such as desktop computers, phones, or tablets.

In This Chapter

Inviting users to connect to Parallels RAS	373
Mass configuring user devices.....	373
Enabling Help Desk support	375
Monitoring devices	375
Windows device groups	377
Managing Windows devices	379
Scheduling Windows devices & groups power cycles.....	385
Client Policies	386
Configuring remote file transfer	408

Inviting users to connect to Parallels RAS

Parallels RAS supports multiple platforms ranging from desktop PCs and Mac computers to mobile devices and ChromeApps. The Invitation Email feature is designed to reduce the complexities involved in the installation and client rollout process. This feature allows the administrator to send client installation and automatic configuration instructions to end users right from the Parallels RAS Console.

Before proceeding, please confirm that you've configured the mailbox as described in **Configuring SMTP server connection for email notifications** (p. 446). To send an invitation email to users, use the **Start** category in the RAS Console. For more information see **Invite users** (p. 37).

Mass configuring user devices

If you need to configure Parallels Client that is already installed on multiple devices in your organization, you can simplify the procedure by using one of the following mass configuration options:

- By exporting Parallels Client settings to a file and then importing them into all other Parallels Client installations.
- Using the Parallels Client URL scheme.

Exporting and importing Parallels Client settings

Parallels Client includes the Export/Import functionality that lets you export RAS or RDP connection settings to a file and then import them into Parallels Client running on another device. This functionality is available on all platforms, including desktop and mobile versions of Parallels Client (except Parallels Client for Chrome App). The Export/Import functionality is accessed in Parallels Client as follows:

- **Windows, Mac, Linux:** On the main menu, click **File > Export Settings** or **File > Import Settings**.
- **iOS/iPadOS:** To export connection settings, tap the [...] icon in the top right corner and choose **Share Connection**. To import, select the file that you exported earlier and choose to open it with Parallels Client.
- **Android:** To export connection settings, tap the menu icon (three vertical dots) in the top right corner and choose **Share connections**. To import, select the file that you exported earlier and choose to open it with Parallels Client.

For more information about exporting and importing connection settings, see the Parallels Client Guide for a desired platform.

Using Parallels Client URL scheme

Parallels RAS uses a URL scheme to perform actions in Parallels Client installed on user devices. Specifically, the URL scheme can be used to configure RAS and RDP connections using predefined settings. For the information about the URL scheme please see **RAS HTML5 Gateway API and Parallels Client URL scheme** (p. 467).

The URL scheme is used in invitation emails when you send an email to your users to install Parallels Client on their devices. An invitation email includes a link, which is a complete URL that uses the Parallels Client URL scheme. When you mass install Parallels Client on user devices, you simply send an invitation email to your users (p. 37). If you need to reconfigure existing Parallels Client installations (and don't want to do it by sending an invitation email), you can do the following:

- 1 Create an invitation email containing configuration profiles for all required platforms and send it to yourself.

- 2 Open the email and copy Parallels Client configuration URLs to a local intranet portal.
- 3 Let your users know where the URLs are.
- 4 To configure Parallels Client, your users will need to simply click a URL for their platform. This will automatically configure Parallels Client on their devices.

Enabling Help Desk support

Parallels Client provides users with the ability to send a support request, together with a problem report, to your organization help desk.

Note: At the time of this writing, this functionality is only available in Parallels Client for iOS and Parallels Client for Android. Support for other clients will be added in future releases.

To enable Help Desk support, do the following:

- 1 In the RAS Console, select the **Features** category.
- 2 Select the **Enable Helpdesk functionality in Parallels Client** option and specify your help desk email address in the field provided. This email address will be updated in Parallels Client every time a user connects to Parallels RAS from it.

Help desk can be accessed in Parallels Client from the Help section (or menu). When the user selects the **Request support from helpdesk** item, a local email client will open. The following information will be prefilled in the email:

- Help desk email address (the one you set in the RAS Console).
- Application name.
- A screenshot.
- User name.
- Application version.
- Operating system version.

The user can provide their own description of the request.

Monitoring devices

Device monitoring allows you to view devices which are connected to the Farm or have established a connection at least once in the past. To monitor devices, select the **Device Manager** category in the Parallels RAS Console and click the **Devices manager** tab in the right pane. The information for a device includes:

- Device name
- IP address
- State (see below for the list of states)
- Last user (who used a device)
- MAC address
- OS version
- Parallels Client version
- Group (if a device is a member of a device group)
- Gateway name (the RAS Secure Client Gateway a device is connected to)
- Gateway IP address

Device states

Devices that connect to Parallels RAS can have any of the following states:

- **Off:** Device is switched off.
- **Connected:** Device is connected.
- **Logged On:** Devices is logged on to the system.
- **Standalone:** Device has previously connected to the Parallels RAS but is not using Parallels Client, therefore it cannot be managed.
- **Not Support:** Device is not supported by the Parallels RAS.
- **Foreign Managed:** Connecting to the Farm but managed by a different Farm.
- **Not Manageable:** Client not manageable due to incompatible client version or uninstalled component.
- **Locked.** Device has an active session in locked status.
- **Pair Pending.** Connection should be refreshed on the client side; port UDP 20009 is blocked from the client to gateway; client management port is disabled on the gateway.

Switching off device monitoring

If you use third-party endpoint management solutions and do not need Parallels RAS device monitoring, you can switch it off on the **Device Manager > Options** tab. It helps you to save computing resources and may improve performance of the RAS Console.

To switch off device monitoring:

- 1 In the RAS Console, navigate to **Device Manager > Options**.
- 2 Clear the **Enable device manager** option.
- 3 Click **Yes** and **Apply**.

As soon as you switch off device monitoring, the RAS Console stops tracking connected devices and deletes device connection history, which is displayed on the **Device Manager > Device Manager** tab. After that you can still view information about current connections in the **Sessions** category.

Getting additional device information

To see the additional device information, right-click a device and choose **Get Device Information** in the context menu. In the dialog that opens, review the following properties:

- **Name:** Device name.
- **IPs:** Device IP address (or multiple addresses if applicable).
- **MAC Address:** MAC address.
- **State:** State (see below for the list of states).
- **Last User:** The user who logged in from this device the last time.
- **Last Logon Time:** The time of last logon.
- **OS Version:** The operating system version running on the device. Windows portable and U3 clients are marked as "Portable".
- **Client Version:** Parallels Client version installed on the device.
- **Gateway IP:** The RAS Secure Client Gateway IP address (the gateway the client is using).
- **Gateway:** The RAS Secure Client Gateway name.
- **Last Activity:** The date and time when any activity was detected from this device.

Windows device groups

The **Windows device groups** tab (**Device Manager** category) allows you to group managed Windows devices and administer them together.

Creating a Windows device group

To create a Windows Device Group:

- 1 Navigate to the **Windows device groups** tab in the **Device Manager** category and click **Tasks > Add**.
- 2 On the **Main** tab page, specify a group name and an optional description.
- 3 On the **OS Settings** tab, set the following options:
 - **Disable removable drives.** Disable mounting of removable drives on managed Windows device.
 - **Disable Print Screen.** Disable the **Print Screen** key.
 - **Replace desktop.** This feature makes a Windows computer behave like a thin client. It limits users from changing system settings or installing new applications. The administrator can add local apps (which are already installed on a computer) to the app list in addition to published resources from Parallels RAS. If you select this option, specify an administrator password in the **Admin Mode Password** field (below) to be used to switch a computer between user and admin modes.
 - **Kiosk mode.** Enable the kiosk mode. This will disable power cycling functions (reboot, shutdown) on computers in the group. Note that power functions will still be available when the computer is switched to the Admin mode.
 - **Use client as desktop.** If this option is selected, Parallels Client will run in full screen mode. A user will not be able to minimize it. Select this option to overcome an issue with Parallels Client breaking out of the kiosk mode on Windows 8.x. The issue may manifest itself in the tile-based UI or while using the "drag to close" feature.
 - **Admin Mode Password.** Specify a password to switch between user and admin modes when a Windows desktop is replaced (see **Replace desktop** above).
- 4 On the **Firewall Settings** tab, enable or disable the firewall and add the inbound ports if necessary.
- 5 On the **Shadowing** tab, select the **Request Authorization** option to prompt a Windows device user before remotely controlling their desktop. If enabled, the user can choose to decline the connection. For more information, see **Managing Windows devices** (p. 379).

Adding a Windows device to a group

To add a Windows device to a group:

- 1 Navigate to the **Device Manager > Device Manager** tab.
- 2 Select one or more devices, then click **Tasks** (or right-click) and choose **Move to Group**.
- 3 Select a group and click **OK** to save the settings.

The administrator can now perform standard Windows power operations (Power On, Power Off, Reboot, Logoff, Lock) on groups of devices.

Managing Windows devices

The Device Manager feature allows the administrator to convert Windows devices running Windows 7 up to Windows 11 into a thin-client-like OS. In order to be managed, Windows devices must be running the latest version of Parallels Client for Windows.

Read the instructions below to learn how to set up Parallels Client on a Windows computer and how to enroll and manage it in Parallels RAS.

Install Parallels Client on a Windows computer

To install and configure Parallels Client for Windows, follow the steps below. You can also read the **Parallels Client for Windows User's Guide** for the complete instructions on how to install and configure Parallels Client.

- 1 Download the Parallels Client for Windows from <https://www.parallels.com/products/ras/download/client/>.
- 2 Double click the `RASClient.msi` or `RASClient-x64.msi` and follow the on-screen instructions to complete the installation wizard.
- 3 Create a new Parallels RAS connection by clicking **File > Add New Connection**.
- 4 Select **Parallels Remote Application Server** and click **OK**.
- 5 Next, configure the following connection properties:
 - **Primary Connection** — Specify the Parallels RAS FQDN or IP address.
 - **User Credentials** — Enter username, password, and domain.
- 6 Click **OK** to create the connection and then double-click it to connect to Parallels RAS.

Upon completion, the Windows device will appear in the Parallels RAS Console in **Device Manager > Devices**.

Windows device enrollment

You can configure Parallels RAS to enroll a Windows device automatically or you can opt to do it manually.

To manually enroll a Windows device in Parallels RAS:

- 1 In the RAS Console, navigate to **Device Manager > Devices**.
- 2 Select a device on the **Devices** tab.
- 3 Click **Tasks > Manage Device**.

The device state will change to **Pair pending** until the device reconnects. Ensure the **Device Manager Port** option is enabled for a gateway. To verify this:

- 1 Navigate to **Farm** > <Site> > **Gateways**.
- 2 Select a gateway and click **Tasks** > **Properties**.
- 3 Click the **Network** tab and make sure that the **Device Manager Port** option is selected

Once the device reconnects, the enrollment process is complete and the device state is updated to **Logged On**, which indicates that it's now managed by Parallels RAS. The user running Parallels Client on their Windows PC can also verify that the PC is managed by clicking **Help** > **About** on the main Parallels Client menu. The information includes the RAS Secure Client Gateway information that the Parallels Client uses to communicate with Parallels RAS.

You can also set Parallels RAS to automatically manage Windows devices. To do so:

- 1 In the RAS Console, select the **Device Manager** category.
- 2 Click the **Options** tab.
- 3 Enable the **Automatically manage Windows devices** option.

The administrator can now check the state of the device and perform power operations, such as Power On, Power Off, Reboot, and Logoff.

Note: Devices running some older versions of Parallels Client cannot be managed and are marked as **Not Supported**.

Lock a Windows device

To lock a Windows device that has an active session, select it in the list and then click the **Lock** item in the toolbar at the bottom. Note that the **Lock** icon is only enabled when the selected device is in the **Logged On** state.

You can also lock a device (or a device group) using the scheduler, which is described in the **Scheduling Windows devices & group power cycles** section (p. 385).

Shadow a Windows device

By shadowing a Windows device, you gain full access to the Windows desktop on the device and can control local and remote applications.

To shadow a Windows device:

- 1 In the RAS Console, navigate to **Device Manager** > **Devices**.
- 2 Select a device and click the **Shadow** item in the toolbar at the bottom.

The Windows user will be prompted to allow the administrator to take control over the device and can choose to deny access. The **Request Authorization** prompt can be deactivated by the administrator. To do so:

- 1 In the Parallels RAS Console, select the **Device Manager** category and click the **Windows Device Groups** tab in the right pane.
- 2 Right-click a group and choose **Properties**.
- 3 In the **Windows Device Group** dialog, select the **Shadowing** tab and clear the **Request Authorization** option.

Desktop replacement

The **Replace desktop** feature limits users from changing system settings or installing new applications. When this feature is enabled, the Windows desktop is replaced by Parallels Client, which converts it into a thin-client-like OS without actually replacing the operating system. This way the user can only deploy applications from Parallels Client, which gives the administrator a higher level of control over connected devices.

Additionally, the Kiosk mode allows you to limit the user from power cycling a device (power actions are still available in the Admin mode; see below for details.).

To enable the **Replace desktop** feature:

- 1 In the **Device Manager** category, select the **Windows Device Groups** tab.
- 2 Right-click a group and choose **Properties**.
- 3 Click the **OS Settings** tab.
- 4 Enable the **Replace desktop** option and optionally the **Kiosk mode** option.
- 5 Click **OK**.

Note: This feature requires an administrative password set to switch between User and Admin mode on the Windows device.

Switching to the Admin mode

In User mode, the user is restricted to use only the applications provided by the administrator. In order to change system settings, switch the device to the Admin mode.

To switch to the Admin mode, right-click on the system tray icon and select **Switch to admin mode**. Type the password when prompted.

The following table outlines features that are available in Admin and User modes.

Feature	User Mode	Admin Mode
Parallels Client Global Options		x
Parallels Client Farm Connection Properties		x
Configuration of Local Applications		x
Add a new RAS Connection		x
Add a new RDP Connection		x
Manage Standard RDP Connections and Folders		x
Display Settings	x	x
Mouse Settings	x	x
Printer Settings		x
Task Manager		x
Control Panel		x
Command Prompt		x
Windows Explorer		x
Import / Export Settings		x

Configuring local applications when using Parallels Client desktop replacement

With the **Replace Desktop** option enabled, the administrator's goal should be to deploy remote applications or remote desktops and use the native OS to simply deploy the software needed to connect remotely. However, in some instances, local applications may be required. The administrator still has the ability to configure local applications to be shown within the Parallels Client Desktop Replacement, however it is necessary to switch to the Admin mode prior to it.

Publish a local application according to the following steps:

- 1 Shadow the user's session or use the user device station directly.
- 2 Switch the Parallels Client Desktop Replacement to admin mode.
- 3 Click **File > Add New Application**
- 4 Fill in the application information

- 5 Applications added will be visible in the Application Launcher.
- 6 Switch back to user mode once all the applications needed are configured.

Windows desktop replacement

This section explains what happens when the **Replace Desktop** option is enabled, and why it is useful to administrators.

When enabled, the Replace Desktop feature allows the administrator to convert a standard desktop into a limited device similar to a Thin Client, without replacing the operating system.

The end user will not have access to Windows Explorer, Taskbar or any other Windows components that usually allow them to install new applications or change system settings. The user can now only deploy applications configured within the Parallels Client, including remote applications, remote desktops, and locally configured applications. Local applications are allowed, so if a specific application is needed, but is not available remotely (e.g. a software which communicates with specific peripherals), the user can still deploy it.

When the **Replace Desktop** option is enabled, the following features take effect on the corresponding versions of Windows (7, 8, 8.1, 10, 11):


Feature	7	8	8.1	10	11
Replace Desktop with Parallels Client	x	x	x	x	x
Disable Start Button	x	x	x	x	n/a
Restrict Control Panel Access	x	x	x	x	x
Disable Windows Key	x	x	x	x	x
Disable the Task Manager	x	x	x	x	x
Disable Quick Access Toolbar	n/a	n/a	n/a	n/a	n/a
Disable Security Manager/Action Center Notifications	x	x	x	x	x
Lock the Taskbar	x	x	x	x	x
Remove Pinned Applications	x	x	x	x	x
Disable Metro Screen (user logs directly to desktop)	n/a	x	x	x	x
Disable Hot Corners	n/a	x	x	x	x
Disable Charm Hints	n/a	x	x	x	x
Disable Help Aids	n/a	x	x	x	x
Disable Windows Sidebar	x	n/a	n/a	n/a	n/a

In this mode, the user also has access to the Mouse and Display Control Panel applets. The user cannot change the Parallels Client Global Options and the Client Farm Connection Options. Advanced management features can be enabled if the device is switched into administration mode.


If the Windows Desktop Replacement feature is switched off, all the restrictions are removed and the standard desktop is made available to the user.

The following are the screenshots of a Windows 10 desktop before and after the **Replace Desktop** option is enabled.

Before



After



Scheduling Windows devices & groups power cycles

The **Scheduler** tab of the **Device Manager** category can be used to schedule automatic power operations on devices.

Adding a new scheduler task

To schedule a task:

- 1 On the **Scheduler** tab, click **Tasks > Add** to open the **Device Scheduler Properties** dialog.
- 2 Select the **Enable this scheduled entry** option.

- 3 Select an action in the **Action** drop-down list:
 - **Device group switch on**
 - **Device group log off**
 - **Device group switch off**
 - **Device group reboot**
 - **Device group lock**
- 4 Select a device group in the **Target** drop-down list.
- 5 Specify the task start date and time.
- 6 Select the **Repeat** option from the following choices:
 - **Never** (a task will run only once, as specified in the **Start** and **Time** fields)
 - **Every day**
 - **Every week**
 - **Every 2 weeks**
 - **Every month**
 - **Every year**
 - **On specific day(s) of the week.** When selecting this option, select the day(s) of the week.
- 7 Enter a task description in the **Description** field.
- 8 Click **OK** to create the task.

Managing scheduled tasks

To modify an existing task, right-click it in the **Schedule List** and click **Properties** in the context menu.

To enable or disable an event, right-click it, click **Properties**, and then select or clear the **Enable this scheduled entry** option.

To execute a scheduled task immediately, right-click it and click **Execute Now** in the context menu.

To delete a task, right-click it and then click **Delete**.

Client Policies

The **Policies** category allows you to manage Parallels Client policies for users connecting to a Farm. By adding client policies, you can group users and push different Parallels Client settings to user devices forcing them to function as your organization requires.

Settings that can be enforced on user devices include RAS connection properties, display, printing, scanning, audio, keyboard, device, and others. Once you create a policy and push it to a client device, the user of the device cannot modify the settings that the policy enforces. In Parallels Client this will manifest itself as hidden or disabled connection properties and global preferences.

Supported Parallels Client versions

Parallels Clients for all platforms are supported.

Note: Starting with Parallels RAS v16.5, a new approach is used to manage client policies. In the previous versions, a client policy would apply the full set of parameters and replace the client settings completely hiding an enforced category. In RAS v16.5 and newer, client policy settings are split into smaller groups with the ability to configure and enforce each group on the client side individually. For the information on how this affects existing client policies that were created in earlier version of Parallels RAS, please read **Client policy backward compatibility** (p. 407).

In this section:

- Add a new client policy (p. 387)
- Configure session settings (p. 388)
- Configure client policy options (p. 401)
- Configure control settings (p. 405)
- Configure gateway redirection (p. 406)
- Client policy backward compatibility (p. 407)

Add a new client policy

To add a new client policy:

- 1** Select the **Policies** category and then click **Tasks > Add** in the right pane. The **Policy Properties** dialog opens.
- 2** The left pane contains a navigation tree allowing you to select a group of options to configure.
- 3** Make sure the **Policy** node is selected and then specify a policy name and an optional description.
- 4** In the **Browse Mode** drop-down list, select how you want to browse for users and groups. The preferred mode is **Secure Identifier** (default). Other options exist for backward compatibility.
- 5** In the **Apply policy to** section, click **Tasks > Add** (or click the plus sign icon) and specify the target users, computers, or groups. Note that in addition to users, user groups and security principles, you can search for and specify Active Directory computer accounts and computer security groups.

Configure criteria for the client policy

By default, a client policy applies to configured users, computers, and groups in all situations. You can optionally define criteria when the policy should be applied. This functionality allows you to create different policies for the same user or computer, which will be applied depending on where the user is connecting from and from which device.

To create new criteria:

- 1 Select **Criteria** (under the **Policy** node) in the left pane.
- 2 In the "gateway criteria" section, select the criteria type in the first drop-down list and then specify the values (if applicable) in the second drop-down list.
- 3 In the "MAC address criteria" section, select the criteria type in the first drop-down list and then specify the values (if applicable) in the second drop-down list.
- 4 In the "Parallels Clients" section, select the version of Parallels Client to which this policy should apply.

Configure session settings

Items under the **Session** node in the **Policy Properties** dialog include connection, display, printing, network, and other settings that will be enforced on a client if defined and enabled.

For a particular group of settings to be enforced on a client device, it must be selected (checked). Unselected groups will not be enforced, so end users will be able to configure them themselves. For example, you can check the **Connection** node, but only check the **Primary connection** and **Secondary connections** groups under it. This will enforce only the two selected groups of settings on client devices.

In this section:

- Connection (p. 389)
- Display (p. 392)
- Printing (p. 392)
- Scanning (p. 395)
- Audio playback (p. 395)
- Keyboard (p. 396)
- Local devices and resources (p. 396)
- Experience (p. 398)
- Network (p. 399)
- Server authentication (p. 399)
- Advanced settings (p. 400)

Connection

To configure connection properties, select the **Connection** node and then go through each child node configuring their respective properties.

Primary connection

The primary connection always defaults to the primary RAS Secure Client Gateway, but you can modify the following connection properties:

- 1 Specify a friendly name for the connection.
- 2 **Auto login:** Enable or disable auto login in RAS HTML5 Client. If the option is disabled, auto login will be disabled in HTML5 Client and the user will not be able to change it. For more information, see **Auto Login** (p. 359).
- 3 In the **Authentication type** drop-down list, select the desired method of authentication:
 - **Credentials.** The user will have to enter credentials to log on.
 - **Single Sign-On.** This option will be included in the list only if the Single Sign-On module is installed during Parallels Client installation. The credentials that the user used to log on will be used to connect to the remote server.
 - **Smart Card.** Select this option to authenticate using a smart card. When connecting to the remote server, a user will need to insert a smart card into the card reader and then enter a PIN when prompted.
 - **Web.** If selected, the SAML SSO authentication is allowed. For more info, see **SAML SSO Authentication** (p. 319).

Note: The allowed authentication type(s) must be specified in the RAS Console in **Connection > Authentication**.

- 4 Select or clear **Save password** as needed (if credentials are used for authentication). This means forcing a client to save the password for this connection.
- 5 Specify the domain name (if credentials are used for authentication).

Secondary Connection

If you have more than one RAS Secure Client Gateway, you can define a secondary connection, which will be used as a backup connection in case the primary gateway connection fails.

To add a secondary connection:

- 1 Select the **Secondary connections** item.
- 2 In the **Secondary connections** pane, click **Tasks > Add** and specify a server name or IP address.
- 3 Select the connection mode and modify the default port number if necessary.

If you have multiple secondary connections, you can move them up or down in the list. If the primary connection cannot be established, Parallels Client will use secondary connections in the order listed.

Reconnection

In this pane, specify what to do if the connection is dropped:

- **Reconnect if connection is dropped.** If this option is selected, Parallels Client will try to reconnect if the connection is dropped. The **Connection retries** property specifies the number of retries.
- **Show connection banner if reconnection is not established within.** Specifies the number of seconds after which the connection banner will be displayed in Parallels Client. This will inform the user that the connection was dropped and will allow them to take actions on their own.

Computer name

Specify the name that a computer will use during a remote desktop session. If set, this will override the default computer name. Any filtering set by the administrator on the server side will make use of the **Override computer name** setting.

Advanced settings

- **Connection timeout.** The Parallels Client connection timeout value.
- **Show connection banner if connection is not established within.** Specifies the number of seconds after which the connection banner will be displayed. This will inform the user that the connection cannot be established and will allow them to take actions on their own.
- **Show desktop if published application does not start within.** If a published application is not launched within the time period specified in this field, the host server desktop will be shown instead. This is helpful if an error occurs on the server side while launching an application. By showing the server desktop, the user can see the error message.

Web authentication

- Select or clear the **Use default OS browser** option. If the option is selected, the SAML SSO login dialog on the client side will open in the default browser. If the option is cleared, the browser built into the Parallels Client will be used.
- The **Open browser window to complete log out** option is used when the built-in browser is used. In this case, there's no control over the SAML log out, so when this option is selected, a URL will open to perform the logout from SAML. By default, this web page will not be displayed, but if you need to interact with the browser, you can enable this option.

For more info, see **SAML SSO Authentication** (p. 319).

Session Prelaunch

When a user opens a remote application, a session must first be launched. Launching a session can take time, which will result in the user waiting for the application to start. To improve user experience, a session can be launched ahead of time, before the user actually opens an application.

To enable (or disable) session prelaunch, choose one of the following in the **Mode** drop-down list:

- **Off.** No session prelaunch is used.
- **Basic.** A session is prelaunched as soon as the user gets the application listing. The assumption is, the user will open an application within the next few minutes. The session will stay active for 10 minutes. If the user doesn't open an application during that time, the client will disconnect from the session.
- **Machine Learning.** When the application listing is acquired, a session is prelaunched based on user habits. With this option enabled, Parallels Client will record and analyze the user habits of launching applications on a given day of the week. A session is started a few minutes before the user usually opens an application.

When a session is prelaunched, it will all happen in the background, so the user will not see any windows or message boxes on the screen. When the user starts an application, it will open using the prelaunched session, so it will start very quickly.

You can configure rules when session prelaunch must not be used. The following options are available:

- Use the **Exclude sessions prelaunch** list to specify dates on which the prelaunch must not be used. Click on the plus-sign icon and select a date. The list can contain multiple entries.
- You can also exclude a published resource from the session prelaunching scheme altogether. This way, the resource is excluded from the analysis and is never considered by Parallels Client when making a decision whether to prelaunch a session. For example, when you have a server on which you never want to prelaunch sessions, you can flag all published resources hosted by that server as to be excluded from session prelaunch. To exclude a published resource from session prelaunch, in the RAS Console, navigate to **Published Resources**, select a resource and then select the **Exclude from session prelaunch** option.

Local proxy address

The setting in this section specifies on which IP address to bind the local RDP proxy. Select the **Use 127.0.01 IP address when using Gateway mode in VPN scenarios** option. You should have this setting enabled. Disabling it may lead to users not being able to open applications or desktops when using a VPN. This setting applies to Parallels Client for Windows only.

Display

To configure display settings, select the **Display** node and then configure the groups of settings described below.

Settings

Select the desired video acceleration mode and color depth.

Multi-monitor

Specify whether all monitors should be used for a desktop session if more than one monitor is connected to the user's computer.

Published applications

Select the **Use primary monitor only** option to start published applications on the primary monitor. Other monitors connected to a user's computer will not be used.

Desktop options

Specify the desktop options as follows:

- **Smart-sizing:** Choose a smart sizing option. The **Scale (fit to window)** option scales a remote desktop to fit the connection window. The **Resize (update resolution)** option updates the resolution dynamically (without the need to reconnect) based on the window size. To disable smart sizing, select **Disabled**.
- **Embed desktop in launcher.** Enable this option to access a published desktop inside Parallels Client.
- **Span desktop across all monitors.** Enable this option to span published desktops across all connected monitors.
- **Connection bar in full screen.** Specify whether the connection bar should be pinned, unpinned, or hidden when connecting in full screen mode.

Browser

This section applies to Parallels HTML5 client only. Specify whether a remote application should open in the same or a new tab in a web browser by default.

Printing

The **Printing** pane allows you to configure printing options.

In the **Technology** section, select the technology to use when redirecting printers to a remote computer:

- **None.** No printer redirection will be used.
- **RAS Universal Printing technology.** Select this option if you want to use RAS Universal Printing technology.
- **Microsoft Basic Printing Redirection technology.** Select this option if you want to use Microsoft Basic printing technology.
- **RAS Universal Printing and Microsoft Basic redirection technologies.** Select this option to use both Parallels RAS and Microsoft technologies.

Note: The following rules apply when using printing in RAS HTML 5 Client. If **None** or **Microsoft Basic Printing** is selected, then no printing redirection will be available in a remote session. If **RAS Universal Printing** or **RAS Universal Printing and Microsoft Basic Printing** is selected, then RAS Universal Printing will be used in a remote session.

RAS Universal Printing

If you selected **RAS Universal Printing technology**, use the **Redirect Printers** drop-down list to specify whether to redirect all printer on the client side, default printer only, or specific printers.

If you select **Specific only** in the step above, click **Tasks > Add**. Type a printer name and then click the **Options** button. In the dialog that opens, specify settings described below.

In the **Choose Format** drop-down list, select a data format for printing:

- **Print Portable Document Format (PDF).** Adobe PDF. This option does not require you to install any local applications capable of printing a PDF document. All the necessary libraries are already installed together with Parallels Client.
- **View PDF with external application.** To use this option you must have a local application installed which is capable of viewing a PDF document. Note that not all applications are supported. For example, the built-in PDF viewer in Windows is not supported, so you must have Adobe Acrobat Reader (or a similar application) installed.
- **Print PDF with external application.** This option works similar to the View PDF option above. It also requires an application capable of printing a PDF document installed locally.
- **Enhanced Meta File (EMF).** Use vector format and embedded fonts.
- **Bitmap (BMP).** Bitmap images.

In the **Client printer preferences** section, select one of the following:

- **Use server preferences for all printers.** If this option is selected, a generic printer preferences dialog will be shown when a user clicks **Print** in a remote application. The dialog has only a minimal set of options that they can choose.

- **Use client preferences for all printers.** With this option selected, a local printer preferences dialog will open when a user clicks **Print** in an application. The dialog will contain a full set of options for a particular printer that the user has installed on their local computer. If they have more than one printer installed, a native preferences dialog will open for any particular printer that they choose to print to.
- **Use client preferences for the following printers.** This option works similar to the **Use client preferences for all printers option** (above), but allows users to select which printers should use it. Select this option and then select one or more printer in the list below. If a printer is not selected, it will use the generic printer preferences dialog, similar to the first option in this list.

Default printer settings

To configure default printer settings, click the **Change Default Printer settings** button.

The default printer list shows printers that can be redirected by the client to the remote computer:

- To disable the default printer, select **<none>**.
- To redirect the default local printer, select **<defaultlocalprinter>**.
- When **<custom printer>** is selected, you can specify a custom printer. The first local printer that matches the printer name inserted in the **Custom** field will be set as the default printer on the remote computer.

Select **Match exact printer name** to match the name exactly as inserted in the **Custom** field. Please note that the remote printer name may not match the original printer name. Also note that local printers may not redirect due to server settings or policies.

The **Force Default printer for** option specifies the time period, during which a printer will be forced as default. If the default printer is changed during this time after the connection is established, the printer is reset as default.

Select the **Update the remote default printer if the local default printer is changed** option to change the remote default printer automatically when the local default printer is changed. Please note that the new printer must have been previously redirected.

A Windows 10 and 11 note

Windows 10 and 11 have a feature that automatically sets the default printer to the one used most recently or more often. This can break the default printer control on RD Sessions Hosts, guest VMs, and Remote PCs. To resolve this issue, the default printer management in Windows 10 and 11 should be disabled. To disable this feature using the Group Policy, do the following:

- 1 Open the group policy editor.
- 2 Navigate to **User Configuration > Administrative Templates > Control Panel > Printers**.
- 3 Find the **Turn off Windows default printer management** policy and enable it.
- 4 Force the group policy to all computers attached to the domain.

You can also disable the default printer management in Windows 10 and 11 locally by using the GUI or the registry editor:

- 1 On a Windows 10 or 11 computer, click **Start**, then click the "gear" icon which will open the **Settings** page.
- 2 On the **Printers and Scanners** tab, set the **Let Windows manage my default printer** option to **OFF**.

Using the registry editor:

- 1 Open the registry editor (regedit).
- 2 Navigate to HKEY_CURRENT_USER\Software\Microsoft\Windows NT\CurrentVersion\Windows.
- 3 Create a new DWORD item and name it **LegacyDefaultPrinterMode**.
- 4 Change the item's Value data to hexadecimal and set the value data to **1**.

In addition to disabling the default printer management, the **Download over metered connections** option should be enabled in **Settings > Devices > Printers & Scanners**.

Scanning

On the **Scanning** pane, you can specify a scanner that should be used when one is required by a published application:

- **Use.** Allows you to select a scanning technology. RAS Universal Scanning uses TWAIN and WIA redirection allowing an application to use either technology depending on the hardware type connected to the local computer. If you select **None**, scanning will be disabled.
- **Redirect Scanners.** Select scanners attached to your computer for redirection. You can select **All** (all attached scanners will be redirected) or **Specific only** (only the scanners you select in the provided list will be redirected).

Audio

This pane allows you to configure remote audio playback and recording settings.

In the **Remote audio playback** section, Use the **Where** drop-down list to select one of the following remote audio playback options:

- **Bring to this computer.** Audio from the remote computer will play on your local computer.
- **Do not play.** Audio from the remote computer will not play on your local computer and will be muted on the remote computer as well.
- **Leave at remote computer.** Audio will not play on your local computer but will play normally on the remote computer.

Use the **Quality** drop-down list to adjust the audio quality:

- **Dynamically adjust based on available bandwidth.** This option will increase or decrease the audio quality based on your connection speed. The faster the connection, the higher audio quality setting will be used.
- **Always use medium audio quality.** The audio quality is fixed at the medium level. You can use this option when you don't require the best possible audio quality and would rather use the available bandwidth for graphics.
- **Always use uncompressed audio quality.** The audio quality is fixed at the highest level. Select this option if you have a very fast connection and require the best possible audio quality.

The **Enable recording (if applicable)** option allows you to enable audio recording on the remote computer. For example, you can speak into a microphone on the local computer and use a sound recording application on the remote computer to record yourself.

Keyboard

On the **Keyboard** pane, select how you want to apply key combinations (e.g. Alt+Tab) that you press on the keyboard:

- **On the local computer.** Key combinations will be applied to Windows running on the local computer.
- **On the remote computer.** Key combinations will be applied to Windows running on the remote computer.
- **In full screen mode only.** Key combinations will be applied to the remote computer only when in the full-screen mode.

Select or clear the **Send unicode characters** as needed.

Local devices and resources

Use the **Local devices and resources** pane to configure how local resources are used in a remote session.

Clipboard

Enable or disable the clipboard in a remote session. In the right pane, choose one of the following clipboard redirection options:

- **Client to server only:** Copy and paste from client to a server app only.
- **Server to client only:** Copy and paste from a server app to client only.
- **Bidirectional:** Copy and paste in both directions.
- **Disabled:** The clipboard is disabled.

Note: When you clear this option, it will also disable the Remote Clipboard functionality for affected users in Parallels HTML5 client. For more information, please see **Using the remote clipboard** (p. 364).

Disk drives

Select the **Allow disk drives redirection** option and select local drives you want to redirect, or select **Use all disk drives available**.

If you select the **Use also disk drives that I plug in later** option, disk drives that you connect to a local computer later will be automatically available in a remote session. Note that this option applies to Parallels Client for Windows only.

Devices

On this pane, specify whether to redirect local devices in general, use all devices available, and also devices that will be plugged in later.

Local devices that can be redirected include supported Plug and Play devices, media players based on the Media Transfer Protocol (MTP), and digital cameras based on the Picture Transfer Protocol (PTP).

Please note that disk drives and smart cards are redirected using dedicated **Disk drives** and **Smart cards** options.

Video capture devices

Specifies video capture devices to redirect from a user device to the remote session. This is a high-level redirection that allows to redirect a composite USB device, such as a webcam with a microphone.

- **Allow devices redirection:** Allows to choose which video capture devices to redirect.
- **Use all devices available:** Redirect all available devices.
- **Use also devices that I plug in later:** A device that is plugged in after a session is started will also be used. Note that if this option is disabled, you will need to restart a session for a newly plugged in device to become available.

Ports

Select whether to redirect LPT and COM ports.

Smart cards

Select whether to redirect smart cards. Note that if smart card is selected as the authentication type in the **Primary connection** pane, the smart card redirection is automatically enabled and this option is grayed out.

Windows touch input

Enables or disables the following functions:

- Pen input redirection with pressure sensitivity support.

Note: Pen input redirection is supported on remote desktops with the following operating systems: Windows Server 2016 and newer, Windows 10 version 1607 and newer.

- Windows touch input redirection. Windows touch input redirection allows users to use Windows native touch gestures from touch-enabled devices, including touch, hold, and release actions. The actions are redirected to remote applications and desktops as corresponding mouse clicks. This option allows you to disable touch input redirection in case of app compatibility issues.

Note: This policy is applicable to Parallels Client for Windows and Parallels HTML5 Client only.

File transfer

Enables file transfer in a remote session. To enable file transfer, select this node and then select a desired option in the **Allow file transfer** drop-down list in the right pane. For additional information, see **Configuring remote file transfer** (p. 408).

Experience

The **Experience** pane allows you to tweak connection speed and compression.

Performance

Choose your connection speed to optimize performance: Choose a connection type according to your situation and then select experience options you want enabled. If you are connecting to a remote server on a local network that runs at 100 Mbps or higher, it is usually safe to have all of the experience options enabled. If you choose **Detect connection quality automatically**, the experience options will be enabled by default, but some may be dynamically disabled depending on the actual connection speed.

Enhance windows move/size: Enable this option if your users experience graphics artifacts (dark squares) while moving or resizing a remote application window on their desktops. The issue may manifest itself when a remote application is hosted on a Windows Server 2016, 2019 or 2022 and when the **Show contents of window while dragging** option is enabled. The issue does not appear with any other versions of Windows.

Compression

It is recommended to enable compression to have a more efficient connection. The available compression options are described below.

Enable RDP Compression: Enables compression for RDP connections.

Universal printing compression policy: The compression type should be selected based on your environment specifics. You can choose from the following options:

- **Compression disabled.** No compression is used.
- **Best speed (uses less CPU).** Compression is optimized for best speed.
- **Best size (uses less network traffic).** Compression is optimized to save network traffic.
- **Based on connection speed.** The faster the connection speed, the lower compression level and the minimum data size to compress are used.

Universal scanning compression policy: This drop-down list has the same options as the universal printing compression above. Select the compression type based on your environment specifics.

Network

Use the **Network** pane to configure a proxy server if you have one.

Select the **Use proxy server** option and then select the protocol from the following list:

- **SOCKS4.** Enable this option to transparently use the service of a network firewall.
- **SOCKS4A.** Enable this option to allow a client that cannot connect to resolve the destination host's name to specify it.
- **SOCKS5.** Enable this option to be able to connect using authentication.
- **HTTP 1.1.** Enable this option to connect using a standard HTTP 1.1 protocol connection.

Specify the proxy host's domain name or IP address and the port number.

For SOCKS5 and HTTP 1.1 protocols, select the **Proxy requires authentication** option. For authentication, select the **Use user logon credentials** option or specify a user name and password in the fields provided.

Server authentication

Use the **Server authentication** pane to specify what should happen if authentication of an RD Session Host, Remote PC, or Guest VM fails.

In the **If authentication fails** drop-down list, select one of the following options:

- **Connect.** The user can ignore the certificate of the server and still connect.
- **Warn.** The user is alerted about the certificate and still has the ability to choose whether to connect or not.
- **Do not connect.** The user is not allowed to connect.

Advanced settings

The **Advanced Settings** pane allows you to customize the default behavior of Parallels Client.

You can specify the following properties:

- **Use client system colors:** Enable this option to use the client system colors instead of those specified on the remote desktop.
- **Use client system settings:** Enable this option to use the client system settings instead of those specified on the RD Session Host.
- **Create shortcuts configured on server:** For each published application, the administrator can configure shortcuts that can be created on the client's desktop and the Start menu. Select this option to create the shortcuts, or clear the option if you don't want to create them.
- **Register file extensions associated from the server:** For each published application, the administrator can create file extension associations. Use this option to either register the associated file extensions or not.
- **Redirect URLs to the client device:** Enable this option to use the local web browser when opening 'http:' links.
- **Redirect MAILTO to the client device:** Enable this option to use the local mail client when opening 'mailto:' links.
- **Always ask for credentials when starting applications:** If this option is enabled, a user will be asked to enter credentials when starting an application even if the session is still active. You can use this option as added security to prevent unauthorized users to access applications. For example, if a user disconnects from a session, no one else will be able to take over the session and run remote applications. As another example, if a user leaves a device with an open HTML5 Client displaying the app listing (with or without running RDP sessions) then any user who tries to open a new application or another instance of a running application will be prompted for credentials. Please note that the **Auto login** option (p. 389) must be disabled for this functionality to work; otherwise saved credentials will be used automatically.
- **Allow Server to send commands to be executed by client:** Enable this option to allow commands being received from the server to be executed by the client.
- **Confirm Server commands before executing them:** If this option is enabled, a message is displayed on the client to confirm any commands before they are executed from the server.
- **Network Level Authentication:** Check this option to enable network level authentication, which will require the client to authenticate before connecting to the server.
- **Redirect POS devices:** Enables the Point of Service (POS) devices such as bar code scanners or magnetic readers that are attached to the local computer to be used in the remote connection.
- **Use Pre Windows 2000 login format:** If this option is selected, it allows you to use legacy (pre-Windows 2000) login format.

- **Disable RDP-UDP for gateway connections:** Disables RDP UDP data tunneling on the client side. You can use this option when some clients experience random disconnects when RDP UDP data tunneling is enabled on the RAS Secure Client Gateway (the **Network** tab in the gateway **Properties** dialog), while other clients are not.
- **Do not show drive redirection dialog:** This option affects Parallels Client for Mac. By default, the **Grant access to Home folder** (drive redirection) dialog opens automatically when a Mac user connects to Parallels RAS. This happens when this option is disabled or when there's no client policy at all. The dialog allows the user to configure which folders on the local disk drive should be available to remote applications. If you enable this option, the dialog will not be shown a user. Read below for more explanation.

Drive redirection cannot be configured via client policies, so Mac users have to do this themselves. By automatically showing the dialog, you can invite the user to go through the local folder configuration procedure. On the other hand, if there's no need for your users to redirect their local drives, you can disable the automatic opening of the dialog. Note that the dialog can still be run manually in Parallels Client for Mac at any time by opening **Connection Properties > Local Resources**, selecting the **Disk drives** option and clicking **Configure**.

When the option is disabled (or when there's no client policy defined), the dialog opens at least once when the user connects to Parallels RAS for the first time. At that time, the user can either configure local folders or select the **Never ask me** again option. In both cases, the dialog will not be shown to the user anymore. The Mac user can reset the **Never ask me** selection by going to **Connection Properties > Advanced** and clearing the **Do not show drive redirection dialog** option.

Configure client policy options

The **Client options** node allows you configure client policy options. Select the node and then select and configure individual items under it as described below.

Connection

On the **Connection** pane, specify the following options:

- **Connection Banner.** Select a banner to display while establishing a connection.
- **Automatically refresh connected RAS connections every [] minutes.** Select this option and specify the time interval to automatically refresh a connection. This will refresh the published resources list in Parallels Client.

Logging

Specify a log level for Parallels Client. Choose from the following options:

- **Standard**
- **Extended**
- **Verbose**

You should normally use the Standard logging. When you have an issue with Parallels Client, you can temporarily raise the log level by selecting Extended or Verbose and setting start date/time and a duration. Note that start date and time correspond to the local client time zone. Parallels Client must be running in order for the logging to take place. If Parallels Client is launched when Extended or Verbose levels should be already in effect, the level will stay on for the remainder of the original duration setting. If a policy changes during this time, the actual log level settings will be reapplied accordingly.

Update

Select **Check for updates on startup** and specify an update URL if you want Parallels Client to check for updates when it starts. The URL can point to the Parallels website or you can store updates on your local network and use this local URL. For the information on how to configure a local update server, please read <https://kb.parallels.com/123658>.

Note: This option works with Parallels Client for Windows only. Parallels Client for Mac can be updated only from the App Store. Parallels Client for Linux does not support this feature.

PC keyboard

To force a particular keyboard to be used, select the Force use PC keyboard and select a keyboard layout from the drop-down list. Note that the selected layout can and will only be used in a Parallels Client version that supports this particular layout.

Single sign-on

Parallels Client for Windows comes with its own SSO component that you can install and use to sign in to Parallels RAS. If you already use a third-party credential provider component on your Windows computers, you first need to try if the single sign-on works right out of the box. If it doesn't, you need to configure Parallels RAS and Parallels Client to use the Parallels RAS SSO component to function as a wrapper for the third-party credential provider component.

To use Parallels RAS SSO as a wrapper, specify a third-party component, select the **Force to wrap third party credential provider component** option and specify the component's GUID in the field provided. You can obtain the GUID in Parallels Client as follows:

- 1 Install Parallels Client on a computer that has the third-party component installed.
- 2 In Parallels Client, navigating to **Tools > Options > Single Sign-On** (tab page).
- 3 Select the "Force to wrap..." option and then select your provider in the drop-down list.
- 4 Click the **Copy GUID to Clipboard** button to obtain the component's GUID.

You will also need to specify the component's GUID when setting up an invitation email in the RAS Console. If you haven't set up an invitation email yet, you can do it as follows:

- 1 In the RAS Console, select the **Start** category and then click the **Invite Users** item in the right pane.
- 2 On the second page of the wizard (target platform and connection options), click the **Advanced** button.
- 3 In the dialog that opens, select the **Force to wrap third party SSO component** option and specify the GUID of the component.

For more information, see the **Invite users** section (p. 37).

After the policies are applied on Windows computers, Parallels Client will be automatically configured to use the specified third-party credentials provider.

Advanced

Use this pane to specify advanced client options, as described below.

Global

- **Always on Top.** With this feature enabled, other applications will no longer mask the launcher.
- **Show connection tree.** Displays the connection tree.
- **Minimize to tray on close or escape.** Enable this feature to place the Parallels Client into the System Tray when you click on the **Close** button or hit escape.
- **Enable graphic acceleration (Chrome client).**
- **Do not warn if server certificate is not verified.** When connected to a RAS Secure Client Gateway over SSL, and the certificate is not verified, a warning message will be displayed. You can disable this warning message by enabling this option.
- **Swap mouse buttons.** When enabling this setting, the mouse buttons will be swapped on the remote computer.
- **DPI aware.** This will force a published application to be DPI-aware depending on the client's DPI settings. This feature works on Windows 8.1 or higher.
- **Add RAS Connection automatically when starting web or shortcuts items.** This option will add the connection preferences in the Parallels Client when starting an item contained in a connection that is not yet listed.
- **Do not show prompt message for auto add RAS connection.** Enable this option to disable prompt messages when adding auto connections.
- **Close error messages automatically.** When a session disconnects because of an error, the error is automatically dismissed after 15 seconds.

- **Clear session cookies on exit.** When a user logs on, a Parallels RAS logon cookie is kept on the client side. This will allow the user to connect again with Parallels RAS without re-authenticating. Check this option to delete any cookies when the user closes the Parallels Client.
- **Enable extended logging.** Enables extended logging.

Language

Specify a language that Parallels Client should use. The **Default** option uses the main language used by the client's operating system.

Printing

- **Install missing fonts automatically.** If automatic fonts are installed on the server, they will be available when a session connects.
- **Redirect vendor paper sizes for RAS Universal Printing.** When enabling this setting, non-standard paper sizes which are not included in the standard options will be redirected to the client. Sizes may vary depending on the vendor.
- **Raw printing support.** When enabling this setting, printing will still work for applications sending data in RAW format.
- **Convert non distributable fonts data to images.** During RAS Universal Printing, if a document includes non-distributable fonts, each page is converted to an image.
- **Cache printers hardware information.** Caching of printer hardware information locally to speed-up RAS universal printer redirection.
- **Refresh printer hardware information every 30 days.** Forces the printer hardware information cache update even if nothing has changed in 30 days. When this option is off, the cache will only be refreshed if there were known changes.
- **Cache RAS Universal Printing embedded fonts.** Caching of embedded fonts locally to speed-up RAS universal printing process time.

Windows client

- **Hide Launcher when application is launched.** If this option is enabled, the launcher will be minimized in the system tray after an application is launched.
- **Launch automatically at Windows startup.** This option will place a shortcut in the start menu folder of the client and the Parallels Client will launch automatically on Windows startup.

RemoteFX USB redirection

- **Allow RDP redirection of other supported RemoteFX USB devices to all users.** This setting applies to Parallels Client for Windows only. Outside Parallels RAS, the standard RemoteFX USB redirection feature must be enabled via Group Policy in order to work. When you select the "Allow RDP redirection ..." option on this screen, it will do the same as GPO, which is update the corresponding registry setting in Windows on a client machine. Parallels Client for Windows relies on this feature to be enabled in Windows registry in order to redirect USB devices. When the policy containing this setting is applied on a client machine, the user will see a message that RemoteFX USB redirection was enabled and that they will need to restart Windows.

Configure control settings

Control settings options allow you to control various actions on the client side. These options affect the following Parallels Clients:

- Windows
- Linux
- Mac
- Android
- iOS

Connections

On the **Connections** pane, select (or clear) the following options:

- **Prohibit adding of RAS connections.** When a user presses the **Add Connection** button, an RDP connection is always created.
- **Prohibit adding standard RDP connections.** When a user presses the **Add Connection** button, a RAS connection is always created

Password

On the **Password** pane, specify the following options:

- **Prohibit saving password.** The option to save the password will not be shown to the user for that particular connection. A password is never saved on a disk, but kept in memory until the user closes the application.
- **Prohibit changing password.** The option to change the password will not be shown in the context menu for that particular connection.

Import and export

On the **Import and Export** pane:

- **Prohibit import/export connection setting.** If this option is selected, the **Import** and **Export** buttons will not be shown to the user.

Configure Gateway redirection

Redirection options allow you to move your existing users from one RAS Secure Client Gateway to another gateway within the same Farm, or you can even redirect users to a gateway in a different Farm.

Note: When setting gateway redirection, make sure that the gateway criteria (the **Criteria** node) does not conflict with it. Read the **Gateway criteria** subsection at the end of this section for the explanation.

To configure redirection options:

- 1 Select the **Redirection** node in the left pane of the **Policy Properties** dialog.
- 2 In the right pane, specify the new connection properties, including:
 - **Gateway address**
 - **Connection mode**
 - **Port number**
 - **Alternative address**

When this policy is applied to user devices, the following will happen:

- Parallels Client connection settings are automatically updated on each device.
- Parallels Client tests the new connection. If succeeded, the current connection policies are removed and new policies are added.
- If Parallels Client cannot connect to Parallels RAS using new settings, the application list will not be shown and an error message will be displayed saying that the redirection policy has failed to apply. The user will be advised to contact the system administrator.

Gateway criteria

If a policy has both **Redirection** and **Criteria** settings enabled and configured, a situation may occur when the policy is applied in an infinite loop on the client side, which will result in an error. Consider the following possible scenarios when this may happen:

- Parallels Client connects to gateway "A" and applies a policy, which redirects it back to gateway "A". This will continue to loop until Parallels Client gives up and displays an error to the user, which will say, "Failed to apply redirection policy....".

- Parallels Client connects to gateway "A" and applies policy "P1", which redirects it to gateway "B". As expected, Parallels Client connects to gateway "B" and applies policy "P2", which redirects it back to gateway "A" where it all began. This will also continue to loop until Parallels Client gives up and displays the same error message as described above.

Once again, this may only happen if the **Criteria** node is enabled and specified gateways conflict with each other. To avoid it, make sure that the **Gateway criteria** option on the **Criteria** pane is set to **if Client is connected to one of the following gateways** and that the same policy is not applied again when Parallels Client is redirected to a new gateway.

Client policy backward compatibility

Starting with Parallels RAS v16.5, a new approach is used to manage client policies. In the previous versions, a client policy would apply the full set of parameters and replace the client settings completely hiding an enforced category. In RAS v16.5 (or newer), client policy settings are split into smaller groups with the ability to configure and enforce each group on the client side individually. For example, the administrator wants to re-design the policies to disable clipboard redirection only, leaving the rest of the local devices and resources settings available for the end users to control. In the previous version, this would not be possible. The new design allows an administrator to easily achieve this goal.

This section explains how the backward compatibility is achieved with older clients and how new clients retain compatibility with older server-side installations.

The new client policies implementation handles compatibility issues as follows:

- All settings found in older policies are sent to the client as if being sent from an older Parallels RAS server. When a client receives the policy, the **Connection properties** and **Options/Preferences** settings are set correctly from the old design point of view. If, however, the policy is configured in such a way that the user cannot change anything, the entire tab will be hidden (no need to display the options if all of them are disabled).
- The Parallels RAS Console handles old-style policy settings as if they are new and displays them using the updated graphical user interface.
- In terms of policies, when a Parallels RAS v16.5 client connects to a previous version of Parallels RAS, the client keeps working normally and all of the policy settings are functioning as expected.

Policy information in Parallels Client

When a policy is applied to a user device, the information about it is displayed in Parallels Client. The information can be used to verify that the correct policy was delivered to a user device. The following information is included:

- **ID:** The policy ID as displayed in the **ID** field in the **Policies** list in the RAS Console.
- **Version:** The policy version number as displayed in the **Version** field in the **Policies** list in the RAS Console.

- **RAS Connection:** The name of the connection through which the policy was delivered. Displayed only on mobile devices and in HTML5 Client.

By comparing the information above in Parallels Client running on a user device and the information in the RAS Console, you can see which policy was applied to a user device.

To see the applied policy information for a connection:

- In Parallels Client for Windows / Mac / Linux, open the **Connection Properties** dialog. The information is displayed at the bottom of a tab page to which the policy was applied.
- In Parallels Client for Android, the information is displayed at the bottom of the **Settings** screen.
- In Parallels Client for iOS, open the Edit RAS Connection screen and tap **View Applied Server Policy** (as the bottom).
- In RAS HTML5 Client, the information is displayed in the **Settings** dialog.

Please note that when all of the connection properties in Parallels Client are managed through client policies, the user can still open the **Connection Properties** dialog, but it will contain a single tab displaying the applied policy information. If only some of the connection properties are managed through policies, the user will be able to see those tabs and the applied policy information that they contain.

When a policy includes global policy options, you can view the applied policy information in Parallels Client as follows:

- In Parallels Client for Windows and Linux, open the **Options** dialog (click **Tools > Options**).
- In Parallels Client for Mac, open **Preferences** (click **Parallels Client > Preferences**).

The applied policy information is displayed at the bottom of the dialog, similar to how it is displayed for the connection.

Configuring remote file transfer

Parallels RAS provides end users with the ability to transfer files remotely to and from a remote server.

Note: At the time of this writing, file transfer is supported in Parallels HTML5 Client and Parallels Client for Chrome only. Bidirectional file transfer is supported in Parallels HTML5 Client only. In Parallels Client for Chrome, you can only enable or disable file transfer.

To make the remote file transfer functionality flexible, Parallels RAS allows you to configure it on the following three levels:

- RD Session Host, VDI provider, or Remote PC (p. 409)
- RAS HTML5 gateway (p. 410)

- Gateway settings in HTML5 Client (p. 349)
- Client policy (p. 410)

File transfer settings that you configure on each level take precedence in the order listed above. For example, if you enable file transfer on a Parallels HTML5 gateway, but disable it on an RD Session Host, file transfer will be disabled for all users who connect to the given RD Session Host through the given HTML5 gateway. As another example, you can enable file transfer on an RD Session Host and then disable it for a particular Client policy (or an HTML5 gateway). This way you can control which clients can use file transfer and which cannot.

Read the subsequent sections to learn how to configure file transfer on each level.

Configure file transfer for a server

To configure remote file transfer for an RD Session Host, VDI provider, or Remote PC, do the following:

- 1 In the Parallels RAS Console, select the **Farm** category and then select a desired server type (RD Session Host, VDI provider, Remote PCs) in the middle pane.
- 2 Right-click a desired server in the right pane and choose **Properties**.
- 3 Select the **Agent Settings** tab.
- 4 Select the **Allow file transfer command** option and click the **Configure** button. A dialog opens where you can specify remote file transfer options as described below.
- 5 In the **Direction** drop-down list, select one of the following:
 - **Client to server only**: Transfer files from client to server only.
 - **Server to client only**: Transfer files from server to client only.
 - **Bidirectional**: Transfer files in both directions.
- 6 In the **Location** field, specify a UNC path to a folder to be used as the default upload location. This path will also be used as the default source location when a user tries to download a file from a remote server. You can select from one of the locations predefined in the drop-down list or you can specify your own. Standard Windows environment variables, such %USERNAME%, %USERDOMAIN%, %USERPROFILE%, can be used. If the location is not found during an upload or download operation, the standard (default) download location will be used.
- 7 The **Do not allow to change location** option prohibits the user to change the UNC path specified in the **Location** field. If the option is enabled, the user cannot select a different location while trying to upload or download a file. If the option is cleared, the user can specify a different location.

Important: Please note that the **Do not allow to change location** option cannot prevent the user from accessing the specified remote location directly. For example, a user can try to upload a file, note the default location's UNC path (to which he/she has access), then open it in File Explorer and copy it to any folder in his/her profile. To prevent such a scenario from happening, you need to implement additional measures to control locations other than the location that you specify here.

Configure file transfer for HTML5 Gateway

To configure remote file transfer for a RAS HTML5 Gateway, do the following:

- 1 In the Parallels RAS Console, navigate to **Farm** > <Site> > **Gateways**.
- 2 Right-click a desired RAS Secure Client Gateway in the right pane and choose **Properties**.
- 3 Select the **HTML5** tab.
- 4 Select the **Allow file transfer command** option and click the **Configure** button. In the dialog that opens, select one of the following:
 - **Client to server only**: Transfer files from client to server only.
 - **Server to client only**: Transfer files from server to client only.
 - **Bidirectional**: Transfer files in both directions.

For more information about configuring RAS HTML5 Gateway, see **Configure HTML5 Gateway** (p. 74).

Configure file transfer for a client policy

To configure remote file transfer for a client policy, do the following:

- 1 In the RAS Console, select the **Policies** category.
- 2 Right-click a desired policy in the right pane and choose **Properties**.
- 3 In the left pane, navigate to **Session** > **Local devices and resources**.
- 4 Select the **File transfer** node.
- 5 In the right pane, select one of the following in the **Allow file transfer** drop-down list:
 - **Client to server only**: Transfer files from client to server only.
 - **Server to client only**: Transfer files from server to client only.
 - **Bidirectional**: Transfer files in both directions.

For additional information about client policies, see **Client Policies** (p. 386).

Reporting

Parallels RAS Reporting is an optional RAS component that allows Parallels RAS administrator to run and view predefined and custom Parallels RAS reports. Predefined reports include user and group activity, device information, session information, and application usage. You can also create custom reports using your own criteria. Read this chapter to learn how to install and configure Parallels RAS Reporting and how to use it.

In This Chapter

System requirements.....	411
Install Microsoft SQL Server	413
Install Parallels RAS Reporting.....	416
Running Parallels RAS reports.....	418
GDPR compliance.....	421

System requirements

Operating system requirements

Parallels RAS Reporting can be installed on a server running one of the following Windows Server versions:

- Windows Server 2022
- Windows Server 2019
- Windows Server 2016
- Windows Server 2012 R2
- Windows Server 2012
- Windows Server 2008 R2
- Windows Server 2008

.NET Framework 3.5 and .NET Framework 4.5 or higher must be installed.

Microsoft SQL Server requirements

Parallels RAS Reporting can be used with the following Microsoft SQL Server versions:

- Microsoft SQL Server 2019
- Microsoft SQL Server 2017
- Microsoft SQL Server 2016
- Microsoft SQL Server 2014
- Microsoft SQL Server 2012
- Microsoft SQL Server 2008 R2 SP1
- Microsoft SQL Server 2008 SP1

Beginning with RAS 17.1, SQL Server Reporting Services (SSRS) and the SQL Server database engine can be deployed on separate hosts.

Using Microsoft SQL Server 2017 and 2019

Microsoft SQL Server 2017 and 2019 allow you to install the database engine and SQL Server Reporting Services (SSRS) on different hosts. Parallels RAS 17.1 (and newer) supports this deployment scenario and gives you the ability to use SQL Server Reporting Services and the SQL Server database engine installed on separate hosts.

Installation locations

RAS Reporting must be installed on the same server where SQL Server Reporting Services are running. Please note that if you have SSRS and the database engine installed on different hosts, RAS Reporting must be installed where the SSRS are installed.

The following table contains RAS and SQL Server version compatibility information and locations where components necessary to use RAS Reporting can be installed:

RAS Reporting version	SSRS version	SQL Server version	Installation locations
17.1, 18.0	2019	2019	SSRS - same host as RAS Reporting SQL Server - can be a different host
17.1, 18.0	2017	2019	SSRS - same host as RAS Reporting SQL Server - can be a different host
17.1, 18.0	2017	2019	SSRS - same host as RAS Reporting SQL Server - can be a different host
17.1, 18.0	2017	2017	SSRS - same host as RAS Reporting SQL Server - can be a different host

17.1, 18.0	2017	2016	SSRS - same host as RAS Reporting SQL Server - can be a different host
17.0, 18.0	2008 R2 - 2016	2008 R2 - 2016	SSRS and SQL Server on the same host

Microsoft SQL Server must be installed as a named instance (not default or unnamed instance), because an instance must have a name for RAS Reporting to work. You can specify an instance name when you install Microsoft SQL Server (or when you create a new SQL Server instance in a multi-instance scenario). For complete details, please read the **Install Microsoft SQL Server** section that follows this one.

Install Microsoft SQL Server

Note: For Parallels RAS installations running on multiple servers, it is recommended that Microsoft SQL Server is installed on a dedicated server.

In this section:

- Install Microsoft SQL Server 2016 or earlier (p. 413)
- Install Microsoft SQL Server 2017 or 2019 (p. 416)

Install Microsoft SQL Server 2016 or earlier

To install a SQL Server instance (SQL Server 2016 or earlier):

- 1 Run the Microsoft SQL Server installation program and select the **Custom** installation type. Wait for the necessary files to be downloaded to your computer.
- 2 Once the files are downloaded, the **SQL Server Installation Center** window opens.
- 3 On the **Installation** page, select **New SQL Server stand-alone installation or add features to an existing installation**.
- 4 Click **Next** and follow onscreen instructions until you get to the **Feature Selection** page.
- 5 On the **Feature Selection** page, make sure that at least the following SQL Server features are selected for installation:
 - **Database Engine Services**
 - **Reporting Services - Native**
- 6 Click **Next**.
- 7 On the **Instance Configuration** page, select the **Named instance** option and enter an instance name. When naming the instance, you have the following options:
 - Enter "RASREPORTING", which is the default instance name used by Parallels RAS Reporting. If you use this name, you don't have to specify it later when installing RAS Reporting and configure it in the RAS Console. This is the recommended option.

- If you would like to use a different name, you can do that, but you will have to make sure that you use this name when installing and configuring RAS Reporting. The instructions for installing RAS Reporting (described later in this chapter) indicate where the instance name must be specified, so if you follow the instructions, you won't miss it. Note that the instance name cannot contain dashes, dots and some other characters.

After entering the instance name, make sure that it is also set in the **Instance ID** field.

- 8** Click **Next** and proceed to the **Database Engine Configuration** page.
- 9** On the **Database Engine Configuration** page, select the **Server Configuration** tab and add the following users to the **SQL Server administrators** list:
 - Local administrator (e.g. Administrator)
 - AD administrator (if you are just testing Parallels RAS Reporting on a local server, you can exclude this account).
 - SYSTEM (click **Add**, type "SYSTEM", click **Check Names** and click **OK**; the account will appear in the list as "NT AUTHORITY\SYSTEM").
- 10** Complete the wizard using default settings on remaining pages.
- 11** Wait for the SQL Server installation to finish. On the **Complete** page, make sure that the installation was successful and exist the wizard.

Install SQL Server Management Tools

You should also install SQL Server management tools, specifically the SQL Server Management Studio. It is not required by RAS Reporting, but it is an essential SQL Server management tool that you might find useful. If you've never worked with SQL Server Management Studio before and not sure whether you need it, we suggest you install it. For example, you can use it to view RAS Reporting database tables, constraints, and stored procedures, which may help you better understand the RAS Reporting database design. The installation link is provided on the **SQL Server Installation Center** window.

Configure Microsoft SQL Server 2016 and earlier

When using Microsoft SQL Server 2016 and earlier version, it must be configured for remote connections as follows:

- 1** Open Microsoft SQL Server Management Studio.
- 2** Right-click on the server and select **Properties**.
- 3** Go to **Connections** and select **Allow Remote**.
- 4** Open SQL Server Configuration Manager and go to **SQL Server Network Configuration > Protocols** for RASREPORTING.
- 5** Right-click on **TCP/IP** and choose **Properties**.
- 6** Make sure the **Enabled** property is set to **Yes**.

- 7 Select the **IP Address** tab and locate the **IPAll** section. Set the **TCP Dynamic Ports** field to be blank and the **TCP Port** field to "1433".
- 8 Restart SQL Server. To do so, in the SQL Server Configuration Manager, right-click the SQL Server service and choose **Restart**.
- 9 After the restart, in the SQL Server Configuration Manager, right-click on **SQL Server Browser** and choose **Properties**.
- 10 Select the **Service** tab and set the **Start Mode** property to **Automatic**.
- 11 Start the SQL Server Browser.

Configure Microsoft SQL Server Reporting Services

To configure Microsoft SQL Server Reporting Services, follow these steps:

- 1 Run the Reporting Service Configuration Manager (**Start > Apps > Microsoft SQL Server 2016 > Reporting Services Configuration Manager**).
- 2 In the **Reporting Services Configuration Connection** dialog that opens, do the following:
 - Make sure the **Server Name** field contains the name of the server hosting the SQL Server instance.
 - Make sure the **Report Server Instance** field contains the name of the SQL Server instance that you've created earlier. If you used the default Parallels RAS name, it will appear as "RASREPORTING". If you used a different instance name, select that name.
- 3 Click **Connect**. If the connection is successful, the **Reporting Services Configuration Manager** window opens.
- 4 Select the **Web Service URL** category (not to be confused with Web Portal URL) in the left pane and set the following properties in the right pane:
 - **Virtual Directory**: Make sure that the directory name is "ReportServer_RASREPORTING". If you used a different name for the SQL Server instance, you should see that name instead of the "RASREPORTING" part.
 - **TCP port**: Set the port number to 8085.
- 5 Click the **Apply** button to apply the settings.
- 6 Select the **Web Portal URL** category in the left pane and then do the following:
 - Make sure that the **Virtual Directory** field is set to "Reports_<InstanceName>", where "InstanceName" is the name of your SQL Server instance. The default Parallels RAS name would be "Reports_RASREPORTING".
 - Examine the **URLs** field. Make sure that the port number after the server name is 8085. If it's not, click the **Advanced** button and change the port number.
- 7 Verify that you can access the Reporting Services Web Portal by clicking the URL on the **Web Portal URL** page. This should open the SQL Server Reporting Services home page in a web browser.
- 8 Click **Exit** to close the Reporting Services Configuration Manager.

Install Microsoft SQL Server 2017 or 2019

Microsoft SQL Server 2017 and 2019 allow you to install the database engine and SQL Server Reporting Services (SSRS) on separate hosts. Parallels RAS 17.1 (and newer) supports this deployment scenario and gives you the ability to use SQL Server Reporting Services and the SQL Server database engine installed on separate hosts.

For step-by-step instructions on how to install and configure Parallels RAS Reporting Service with SQL Server 2019 and Microsoft SSRS 2019, please read the following Parallels KB articles:

- **Microsoft SQL Server 2017 and 2019 single server installation:**
<https://kb.parallels.com/125164>.
- **Microsoft SQL Server 2017 and 2019 multi-server installation:**
<https://kb.parallels.com/125156>.

Install Parallels RAS Reporting

To install Parallels RAS Reporting:

- 1 Log in to the server where you have Microsoft SQL Server Reporting Services installed. Make sure you use the account with administrative privileges (AD).

Note: As was mentioned earlier, SQL Server 2017 and newer allow you to install SQL Server database engine and SQL Server Reporting Services (SSRS) on different hosts. You need to be logged in to the server where you have SSRS installed.

- 2 Download the latest version of Parallels RAS Reporting from <https://www.parallels.com/products/ras/download/links/>.
- 3 Once downloaded, double-click the `RASReporting-xxx.msi` file to run the installation wizard.
- 4 Follow the onscreen instructions and proceed to the **Database connection** page. Specify the SQL Server database engine location:
 - **Location:** If the SQL Server database engine is installed on the local server (together with SSRS), select **Localhost**. If the SQL Server is installed on a different server, select **Remote** and then specify the server connection properties (see below).
 - **Server:** If you selected **Remote**, specify the FQDN or IP address of the server where you have SQL Server installed.
 - **Username:** Specify the username to log in to SQL Server.
 - **Password:** Specify the password.
- 5 On the same page, specify the SQL Server instance name. The default instance name is RASREPORTING. If you would like to use a different instance, you can specify it on this page. If the instance doesn't exist, you need to create it first.

- 6 Click **Next**.
- 7 On the **Viewing Reports User** page, you need to specify an Active Directory user who will be granted permissions to access the RAS reporting database. The default user is "rasreportingview" (note that the user must be created in Active Directory before it can be used here). You can specify a different Active Directory user if you wish, but you will need to change the reporting settings in the RAS Console before you can view reports (this change is described later in this chapter when the RAS reporting configuration is explained).
- 8 Click **Next** to install Parallels RAS Reporting.

Configure RAS Reporting in the RAS Console

To configure Parallels RAS Reporting:

- 1 Log in to the Parallels RAS Console.
- 2 Select the **Administration** category and click the **Reporting** tab in the right pane.
- 3 In the **Reporting** tab, select the **Enable RAS Reporting** option.
- 4 In the **Server** field, specify the FQDN or IP address of the server hosting your SQL Server instance. The value in the **Port** field is used by the service which receives data from the RAS Publishing Agent. The default port is 30008.
- 5 Specify a user login option by selecting one of the following:
 - **Prompt user for login details** — If this option is selected, the Parallels RAS Console user will be prompted to enter credentials before they can run a report.
 - **Use following credentials** — If this option is selected, the specified username and password will be used. The default (built-in) user name is RASREPORTINGVIEW. If you specified a different user when you installed RAS Reporting, specify that user credentials here.
- 6 To test the database connection, click the **Test connection** button.

Configure advanced settings

These settings are optional, so you can configure them according to your needs.

To access advanced settings:

- 1 On the **Administration > Reporting** tab page, click the **Tracking Settings** button. The **Advanced Setting** dialog opens.
- 2 In the **Session Information** section, specify the following options:
 - **Enable Tracking**. Records sessions data (affects all reports except server reports).
 - **Retain information for**. Select for how long the information should be kept in the database.

- 3 In the **Server Counters Information** section, specify the following:
 - **Enable Tracking.** If selected, server counter data is recorded (affects server reports only).
 - **Retain information for.** Select for how long the information should be kept in the database.
 - **Track CPU / Memory counter when change is more than (%).** Use these two options to set the minimum CPU and Memory resource usage required to record data.
- 4 The **Custom reports** section is used to enable custom reports in the Parallels RAS Console. Select the **Enable custom reports** option and specify a folder name where custom reports will be stored (or use the default "Custom reports" name). Note that this is a virtual folder located on the SQL Server Reporting Services side, so you need to specify just a name (not a traditional path). You will see the folder in the Parallels RAS Console in the **Reporting** category together with other (predefined) folders that contain reports.

Running Parallels RAS reports

To view Parallels RAS reports, select the **Reporting** category in the RAS Console. The report information is displayed as follows:

- The middle pane lists the available reports. See the **Predefined reports** subsection below for the complete list. The blue "folders" icon (at the top of the list) groups reports by type or displays all of them as a flat list. The "refresh" icon refreshes the report list by retrieving it from the database (this can be useful when you enable/disable the reporting functionality or when you add custom reports, which may not appear in the list automatically).
- When you initially open the **Reporting** category, the right pane contains just the **Information** tab page, which informs you whether Parallels RAS Reporting is active. If it's not, you need to make sure that it is installed and enabled.
- The blue "square" icon in front of the **Tasks** drop-down menu (upper right-hand side of the RAS Console) expands the reporting interface into full screen. The **Tasks** drop-down menu allows you to perform the following actions: **Duplicate** (duplicates a report tab page), **Full screen** (on/off), various **Close Report** options, **Delegate Permissions** (allows you to grant permissions to view reports to other RAS administrators, such as Power and Custom administrators who don't have these rights).

To run a report, double-click it in the middle pane. The report opens in a tab page in the right pane:

- Most of the predefined reports include controls that you can interact with, such as **From/To** dates, **Sort By**, **Sort Order**, **Chart Type**, **Server Name**, and others depending on the report type. When you change a value in any of these controls, click the **View Report** button to apply the new values/options and re-run the report.
- The main report area (lower portion where the data is represented as a graph, text, or numbers) includes a menu bar with icons that allow you to change the magnification, list through report pages (if more than one is included), search for text, save a report to a file, print a report, and export it to one of the available formats (Word, Excel, PowerPoint, PDF, or a data feed).

Note: The first time the reports are viewed, you may be requested to add https://<server domain/ IP> as a trusted website. This will appear depending on the Parallels RAS machine's Internet Explorer Enhanced Security Configuration.

Predefined reports

Parallels RAS Reporting includes a number of predefined reports in the following groups:

1 Users reports.

This group includes reports about how end users are interacting with Parallels RAS:

- **Sessions activity for all users** — shows all sessions produced by all users in the system. The report shows information about each session and includes active time, idle time, disconnected and total time. A user is identified by username and IP address. The Gateway information is also included.
- **Sessions activity for user** — shows all sessions produced by a single user. The report shows information about each session and includes active time, idle time, disconnected and total time.
- **Application usage for user** — shows applications used by a specified user, including number of times used and total time.
- **Device usage for user** — shows information about devices used by a user. The report includes information such as device vendor, device model, and total time used.
- **Operating system usage for user** — shows the operating system being used by a specified user.
- **Full user information** — shows detailed information about a specified user.

2 User groups reports.

These reports obtain information about how groups of users are interacting with Parallels RAS:

- **Sessions activity for all groups** — shows all sessions produced by all groups in the system. The report includes active, idle, and disconnected time.
- **Sessions activity for group** — shows all sessions produced by a group in the system. The report shows information about each session produced by each user in the group and includes start, end, active, idle, disconnect and total time.
- **Application usage for group** — shows applications used by a specified group, including number of times used and total time.
- **Device usage for group** — shows information about devices used by users as members of a specified group. The report includes device vendor, model and total time used.
- **Client operating system usage for group** — shows the operating system used by members of a particular group.

3 Devices reports.

This group includes reports about the devices that are connecting to Parallels RAS:

- **Devices used** — shows all devices using the system. The report includes a device manufacturer, model, and the number of sessions opened by the device.

- **Client operating system used** — shows devices and corresponding operating systems that are using the system.
- **Parallels client version used** — shows information about a device model, Parallels Client version used, and session information.

4 Servers reports. This group includes reports about the activity of Parallels RAS server components:

- **Sessions activity for RD session hosts** — shows the session activity of users on a particular RD Session Host. Report includes start, end, active, idle and disconnect time.
- **Sessions activity for VDI provider** — shows the session activity of users on a particular VDI provider. Report includes start, end, active, idle and disconnect time (standalone Hyper-V and VMware ESXi only).
- **Sessions activity for AVD provider** — shows the session activity of users on a particular AVD provider. Report includes start, end, active, idle and disconnect time.
- **RD Session hosts health** — shows server CPU and RAM usage for a specified server in the Farm.
- **VDI providers health** — shows server CPU and RAM usage for a specified provider in the Farm.
- **Remote PCs health** — shows server CPU and RAM usage for a specified Remote PC in the Farm.
- **Gateways health** — shows server CPU and RAM usage for a specified Gateway in the Farm.
- **Publishing agents health** — shows server CPU and RAM usage for a specified Publishing Agent in the Farm.
- **Enrollment servers health** — shows server CPU and RAM usage for a specified Enrollment server in the Farm.
- **Gateway tunnelled sessions** — shows tunnelled session information for a specified Gateway.

5 Application reports. Reports related to applications.

- **Activity for all applications** — shows information about applications used in the system. Report includes information such as application name, number of times used and the total usage time. When viewing this report, select "All applications" or "RAS published applications" depending on your needs. When the second option is selected, the report will not include non-published applications and duplicates.
- **Activity for application** — shows usage of an application by individual users during a specified time period. The information includes start time, end time, and total time for each session. Other information, such as host server name and session ID is also shown.

GDPR compliance

The Parallels RAS reporting database contains information about users, which may possibly include personal user information. To conform to GDPR, Parallels RAS gives you the ability to clear user data from the database at any time. **Parallels RAS Reporting Tools** is a simple application that you can use to perform this task. The tool is installed automatically when you install Parallels RAS.

To clear user data:

- 1** On the computer where you have Parallels RAS installed, navigate to `C:\Program Files (x86)\Parallels\RAS Reporting`.
- 2** In the folder specified above, locate and run the **RASReportingTools** application.
- 3** When the application starts, enter a user name in the **User data** field and click **Find user**. If the user is found, the user information is displayed. If the user is not found, it means that the RAS reporting database doesn't have any information about that user.
- 4** To see the user information contained in the RAS reporting database, click the **Show full user information** button. This will open the **Full User Information** report in a web browser (note that this report is also available in the **Reporting** category in the RAS Console). Examine the report to determine if any of the user information is subject to GDPR requirements.
- 5** To clear the user data, go back to the **Parallels RAS Reporting Tool** app and click the **Clear user data** button. When asked, confirm that you want to clear the data.

Performance Monitor

In This Chapter

Overview	422
Install RAS Performance Monitor	423
Using Parallels RAS Performance Monitor	423
Configure RAS Performance Monitor Security	426
Updating RAS Performance Monitor	427

Overview

Parallels RAS Performance Monitor is a browser-based dashboard designed to help administrators analyze Parallels RAS deployment bottlenecks and resource usage. The dashboard provides a visual display of performance metrics, which can be viewed in the Parallels RAS Console or in a web browser.

Components

Parallels RAS Performance Monitor consists of the following components:

- **InfluxDB database** — a database for storage of system performance data.
- **Grafana dashboard** — a browser-based dashboard providing a visual display of performance metrics.
- **Telegraf service** — a service that collects performance data on a server where it is installed. The service is installed automatically when you add a server to a Parallels RAS Farm and install a corresponding RAS Agent on it (e.g. RAS Secure Client Gateway Agent, RD Session Host Agent, Remote PC Agent, etc.).

How it works

The Telegraf service is stopped by default, so it doesn't collect any data. To start the service on each server in the Farm, the performance monitoring functionality must be configured and enabled in the Parallels RAS Console. Once enabled, the Telegraf service begins collecting a predefined set of performance counters at a fixed time interval (10 seconds). It then sends the collected data to the InfluxDB database for storage. To view performance metrics, the Parallels RAS administrator uses the dashboard (Grafana), which displays the visual representation of performance counters in real time.

The performance metrics are grouped in the dashboard by type (Session, CPU, Memory, Disk, etc.), so the administrator can view each group of metrics separately. The administrator can also select whether to view performance metrics for one or more specific servers or for all servers in the Farm or Site. In addition, the administrator can select a specific Site for which the data should be displayed.

Install RAS Performance Monitor

Requirements

Parallels RAS Performance Monitor is a separate component of Parallels RAS with its own installer. It can be installed on a dedicated server or on a server hosting any of the Parallels RAS components. When you run the installer, the InfluxDB database and the Grafana dashboard service are automatically installed. For additional info, see the **Installation** subsection below.

The following firewall rules (open ports) are automatically added on the server where you install Parallels RAS Performance Monitor:

- TCP port 8086 (used by the InfluxDB database).
- TCP port 3000 (used by the Grafana performance dashboard).

Installation

To install Parallels RAS Performance Monitor:

- 1 Download the Parallels RAS Performance Monitor installer from <https://www.parallels.com/products/ras/download/links/>.
- 2 Run the installation wizard (the RASPerformanceMonitor.msi file) and follow the onscreen instructions.
- 3 Close the wizard when finished.

Using Parallels RAS Performance Monitor

Configure access to Parallels RAS Performance Monitor

To enable data collection and view the dashboard:

- 1 In the RAS Console, navigate to **Administration > Reporting**.
- 2 Select the **Enable RAS Performance Monitor** option (the **RAS Performance Monitor configuration** section).

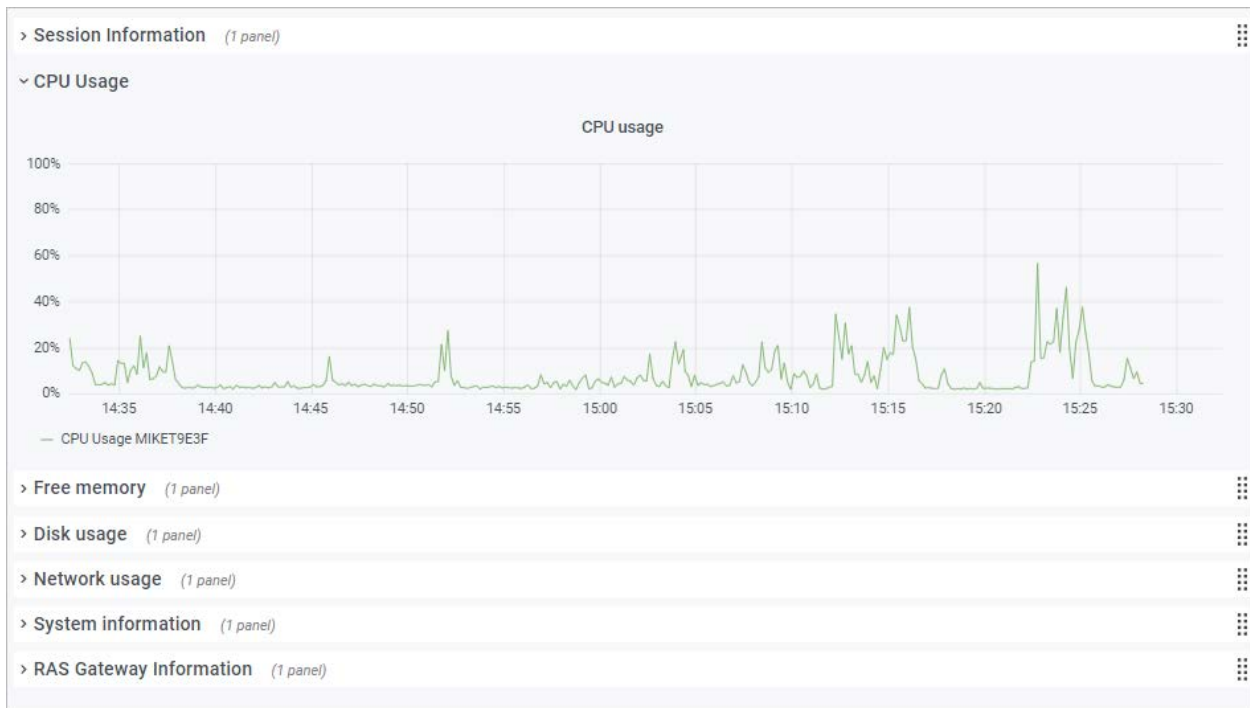
- 3 Enter the FQDN or IP address of the server where you have the InfluxDB database and Grafana dashboard installed.
- 4 Click **Apply** to commit the changes.

Once you perform the steps above, the Telegraf service is started on each server in the Site and the data collection begins.

Viewing performance metrics

Note: You should give Parallels RAS Performance Monitor some time to collect performance data before you can view it (about 1 hour on initial installation).

To view performance metrics, select the **Monitoring** category in the RAS Console. The data is displayed in the right pane. The logon is performed automatically, so no logon credentials are required.



The buttons on the **Monitoring** tab (below the dashboard area) are as follows:

- **Home.** Displays the home page. The button is useful when you click on an external link in the dashboard, which may take you to an external web page.
- **Refresh.** Reloads the current page.
- **Open in browser.** Opens the performance dashboard in a web browser.


To view metrics of a specific type, expand the desired category in the main area of the dashboard. The categories include:

- **Session Information.** Displays the information about active sessions and disconnected sessions.
- **CPU usage.** CPU counters.
- **Free memory.** Physical memory counters.
- **Disk usage.** Disk I/O counters.
- **Network usage.** Network interface I/O counters.
- **System information.** System information counters.
- **RAS Gateway information.** RAS Gateway counters.

Performance metrics are displayed in the dashboard as a graph. Different counters are displayed using different colors.

To zoom in on a particular area of a graph, select a rectangular block with a mouse. You can also use the **Zoom** controls at the top of the dashboard for time range zoom out, shift time forward, or shift time backwards. To select a specific time range, click the "clock" icon at the top and then specify a time range.

By default, the dashboard opens in kiosk mode. To exist it, press "Esc". To cycle view mode, click the "monitor" icon in the upper right. When you exist kiosk mode, the **RAS Infrastructure Health** page is displayed:



The menu at the top has the following items:

- **Hosts.** Allows you to select one or multiple servers for which the performance metrics should be displayed. To display the data for all servers in the Site, select **All**. Please note that if you don't see any servers in the list, you need to wait for Parallels RAS Performance Monitor to collect the initial set of statistics. This only happens on initial installation.
- **Instance.** This item allows you to select a specific counter instance (if there's more than one). For Network counters it is usually the name of a network interface. For Disk counters it is a disk name. Other types of counters don't usually have multiple instances.
- **Farm & Site.** Select a Site for which to display the data. Selecting **All** displays the data for all sites in the Farm. If you have another RAS Farm, and the RAS Performance Monitor is configured and enabled in it, you can also select a Site from that Farm.
- **Agent Type.** Select a RAS agent type.
- **Group.** Select an RDS group.

For more information about performance metrics and their meaning, please refer to the following articles from Microsoft:

- <https://technet.microsoft.com/en-us/library/cc976785.aspx>
- <https://technet.microsoft.com/en-us/library/2008.08.pulse.aspx>

See also **RAS Performance Counters** (p. 483).

Configure RAS Performance Monitor Security

By default, any user can access the Performance Monitor page and view performance metrics. To increase security, you can set up the RAS Performance Monitor to use credentials so that only authorized users can view it.

First, remove anonymous authentication from the Grafana configuration file as follows:

1 Open file C:\Program Files\Parallels\RAS Performance Monitor\conf\defaults.ini.

2 In the file, look for the following:

```
##### Anonymous Auth
#####

[auth.anonymous]
# enable anonymous access
enabled = true
```

3 Change "enabled = true" to "enabled = false".

Note: The user will be prompted to change the admin password automatically after disabling the anonymous access. After that, the password can be changed following the Grafana official documentation: <https://grafana.com/docs/grafana/latest/manage-users/user-admin/change-your-password/>.

4 Restart the Grafana service.

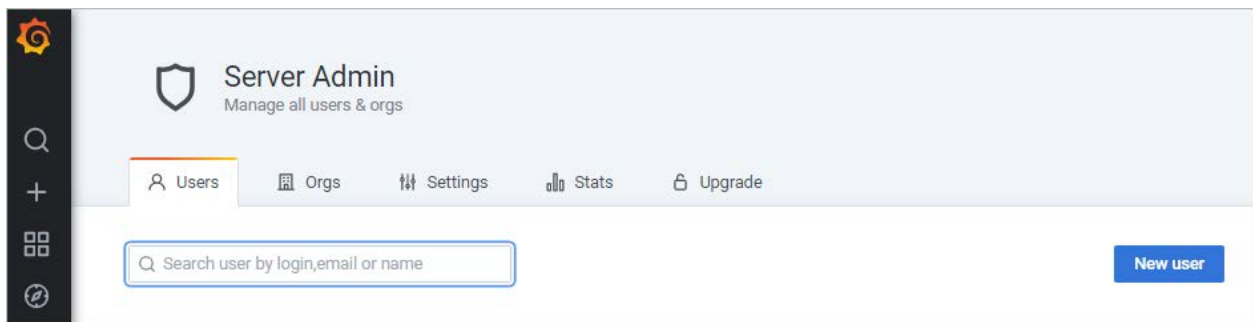
5 Select the **Monitoring** category and log in to Grafana using the following credentials:

- **User:** admin
- **Password:** admin (if you changed the password before, use the current password).

- 6 Once logged in, press "Esc" and then click the "shield" icon > **Users**.



- 7 Click **New user** and create a new user.



- 8 You now need to add the user to your organization's list. To do so, in the **Users** list, click **Edit** to edit the user and then set the organization and make the user a **Viewer**.
- 9 Click **Add** to add the user to your organization's list. The user can now view the RAS Performance Monitor statistics.

Updating RAS Performance Monitor

Newer versions of Grafana may be incompatible with a browser installed in the OS. A modern Edge browser is supported and will be embedded automatically in the RAS Console. But if it's not so, we suggest using a redistributable browser that comes in a form of a plugin. Plugins are available as separate msi packaged and distributed the same way as all other Parallels RAS distributions.

To avoid potential issues, Parallels RAS monitors which versions of Grafana and browser are installed and informs the RAS administrator about it when necessary. In a situation like that, when you select the **Monitoring** category, you will not see a dashboard as usual. Instead, a message is displayed that an updated browser plugin must be installed. To install the plugin, follow the instructions below.

- 1** In the main menu (at the top), click **Tools** and choose **Plugins**.
- 2** In the **Plugins** dialog, locate the **Browser Engine** plugin. The **State** column should indicate that it is not installed.
- 3** Select the plugin and click one of the following buttons:
 - **Install online** — Downloads the plugin from the Internet and installs it. Click this button if you can access the Internet from your environment.
 - **Install offline** — In an offline environment, admin can't get the list of plugins and instead uploads the plugin zip from a local file server. Only installed plugins are listed in the plugins table.
- 4** After the plugin is installed, you can go back to the Monitoring category and this time you will see the Grafana dashboard.

To update or remove a plugin (if necessary), use the same instructions as above, but select a plugin that is marked "installed" and then click **Update online** (or offline) or **Uninstall**. Please note that custom dashboards are now maintained after upgrades from Parallels RAS 18 to future versions.

Common Management Tasks

This chapter describes common Parallels RAS management tasks, including Farm status monitoring, license management, backup management, and others.

In This Chapter

Recovery - add a root administrator	429
Host name resolution	430
Computer management tools	431
Site information	433
Site settings	433
Settings audit	435
Upgrading RAS agents.....	437
Licensing.....	438
Configure HTTP proxy settings	440
System event notifications	440
RAS session variables	446
Maintenance and backup	447
Problem reporting and troubleshooting.....	450
Logging.....	451
Suggest a feature	453

Recovery - add a root administrator

This topic addresses a possible issue when the root administrator is not available or the domain is changed. In such events, the system becomes inaccessible. If you encounter this issue, you can quickly add a root administrator by executing the following command on the server hosting the primary RAS Publishing Agent:

```
2XRedundancy -c -AddRootAccount user [domain]
```

Please note that an open Parallels RAS console will not be notified about the new account since this is an emergency recovery. You need to log out and then log in again to see the new account in the **Administration** area.

Host name resolution

When you add a server component (Publishing Agent, Gateway, RD Session Host, VDI provider, etc.) to a RAS Farm you have to specify its FQDN or IP address. It is normally up to you whether to use FQDN or IP address. On the other hand, the server IP address can change in the future. If that happens, you will have to reconfigure the corresponding component in the RAS Farm. On the other hand, the server FQDN usually stays the same, so if you used it instead of the IP address, no RAS configuration changes will be necessary. For this reason, Parallels RAS gives you an option to always resolve IP addresses to FQDNs for all server components in a Farm.

To always use name resolution, do the following:

- 1 In the RAS Console, click **Tools > Options** on the main menu (that's the menu at the top of the RAS Console window).
- 2 In the **Options** dialog, select the **Always attempt to resolve to fully qualified domain name (FQDN) when adding hosts** option.
- 3 Click **OK**.

When you now try to add a component to a Farm and enter its IP address instead of a name, it will be automatically resolved to FQDN. If the FQDN cannot be determined, you will see an error message and will be asked if you would like to use the IP address instead.

The examples below demonstrate how the automatic name resolution works for different components.

Adding a RAS Publishing Agent

- 1 On the **Publishing Agents** tab, click **Tasks > Add**.
- 2 In the **Server** field, enter the server IP address.
- 3 Click **Next**.
- 4 In the dialog that opens, observe that the IP address has been resolved to FQDN and the **Server** field contains the FQDN.

Adding a RAS Secure Client Gateway

- 1 On the **Gateways** tab, click **Tasks > Add**.
- 2 In the **Server** field, enter the server IP address.
- 3 Click **Resolve**. This will copy the IP address to the **IP(s)** field and will enable the **Next** button.
- 4 Click **Next**.
- 5 In the **Installing RAS Secure Client Gateway** dialog, observe that the server IP address is replaced with the FQDN.

Adding an RD Session Host

- 1 On the **RD Session Hosts** tab, click **Tasks > Add**.
- 2 On the first page of the wizard, enter the server IP address and click the plus-sign icon.
- 3 Observe that the server is added to the list, but the IP address is substituted with the FQDN that was automatically resolved.

Add a VDI provider

- 1 On the **VDI > Providers** tab, click **Tasks > Add**.
- 2 Select **Virtualization** and click **Next**.
- 3 In the **Address** field, enter the IP address of a VDI provider.
- 4 Enter the remaining properties and click **Next**.
- 5 Observe that the VDI provider address is replaced with the FQDN.

Computer management tools

When you need to perform standard Windows computer management tasks, you can do it without leaving the RAS Console. The tasks include Remote Desktop Connection, Computer Management, Service Management, Event Viewer, PowerShell, Reboot, and others. To perform these tasks, use the **Tools** menu, which is accessible from the **Site** menu and individual Parallels RAS infrastructure servers and session hosts.

Requirements for using computer management tools

Some of the tools require an appropriate target host configuration before you can use them in the RAS Console. Please read the following requirements and make sure they are met.

To use Remote Desktop, remote connections must be enabled on a target host. You can verify that by using the standard Windows Remote Desktop Connection application and see if you can connect to a remote server.

PowerShell related tools require PowerShell remoting enabled on a target server. To enable PowerShell remoting, run the `Enable-PSRemoting` cmdlet on a target computer in PowerShell window with administrator privileges. Please note the following:

- The cmdlet configures a computer to receive PowerShell remote commands.
- The cmdlet starts the WinRM (Windows Remote Management) service, among other tasks. To see if the WinRM service is running, use the `Test-WSMan` cmdlet.
- When you execute the cmdlet, it will ask you to confirm every task that it wants to perform. To execute the command silently, use the `-Force` option.

Common Management Tasks

- If you receive an error saying that "WinRM firewall exception will not work since one of the network connection types on this machine is set to Public", you can try to execute the cmdlet with the `-SkipNetworkProfileCheck` option, or you can change the network connection type on this host to Domain or Private.

To use PowerShell to manage a remote host, you also need to add the host to the TrustedHosts list on the computer where you have the RAS Console installed. To view the current TrustedHosts list, execute the following command in PowerShell window:

```
Get-Item WSMan:\localhost\Client\TrustedHosts
```

To add a host to the TrustedHosts list, use one of the options described below. Please note that all examples below, except the last one, always overwrite an existing TrustedHosts list. To add a specific computer to an existing list, use the last example (the one with the `-Concatenate` parameter).

Add all computers to the list:

```
Set-Item WSMan:\localhost\Client\TrustedHosts *
```

Add all domain computers:

```
Set-Item WSMan:\localhost\Client\TrustedHosts *.domain-name.dom
```

Add specific computers:

```
Set-Item WSMan:\localhost\Client\TrustedHosts <computer-name> , [ <computer-name> ]
```

Add a computer to an existing list (this is the only example that will not overwrite an existing TrustedHosts list):

```
Set-Item WSMan:\localhost\Client\TrustedHosts -Concatenate <ComputerName>
```

Available tools

The table below describes the tools available in the **Tasks > Tools** menu and their execution strings.

Tool	Execution string	Description
Remote Desktop	<code>mstsc.exe /v:<selectedRDShostName>:<port> /admin</code>	Launch a standard RDP connection to the selected RDS host.
Computer Management	<code>compmgmt.msc /computer:<selectedRDShostName></code>	Launch Computer Management locally with connection to the selected host.
Service Management	<code>services.msc /computer:<selectedRDShostName></code>	Launch Services Management locally with connection to the selected host.
Event Viewer	<code>eventvwr.msc /computer:<selectedRDShostName></code>	Launch Event Viewer locally with connection to the selected host.

Shared Folders	smgmt.msc /computer:<selectedRDShostName>	Launch Shared Folders locally with connection to the selected host.
Powershell	Enter-PSSession -ComputerName <selectedRDShostName> [-Credential username]	Launch Powershell locally with connection to the selected host.
IPconfig	- Powershell remote connection to selected host - Get-NetIPConfiguration	Provides network configuration for the selected host.
Ping	- Powershell remote connection to selected host - Test-NetConnection -ComputerName www.microsoft.com Select -ExpandProperty PingReplyDetails FT Address, Status, RoundTripTime	Provides ICMP reply with status and RTT for the selected host.
Netstat	- Powershell remote connection to selected host - Get-NetTCPConnection	Displays network connections for Transmission Control Protocol on the selected host.
Reboot	shutdown /m \\<selectedRDShostName> /f /r /t 0	Reboot the selected host.
Shutdown	shutdown /m \\<selectedRDShostName> /f /s /t 0	Shutdown the selected host.

Note that individual tools availability depends on the server type. For example, HALB has only **Ping** in the **Tools** menu.

Site information

To view the Site information, select the **Information** category in the RAS Console.

The **Site Information** tab displays information about available servers, Publishing Agents, Secure Client Gateways (see **Viewing Gateway summary and metrics** (p. 81)), and sessions on the local computer. To view information about running applications, select the **Show application information** option (at the bottom of the page).

The **Local Information** tab shows the status of RAS components running on the local server.

Site settings

To view and configure common Site settings in the RAS Console, navigate to **Farm > <Site> > Settings**.

Auditing

The **Auditing** tab allows you to configure application auditing. When enabled, application auditing monitors processes running in the Site and records this information in the audit file. To view the information, click the **View Audit** button (at the bottom of the page). The information is also displayed on the **Information > Site** page and in RAS Reports.

To enable or disable application auditing, use the **Auditing** drop-down list (at the bottom of the page). The **Clear Audit File** button clears the current audit.

The **Filtering the following processes** list allows you to specify processes that will be excluded from the audit. Use the **Tasks** drop-down menu to add or delete a process. You can also use the **Task** menu to import and export a process list from/to a CSV file. The **Task > Properties** menu item allows you to edit a process name. The **Default** menu item resets the list to contain the default set of standard processes.

Global logging

The **Global logging** tab allows you to specify the log level for Parallels RAS components. Logs are used by Parallels RAS support engineers to analyze possible issues with a Parallels RAS installation. To specify the log level, select one or more servers in the list and click the **Configure Logging** item. In the dialog that opens, select one of the following:

- **Standard** — This is the standard log level that records only the most important events. Unless you are asked by Parallels RAS support to use one of the log levels described below, you should always use this one.
- **Extended** — This logging involves more information than the standard logging, but it slows down the system because of the additional information that it needs to collect.
- **Verbose** — Verbose logging involves even more information than the extended logging and can slow down your system significantly.

Please note that to avoid degraded performance, extended and verbose logging should only be enabled for a limited time period (enough to collect the necessary information for analysis). You can set this time period using **Reset to the standard level after** option. The default value is 12 hours. In specific cases, a Parallels support engineer will advise you whether this time period should be set to a different value. Once this time period is over, the log level will be reset back to standard.

To retrieve a ZIP archive containing the collected log files, click the **Retrieve** item and then specify a location where you want the file to be saved. The **Clear** item clears all logs.

You can also set the log level for an individual server by navigating to the page where servers of that type are listed (e.g. RD Session hosts, Gateways, etc) and clicking **Tasks** (or right-click) > **Troubleshooting > Logging**. The context menu that opens has the same **Configure**, **Retrieve**, and **Clear** options as described above. The **Log Level** column in the server list indicates the currently set level.

URL redirection

The **URL redirection** tab allows you to specify URLs that will not be redirected when the **Allow Client URL/Mail redirection** option is enabled for an RD Session Host, VDI provider, or Remote PC (**Agent Settings** tab in the corresponding server properties).

To add one or more URLs that should not be redirected, click **Tasks > Add** and type a URL in the **Do not redirect the following URLs** list box.

Client settings

See **Specifying client settings** (p. 254)

Settings audit

Parallels RAS gives you the ability to audit the modifications that were done to a Parallels RAS Farm, including changes to any of the components, objects, resources, and users. This information is stored in a database, so it can be reviewed and possibly reverted, if needed. The information is stored in the primary database but is replicated in a local database on the computer where Parallels RAS Console is running.

You can view the list of modifications using one of the following options:

- By navigating to **Administration > Settings audit**. The tab displays the main list of all changes to any components/objects in the Farm. If a modification can be reverted, you can do it here.
- By clicking **Tasks > Settings audit** on any pane in the RAS Console that supports this functionality. Compared to the main list (described above), you will only see modifications to the same types of components or objects that are managed on a given pane. You can also revert a modification here if it can be reverted. If the **Settings audit** menu option is not available on a particular pane, it means that the functionality is not available for the types of components or objects that this pane manages.

The following describes in detail how to view and revert Farm modifications.

View the main settings audit list

To view the main list of all modifications for a Farm, do the following:

- 1** In the Parallels RAS Console, select the **Administration** category and then click the **Settings audit** tab.
- 2** The sync process will check that the local audit database is in sync with the primary database and will do an update if necessary (you may see a progress indicator while the syncing is in progress).

- 3 Once the syncing is complete, the **Settings audit** tab will be populated with data. Each entry in the list corresponds to a modification that was done either by a RAS administrator or a system service.

The information for each entry in the list includes the following:

- **Date:** Date and time of the modification.
- **Session:** Session ID.
- **Username:** The name of the administrator or RAS service that was responsible for the modification. RAS services may include System (redundancy service) and Publishing Agent (controller service).
- **Action:** The action that was performed, such as Connect, Disconnect, Create, Update, Switch site, and others.
- **ID:** The affected object's ID.
- **Site:** The number and name of the affected Site. "Global" means the change affected all sites.
- **Type:** The modification type. This usually makes sense when viewed together with the **Action** value.
- **Name:** The value in this column is displayed for some entries and can provide additional information, such as the name of the changed object.

You can perform the following actions on the list:

- To refresh the list, click the "recycle" icon (top right).
- To view details for an entry, double-click it (or select an entry and click **Tasks > View entry**).
- To search for a specific entry (or entries), click the magnifying glass icon (top right). An extra row is added at the top of the list allowing you to enter the search criteria. You can type a string to search for in one or multiple columns. The search is performed as you type and the list is filtered to include only the matching entries. To cancel filtering and display the complete list, click the magnifying glass icon again.

Reverting a modification

To revert a modification in the main list:

- 1 Double-click a desired entry on the **Settings Audit** tab.
- 2 The **Audit Entry** dialog opens. While here, you can click **Next** and **Previous** buttons to go to the next or previous item as they are displayed in the main list.
- 3 To revert the change, click the **Revert** button. If the button is disabled, it means that the change cannot be reverted.

Changes that can never be reverted include the following:

- Any changes done by System or Publishing Agent (as displayed in the **Username** column).

- Changes that were done in previous versions of Parallels RAS where this feature did not exist.
- Changes related to administrator accounts.

View a local settings audit list

You can also view and revert configuration changes for a specific type of RAS components or objects. When you are on a particular pane (or tab) in the RAS Console, look for the **Tasks > Settings audit** menu option (or right-click > **Settings audit**). If it's there, then you can view the changes and revert them if needed. Consider the example below.

Let's say you want to see changes that were done to RD Session Hosts. To do so:

- 1** In the RAS Console, navigate to **Farm > <Site> > RD Session Hosts**.
- 2** Click **Tasks > Settings audit**.
- 3** The **Settings Audit** dialog opens listing all known modifications that were done to RD Session Hosts. The modifications may include creating, moving, deleting, or updating an RD Session Host. The type of the modification is displayed in the **Action** column in the list.
- 4** To revert a modification, select it and click the **Revert** button (in the lower right of the dialog). If the button is disabled when you select a particular entry, it means that the modification cannot be reverted.

The local settings audit functionality is available for most of the major components and objects in the Parallels RAS Console. This includes RD Session Hosts (including Groups and Scheduler), VDI, Remote PCs, Gateways, Publishing Agents, Themes, Publishing, Quick Keypad, and many others. Once again, when you view a particular pane, look for the **Tasks > Settings audit** menu option (or right-click > **Settings audit**). If it's there, then you can view the changes and revert them if needed.

Upgrading RAS agents

When you add Parallels RAS components to a Farm, you install a corresponding RAS Agent on them. This includes RAS Publishing Agent, RD Session Host Agent, VDI Agent, Guest Agent, Remote PC Agent. In addition to the functionality that allows you to check agent status, and update it if necessary, you can do a bulk agent update or upgrade.

There are two ways you can find out if agents need to be updated. You can be notified by Parallels RAS or you can check the status and initiate the update procedure manually.

When the Parallels RAS Console starts, you may see a message box saying that Agents need to be installed or updated. You can start the update procedure by clicking **Yes** in this dialog. You will then see a list of all servers on which an Agent needs to be updated where you can decide whether to include a server in the bulk update procedure or exclude it. Once you've made your selection, follow the onscreen instructions and update the Agents.

To initiate the procedure manually, click the **Task > Upgrade all Agents** in the RAS Console where this menu is available (most of the views where it makes sense). You can also right-click inside the view and choose **Upgrade all Agents**. Follow the onscreen instructions and select the servers on which an Agent requires an update or upgrade. Please note that if all Agents on all servers displayed on a given pane are up to date, the menu option will be disabled.

For example, to upgrade all primary Publishing Agents in all sites, select **Farm > Farm** and then click **Tasks > Upgrade all Agents** (or right-click inside the pane and choose **Upgrade all Agents**). To upgrade all Agents on all servers in a Site, select **Farm > <Site>** and click **Tasks > Upgrade all Agents**. Similarly, to upgrade Agents on all RAS Secure Client Gateways, select **Farm > <Site> > Gateways** and use the same **Tasks > Upgrade all Agents** menu item. For other components, do exactly the same. Note that if you use the same credentials on all servers, you will have to enter them only once. The update procedure will remember the last entered credentials and will try to use them on all servers. If the credentials don't work on a server, you'll be asked to enter them again.

Please note that after you click the **Tasks > Upgrade all Agents** menu, the dialog that opens will contain the hosts on which an Agent needs updating or upgrading. The **Status** column in the list will indicate that and the host will be preselected for the upgrade. Unverified Agents will also be included in the list but will not be preselected. You can select them if you believe that an Agent has to be upgraded on them too.

Note: When updating an agent in a template (VDI), full and linked clone templates are updated differently. For important information, please read the **Template maintenance** section (p. 163).

Licensing

The **Licensing** category allows you to manage your Parallels RAS license. When you click on the **Licensing** category, the **License details** tab displays the license information. Note that depending on the license type (prepaid subscription, SPLA, NFR, trial), different information is displayed. Please also note that an NFR (Not for resale) license can be either a prepaid subscription or SPLA, so different information can also be shown for an NFR license.

The information that can be displayed on the **License details** tab includes the following:

- **License Type:** The type of your Parallels RAS license (e.g. prepaid subscription, SPLA, NFR, trial).
- **License Key:** The license key that was used to activate the Farm (only last characters are shown).
- **Support Expiration Date:** The support plan expiration date.
- **Upgrade Insurance:** The upgrade insurance expiration date which, in case of using subscription based licenses, is in line with Expiration date, since such licenses are automatically eligible for upgrades.
- **Expiration date:** The license expiration date (or the number of days remaining for a trial license).

- **First Activation:** The date the Farm was first activated.
- **Peak Users:** Peak concurrent users to date. For a prepaid subscription, you can use this value to evaluate whether you might need to upgrade your subscription to include more concurrent users.
- **Usage Today:** SPLA license only. The number of maximum concurrent users recorded for today. Note that "today" starts at midnight UTC.
- **Current Period Usage:** SPLA license only. The combined usage for all Farms activated with the same license key (SPLA licenses allow the activation of multiple farms with the same key).
- **Billing period started:** SPLA license only. The current billing period start date.
- **Billing period ends:** SPLA license only. The current billing period end date.
- **Current Users:** The number of users currently connected to your Parallels RAS Farm.
- **Maximum allowed concurrent users:** Prepaid subscription and NFR licenses only. The maximum number of concurrent users that your license allows. For example, if you own a prepaid subscription and need more concurrent connections, you need to upgrade your subscription.
- **Parallels Account user email:** Parallels business account email that was used to activate the Farm.
- **Parallels Account user name:** Parallels business account user name.
- **Parallels Account company:** Parallels business account company name.

Please note that you can also see these values (and more) in your Parallels Account. For more information, please read the **Parallels RAS Licensing Guide** and **Parallels RAS SPLA Guide**, which are available on the Parallels website.

The **View Active Users** button opens a dialog where you can view currently active users and license usage. Use the toolbar buttons to refresh the list and to copy the information to the clipboard.

The **Manage license** button allows you to switch to a different Parallels account and to activate Parallels RAS using a different license key. Click the button to open the **Sign in to Parallels My Account** dialog. Use the dialog to sign in using an existing account or click **Register** to create a new account. If you are creating a new account, you'll also have to register a Parallels RAS license key in it and activate your Parallels RAS Farm using that key (see below).

To activate Parallels RAS using a different license key:

- 1 In the **Sign In to Parallels My Account** dialog, type the email address and password you used to register your account and click **Sign In**. You'll see the **Activate Product** dialog.
- 2 Select the **Activate using license key** option and enter the key in the field provided. You can click the button next to the field to see the list of subscriptions and/or permanent license keys you have registered with Parallels My Account. If the list is empty, it means that you don't have a subscription yet and need to purchase one first.

- 3 To purchase a subscription online, click the **Purchase a license** link.
- 4 After entering a license key, click **Activate**. You should see the confirmation message that your Parallels RAS was activated successfully.

Configure HTTP proxy settings

If you use an HTTP proxy server on your network, you need to configure it in the RAS Console. The proxy server settings will be used during Parallels RAS license updates and by other features that communicate with the Parallels cloud.

To configure a proxy server:

- 1 In the RAS Console, navigate to **Administration > Settings**.
- 2 In the **HTTP Proxy settings** section, click the **Configure Proxy** button.
- 3 In the dialog that opens, select one of the following options:
 - **No Proxy server** — if you don't use a proxy server.
 - **Manual HTTP proxy configuration** — select this option to specify the settings manually. The **Detect Settings** button will attempt to detect the proxy settings automatically.

The **Proxy requires authentication** option allows you to specify or omit credentials for the proxy server. If your proxy server uses an IP address to authenticate clients, you can omit the credentials. Otherwise, select this option and specify a user name and password.
- 4 Click **OK** to save the settings.

System event notifications

You can configure system event notifications on the **Farm > Site > Settings > Notifications** tab. Notifications are used to alert the administrator about system events via email. When you configure notifications, the settings apply to all servers in the Farm.

To configure notifications, you first need to configure notification handlers where you can specify threshold values (where available) and whether an administrator should be notified via email. You can also configure notification scripts, which will be automatically executed when an event occurs.

Configuring notification handlers

To configure notification handlers:

- 1 In the RAS Console, navigate to **Farm > Site > Settings**.
- 2 Select the **Notifications** tab.

- 3 Click **Tasks > New** (or click the plus-sign icon) and choose an event for which to create a handler. For the list of events and their descriptions, please see the **System Events** subsection below.
- 4 A dialog opens where you can specify the event handler setting.

On the **General** tab, specify the following options:

- The threshold value (a number or percentage). Not available for some events (such as Licensing, Agent, and some other events).
- The direction (whether the event should trigger when the value rises above or drops below the specified value). Not available for some event (same as above).
- Whether to notify the administrator via email.
- Additional emails (separated by commas or semi-colons) to which to send event messages.
- Whether to execute a script when the event triggers. Here you need to select the **Execute a notification script** option and then choose a script from the drop-down list. Before you can use this option, you need to create one or more scripts as described in **Configuring notification scripts** (p. 443).

On the **Settings** tab, specify the following:

- **Use default settings:** Select this option to use default settings. To edit defaults, click the **Edit Defaults** link. To use custom settings, clear this option and specify the options as described below.
- **Notification handler grace period:** Specify a time period (in minutes) to wait from the event occurrence until the notification is triggered. Some events may trigger but last for a very short period of time. For example, a CPU usage can sharply jump above the specified threshold but quickly return to normal. For such events, it would probably make sense not to trigger the notification right away. This option allows you to specify the delay.
- **Notification interval:** Specify the minimum time interval (in minutes) between the last and the next notification. Allows to prevent multiple notifications to be emailed to administrators in rapid succession (i.e. prevents spamming).
- **Send one notification and suspend further notifications until recovered:** When this setting is enabled, a notification will be raised only once, and after that it will be suspended until the values monitored by the notification have recovered. For example, if the CPU usage is above the threshold for the whole day, instead of executing the notification handler multiple times, RAS would execute it only once.

- 5 When done, click **OK** to save the notification handler.

Please note that the mailbox should be configured in the RAS Console for the outgoing email functionality to work. This mailbox is usually set up when you run the RAS Console for the first time and then use the **Start** category to set up your RAS environment. You can also set up a mailbox as described in **Configuring SMTP server connection for event notifications** (p. 446).

To enable or disable an event handler, select or clear the checkbox in the first column, or right-click an event and choose **Enable** or **Disable**. To modify a handler, right-click it and choose **Properties**. To delete a handler right-click and choose **Delete**.

System Events

You can create notification handlers for the following system events:

- **CPU utilization.** Triggers when CPU utilization rises above or drops below a specified value.
- **Memory utilization.** Triggers when memory utilization rises above or drops below a specified value.
- **Number of RDSH sessions.** Triggers when the number of active sessions rises above or drops below a specified value.
- **Number of disconnected RDSH sessions.** Triggers when the number of disconnected sessions rises above or drops below a specified value.
- **RDSH session utilization.** Triggers when the number of RDSH sessions rises above or drops below a specified percentage of the maximum number of sessions.
- **RDSH disconnected sessions utilization.** Triggers when the number of RDSH disconnected sessions rises above or drops below a specified percentage of the maximum number of sessions.
- **Number of gateway tunneled sessions.** Triggers when the number of gateway tunneled sessions rises above or drops below a specified value.
- **Failed gateway tunneled sessions.** Triggers when a connection between a gateway and a resource object cannot be established.
- **RAS Agents events.** Triggers when an agent event occurs (e.g. agent disconnects or reconnects).
- **Licensing events.** Triggers when a licensing event occurs. One notable event here is the license usage reaching a predefined threshold. Specifically, when the license usage reaches 90% of all available licenses, you will receive an email, so you have time to decide whether you have enough licenses or need to add more. Other events include license activation/deactivation, license expiration, grace period starting/ending, license information changes, problem communicating with the licensing server.
- **Authentication server events.** Triggers when a connection issues occurs with an authentication server.
- **Published items events.** Triggers when a published item event occurs (e.g. the concurrent instance limit for an application is reached).
- **VDI events.** Triggers when a VDI event occurs (e.g. a template is not found).
- **Tenant events.** Triggers when a Tenant event occurs. For more info, see **RAS Multi-Tenant Architecture > Configuring notifications** (p. 316).

Please also see the **Notification Types** table in the **Configuring notification scripts** section (p. 443).

Configuring notification scripts

To configure notification scripts:

- 1 On the **Notifications** tab, click **Tasks > New** (or click the plus-sign icon) in the **Notifications scripts** section.
- 2 In the dialog that open, specify the following options:
 - **Script name:** Enter a friendly name for the script.
 - **Command:** The command to execute.
 - **Arguments:** Command line arguments to pass to the command. An argument can be one of the predefined variables, which Parallels RAS will automatically replace with an actual value. See the **Command Line Variables** table below (the ID column contains the values that can be used here).
 - **Initial directory:** The full path to the current directory for the process. The string can also specify a UNC path.
 - **User name, Password:** These are optional fields that you can specify if you would like to execute the command under a specific user account.
- 3 When done, click **OK** to save the notification script item.

To modify a notification script, right-click it and choose **Properties**.

To delete a script, right-click and choose **Delete**. Please note that if a script is used by a notification handler, you will see a warning message. If you choose to delete it anyway, the script association will be removed from all notification handlers where it is used and all affected handlers will be automatically configured to send an email alert.

Command line variables

The following table lists command line variables that you can use as arguments when executing a script (see the **Arguments** option description above):

Variable	Description
(\$FARM-NAME)	Name of the RAS Farm which has raised the notification.
(\$SITE-NAME)	Name of the RAS Site which has raised the notification.
(\$SERVER-ADDRESS)	IP address or FQDN of the server which has raised the notification. It could be an RDSH server, the server hosting a RAS Publishing Agent, RAS Secure Client Gateway, etc.
(\$TRIGGER-ADDRESS)	IP address or FQDN of the Publishing Agent that has raised the notification.
(\$THRESHOLD-VALUE)	The threshold value that is assigned to the notification handler. If a notification type doesn't support thresholds, the argument should be replaced with an empty string.
(\$NOTIFICATION-TIME)	GMT time and date of when the event has occurred. String format shall use the "R" or "r" format specifier. Please see the following article from Microsoft for details:

	https://docs.microsoft.com/en-us/dotnet/standard/base-types/standard-date-and-time-format-strings Note: The time should represent the time when the notification has occurred, and not when the notification handler has been executed. The notification handler may be executed with a delay if a grace period is enabled.
(\$NOTIFICATION-TYPE)	A numeric value that is assigned to each particular notification type. Notification type values are listed in the Notification Types table below.

Notification types

The following table lists supported notification types (the ID column represents values that are passed to the (\$NOTIFICATION-TYPE) command line variable):

Notification type	ID	Description
CPU utilization	1	This notification is triggered when CPU utilization rises above or decreases below a certain value.
Memory utilization	2	This notification is triggered when memory utilization rises above or reaches below a certain value.
Number of active session	3	This notification is triggered when the number of active sessions rises above or decreases below a certain value.
Number of disconnected sessions	4	This notification is triggered when the number of disconnected sessions rises above or decreases below a certain value.
RAS Agent reconnect	5	Agent reconnected.
RAS Agent disconnect	6	Agent disconnected.
VDI template is missing	7	This notification is triggered when an a VDI event occurs (e.g. a template is not found).
Published application limit exceeded	8	This notification is triggered when a published item event occurs (e.g. an application's instance limit is exceeded).
Multi PA communication error	9	Multiple PA communication error.
Authentication provider not reachable	10	This notification is triggered when a connection issue occurs with an authentication server.
% of RDSH session out of the maximum specified value	11	This notification is triggered when the number of RDSH sessions rises above or decreases below the specified percentage of the maximum number of sessions.
Gateway is tunneling X number of sessions	12	This notification is triggered when the number of gateway tunnelled sessions rises above or decreases below a certain value.
% of RDSH disconnected session out of the maximum specified value	13	This notification is triggered when the number of RDSH disconnected sessions rises above or decreases below the specified percentage of the maximum number of sessions.
Publishing Agent auto promotion	20	Publishing Agent auto promotion
Publishing Agent auto promotion failed	21	Publishing Agent auto promotion failed.
Publishing Agent auto promotion fallback	22	Publishing Agent auto promotion fallback.
CA not available	30	This notification is triggered when a connection issue occurs Certificate Authority.
License site switched to failed over mode	50	Publishing Agent failover mode.
License site is offline	51	Licensing site is offline.

License site reconnected	52	Licensing site is back online.
IP of Licensing PA changed	53	Licensing Publishing Agent IP change.
Hostname of Licensing PA changed	54	Licensing Publishing Agent Hostname change.
IP of secondary PA changed	55	Non Licensing Publishing Agent IP change.
Hostname of secondary PA changed	56	Non Licensing Publishing Agent Hostname change.
Templates max guests reached	60	Template max guest limit reached.
Template max servers reached	61	Template maximum server limit reached
Templates cloning failed	62	Template cloning failed.
License activated	100	This notification is triggered when a licensing event occurs (e.g. a farm has been successfully activated).
License deactivated	101	License was deactivated.
License max usage	102	The maximum license usage has reached x%.
License about to expire	103	If license is about to expire, notify every day saying how many days left.
License expired	104	License expired.
License trial expired	105	Trial period expired
License grace period start	106	Grace period started.
License grace period end	107	Grace period ended.
License disabled	108	License was disabled.
License information changed	109	License information changed
License failed to communicate with server	110	Failed to communicate with licensing server.
License no file	111	Failed to load license file.
License invalid version	112	Invalid license file version.
License invalid signature	113	Invalid license signature.
License invalid license	114	System errors.
License invalid MAC address	115	Invalid MAC address (hardware change).
Licensing unsigned grace period	116	Migration grace period started.
Tenant enrolled	200	This notification is triggered when an event related to any registered tenant occurs (e.g. a new tenant was added to Tenant Broker or tenant becomes unavailable).
Tenant status changed	201	Tenant status changed
Broker status change	202	Tenant Broker status changed
Tenant disenrolled	203	A Tenant has unjoined the broker.
Standard Farm tunnel session failed	220	Standard tunnel session failed.
Broker Farm tunnel session failed	221	Tenant Broker tunnel session failed.

Configuring SMTP server connection for event notifications

The **Mailbox** tab in the **Administration** category allows you to configure an SMTP server for outgoing emails. The SMTP server is required for the administrator to receive system event alerts (as described in the previous sections) and to send invitation emails to users.

To configure an SMTP server:

- 1 In the RAS Console, select the **Administration** category and then click the **Mailbox** tab.
- 2 In the **Mail Server** field, type your mail server FDQN or IP address.
- 3 In the **TLS / SSL** drop-down list, select whether to use it the protocol.
- 4 Select the **SMTP server requires authentication** option if required and then type the SMTP server username and password in the fields provided.
- 5 In the **Sender information** section, type the sender email address (e.g. your email).
- 6 The **Test mailbox settings** section can be used to test your SMTP server configuration. Enter one or more email addresses separated by a semicolon. Click **Send Test Email** to test the settings.

RAS session variables

When a remote user starts a published application or desktop, a set of session variables is created by Parallels RAS on the host server. The variables contain information about the client machine, which you can examine if needed. The variables are always updated, so on connect/reconnect they always contain the latest values.

The following RAS session variables are available:

Variable Name	Description
TUX_REMOTECLIENT_PLATFORM	Name and version of the operating system running on the client machine. For example, "Windows 8.1 Enterprise Edition (WOW 64)", "iPhone OS 9.2.1", "Android 6.0", etc.
TUX_REMOTECLIENT_MAC	MAC address of the client machine.
TUX_REMOTECLIENT_IP	IP address of the client machine as seen by the client.
TUX_REMOTECLIENT_LANG	Language used by the GUI on the client machine: EN, FR, RU, DE, ES, IT, PT, NL, JP, CS (Chinese Simplified), CT (Chinese Traditional), KR (Korean). Note that on macOS, iOS, and Android devices, the language is reported as the one used in the OS but only if it's a supported language. If it's not supported, it will default to EN.
TUX_REMOTECLIENT_MACHINE	Client's computer name. For example, "Bob's iPad mini 1st generation", "BobPC", "Bob's iMac", etc.

TUX_REMOTECLIENT_LOGIN	The username (including domain) that was used to log in to Parallels RAS. For example, myuser@somedomain.
TUX_REMOTECLIENT_VERSION	Parallels Client version.
TUX_REMOTECLIENT_VENDOR	Device vendor name. For example, "Asus", "Apple", "Google", etc.
TUX_REMOTECLIENT_MODEL	Device model name. For example, "Nexus 5", "iPad2.6", etc.

You can view RAS session variables and their values using one of the following two methods:

- By examining the Windows registry on the host server.
- By executing the GetRASVariable.exe utility (provided by Parallels RAS).

Each method is described below.

Examine the registry

To see the variables, run `regedit` and navigate to `HKEY_CURRENT_USER\Software\Parallels\Shell\<Session-ID>`, where `<Session-ID>` is the ID of a session as displayed in the RAS Console (e.g. 2, 3, 4, etc.) The variables for a particular session are listed under the session ID node. On user connect/reconnect they are updated to reflect the actual client configuration. The variables exist for the duration of a session and are removed from the registry once the session is terminated.

Please note that in addition to the variables listed in the table above you may see other (undocumented) variables under a session ID. Those are for internal Parallels RAS use only and should be ignored.

Using GetRASVariable.exe utility

The `GetRASVariable.exe` utility is located in the Parallels RAS installation folder (e.g. `C:\Program Files (x86)\Parallels\ApplicationServer`). To obtain a value of a variable, execute the utility from the command line passing the variable name as parameter (see the table above). The utility will output the value to the screen.

The following example displays the value of the `TUX_REMOTECLIENT_MACHINE` variable:

```
GetRASVariable.exe TUX_REMOTECLIENT_MACHINE
```

Maintenance and backup

Keeping Parallels RAS up to date

By default, Parallels RAS checks for updates each time the RAS Console is started. If you wish to change this behavior:

- 1 Select the **Administration** category and click the **Settings** tab.
- 2 Select or clear the **Check for updates when launching Parallels RAS Console** option according to your needs.
- 3 To check for updates manually, click the **Check Now** button.

Backing up the Parallels RAS Farm configuration

To backup the Parallels RAS Farm configuration:

- 1 Select the **Administration** category and then click the **Settings** tab.
- 2 Click the **Export Settings** button.
- 3 You'll see a message box saying that all sites will be synchronized. Click **Yes** to continue with export or click **No** to abort it.
- 4 Specify the file name and target folder and click **Save**.

Note: The export procedure only exports the Parallels RAS Farm configuration data. Unrelated objects, such as downloaded OS, etc. are not included in the exported file.

To restore a Parallels RAS Farm configuration from a backup file, click the **Import Settings** button and select a backup file (the default file extension is `.dat2`). When you import a configuration from a file, your existing Farm configuration will be completely replaced with it.

You can also export and import a Parallels RAS Farm configuration from the command line. For complete instructions, please read on.

Exporting and importing Farm settings from the command line

Parallels RAS PowerShell allows you to perform the majority of Parallels RAS administration tasks from the command line.

This section contains information about using PowerShell to export and import Farm settings. To learn more about Parallels RAS PowerShell, please visit <https://www.parallels.com/products/ras/resources/> and download (or view online) the **Parallels RAS PowerShell Guide**.

One of the uses of exporting and importing Farm settings is running automated tests. Specific configurations can be created, exported, and then imported for specific test scenarios. You can also use this functionality with Windows task scheduler for regular backups of Farm settings.

Installing Parallels RAS PowerShell

RAS PowerShell is installed by default when you run the default Parallels RAS installation. If you haven't installed it (or to install it on a different computer), do the following:

- 1 Run the Parallels RAS installer.
- 2 Select **Custom** and then select the **Parallels RAS PowerShell** component.
- 3 Complete the wizard and install Parallels RAS PowerShell.

Using Parallels RAS PowerShell

The complete up-to-date information about Parallels RAS PowerShell can be found in the **Parallels RAS PowerShell Guide**. The guide includes the **Getting Started** chapter to help you quickly get started with Parallels RAS PowerShell, as well as the complete reference and code samples. Please visit <https://www.parallels.com/products/ras/resources/> to view or download the guide.

Use the instructions below to export and import Parallels RAS Farm settings.

To import the Parallels RAS PowerShell module, open the PowerShell console and execute the following command:

```
Import-Module PSAdmin
```

Create a Parallels RAS session (use the name or IP address of the server where you have Parallels RAS installed):

```
New-RASSession -Server "server.company.dom"
```

To export Farm settings, execute the following command (substitute the path and filename of the backup file with your own):

```
Invoke-ExportSettings "C:\Backup\RAS-backup.dat2"
```

To import Farm settings:

```
Invoke-ImportSettings "C:\Backup\RAS-backup.dat2"
```

Problem reporting and troubleshooting

If you are experiencing an issue with Parallels RAS, you can search for a solution right from the RAS Console. If you can't find a solution, you can send a support request to Parallels. This section describes how to accomplish these tasks.

Search for a solution

To search for a solution from the RAS Console:

- 1 In the console, click **Help** on the main menu and choose **Troubleshooting and request support**.
- 2 The **Troubleshooting** dialog opens.
- 3 In the **Select Category** drop-down list, select a category you are having a problem with. The area in the middle of the dialog will be populated with a list of existing KB articles related to that category.
- 4 Click an article of interest to read in a web browser.
- 5 You can also click **Knowledge Base Index** or **Forums** links to go to the Parallels knowledge base or Parallels forums.

Request support

If you can't find a solution for your problem using the options described above, you can send a support request to Parallels. When you do, the collected logging information is retrieved and attached to the email, so that Parallels Support can analyze it. See **Logging** (p. 451) for more information.

Note: A support request creates a support ticket, which is then sent to Parallels Support. If you already have a request support ticket, you can send just the system report to Parallels without creating an additional (and identical) ticket. See the **Send a report** subsection below. Please note that if you don't have a valid RAS subscription or a support contract, the ticket will not be created. In order to receive support, you will need to purchase a subscription or support contract.

Before you request support, please make sure that you have a mailbox setup in the RAS Console. If you haven't set up a mailbox yet, do it as follows:

- 1 In the RAS Console, navigate to **Administration > Mailbox**.
- 2 Enter your outgoing email server information, your email address, and the security/authentication information if needed.
- 3 You can send a test email by entering an email address in the field provided and clicking the **Send Test Email** button.

To send a support request to Parallels:

- 1 In the **Troubleshooting** dialog, click the **Send Support Request** button.
- 2 The **Contact Support** dialog opens.
- 3 Enter your full name and your company name.
- 4 Enter the subject. This will be used as a subject in the email that will be sent to Parallels Support.
- 5 In the **Enter your query** box, describe the issue the best you can.
- 6 Use the **Attachment** field to attach a file to the email. Click the [...] button to browse for a file. You can attach a picture or any other file that you think might help the Parallels Support to find a solution. Please note that the log files and the Parallels RAS settings are collected and attached to the email automatically, so you don't have to do it yourself.
- 7 In the drop-down list at the bottom of the dialog, select whether you want to send the email or save it (including the automatically collected information) as a zip file.
- 8 Depending on the action selected in the previous step, click **Send** to send the email or **Save** to save it as a zip file on your local drive or a network folder.

Send a report

If you already have a support request ticket, you can send just a system report to Parallels without creating a new ticket.

To send a report:

- 1 In the console, click **Help** on the main menu and choose **Upload System Report to Parallels**.
- 2 A dialog opens displaying the progress bar.
- 3 Once the system report data is collected and sent to Parallels, a message box is displayed containing the report number.
- 4 Click **OK** to finish.

Logging

Parallels RAS components are monitored and logs are created containing relevant information. Logs are used by Parallels RAS support engineers to analyze possible issues with a Parallels RAS installation. As a Parallels RAS administrator you have the ability to set the log level for a specific component or multiple components. By default, the standard log level is used, which collects and saves only the essential information. A Parallels RAS support engineer can ask you to enable the extended or verbose log level when an additional information is required to analyze an issue.

To set the log level for a specific component/server, navigate to the page in the RAS Console where the components of that type are listed (e.g. RD Session hosts, VDI, Gateways, Publishing Agents), select a component and then click **Tasks** (or right-click) > **Troubleshooting** > **Logging** > **Configure**. This opens the **Set Log Level** dialog where you can choose a log level from the following:

- **Standard** — This is the standard log level that records only the most important events. Unless you are asked by Parallels RAS support to use one of the log levels described below, you should always use this one.
- **Extended** — This logging involves more information than the standard logging, but it slows down the system because of the additional information that it needs to collect.
- **Verbose** — Verbose logging involves even more information than the extended logging and can slow down your system significantly.

Please note that to avoid degraded performance, extended and verbose logging should only be enabled for a limited time period (enough to collect the necessary information for analysis). You can set this time period using **Reset to the standard level after** option. The default value is 12 hours. In specific cases, a Parallels support engineer will advise you whether this time period should be set to a different value. Once this time period is over, the log level will be reset back to standard.

To retrieve a ZIP archive containing the log files, click **Tasks** (or right-click) > **Troubleshooting** > **Logging** > **Retrieve** and then specify a location where you want the file to be saved. The **Clear** item in the same context menu clears all logs.

Note that you can also set the log level on the **Farm** > <Site> > **Settings** > **Global logging** tab, where you can see RAS components of all types in one list. For more information, see **Site settings** (p. 433).

Log rotation

Parallels RAS log rotation works as follows:

- 1 When the total size of all log files reaches a predefined size (200 MB by default), the logs are archived. This is done log by log by appending the current timestamp to the filename and starting a new empty log file.
- 2 A new ZIP file is created for each old log named %logname%_%DATE%.zip . (e.g. console_10.06.2018.zip, controller_10.06.2018.zip).
- 3 Renamed old logs are moved to the ZIP file. Parallels RAS keeps five ZIP files by default.
- 4 When the maximum number of archived files is exceeded, the oldest file is deleted.

- 5 This log rotation mechanism guarantees that the total log file size never exceeds $X * Y * Z$ MB, where X is the total size of all log files (200 MB by default), Y is the maximum number of ZIP files (5 by default) and Z is the number of RAS components.
- 6 The X and Y values from the example above are pre-configured in Windows registry on a computer hosting a given RAS component. The default values are the same for every RAS component. To modify the values, navigate to HKEY_LOCAL_MACHINE > SOFTWARE > Wow6432Node > Parallels and set the LogMaxSize and LogMaxBackups values for a RAS component.

Suggest a feature

If you have an idea of a new feature for Parallels RAS, we would like to hear from you! To suggest a feature, in the RAS Console, click **Help** on the main menu and choose **Suggest a Feature**. This will take you to the **Parallels RAS Feature Suggestion** web page where you can communicate your ideas to us. Please note that you must be signed in using your Parallels account email address and password to post in the feature suggestion forum.

Parallels RAS Management Portal

This chapter gives you an overview of Parallels RAS Management Portal. For the complete information, please refer to **Parallels RAS Management Portal Guide**, which is available on the Parallels website at <https://www.parallels.com/products/ras/resources/>.

In This Chapter

Overview	454
Prerequisites	455
Installation	455
Log in to RAS Management Portal	456
Configure RAS Web Administration Service	456
RAS Management Portal user interface	457

Overview

Parallels® RAS Management Portal is a modern web-based configuration and administration console designed for Parallels RAS administrators using a desktop/laptop computer or a mobile device to carry out configurations and day-to-day activities.

Parallels RAS Management Portal provides administrators with ability to:

- Centrally deploy, manage, and configure essential Parallels RAS components such as RD Session Hosts, Publishing Agents and Secure Client Gateways.
- Publish various resources from RD Session Hosts.
- Configure FSLogix Profile Container settings.
- Configure printing and scanning settings.
- Manage SSL certificates.
- Configure connection settings and MFA (Google Authenticator or other Time-based One-time Password (TOTP) apps such as Microsoft Authenticator).
- Monitor and manage user sessions.
- Manage administrative accounts and sessions
- Configure mailbox.

- Manage your license.
- Contact support and provide necessary system reports.

Note: More features and capabilities that are currently available in the desktop-based Parallels RAS Console will be included in Parallels RAS Management Portal in future releases until it becomes the main management tool for Parallels RAS.

Management of Azure Virtual Desktop capabilities included in Parallels RAS Management Portal are experimental and expected to be released in upcoming versions.

Prerequisites

RAS Management Portal can run in any modern web browser supporting HTML5 such as Microsoft Edge (Chromium-based), Google Chrome, Mozilla Firefox, Safari, etc.

Make sure your Windows Server has the following updates installed (RAS Management Portal depends on them):

- Windows Server 2008 R2: KB2999226 and KB2533623
- Windows Server 2012 R2: KB2999226

Newer versions of Windows Server do not require any specific updates.

The web service listens to web requests on the following ports by default:

- HTTPS: 20443
- HTTP: 20080

Installation

To enable RAS Management Portal in a RAS Farm, you need to install the RAS Web Administration Service component. The component is installed automatically when you do a clean Parallels RAS install using the "Typical" installation option. You can also install the component using the "Custom" installation option and choosing the "RAS Web Administration Service" as the component to install. For example, if you want to install RAS Management Portal on a dedicated machine, you should use the "Custom" installation option and select "RAS Web Administration Service" as a component to install.

After the RAS Web Administration Service is installed, you need to configure it. Specifically, you need to specify a RAS Farm that the RAS Management Portal will be used to manage, and you also need to configure a number of other parameters. For complete instructions, please see **Configure RAS Web Administration Service** (p. 456).

Log in to RAS Management Portal

To open RAS Management Portal on the machine where you've installed the RAS Web Administration Service, navigate to **Apps > Parallels** and click **Parallels Remote Application Server Management Portal**.

To log in to RAS Management Portal from a remote computer, enter the following URL in a web browser:

```
https://<server-address>:20443
```

The <server-address> is the FQDN or IP address of the server where the RAS Web Administration Service is installed. By default, port 20443 is used for HTTPS connections. You can change the port number if needed as described in **Configure RAS Web Administration Service (p. 456)**.

On the **Welcome** page, enter your RAS administrator username and password and click **Sign in**.

Configure RAS Web Administration Service

Before you begin, you may need to configure the RAS Web Administration Service as described below:

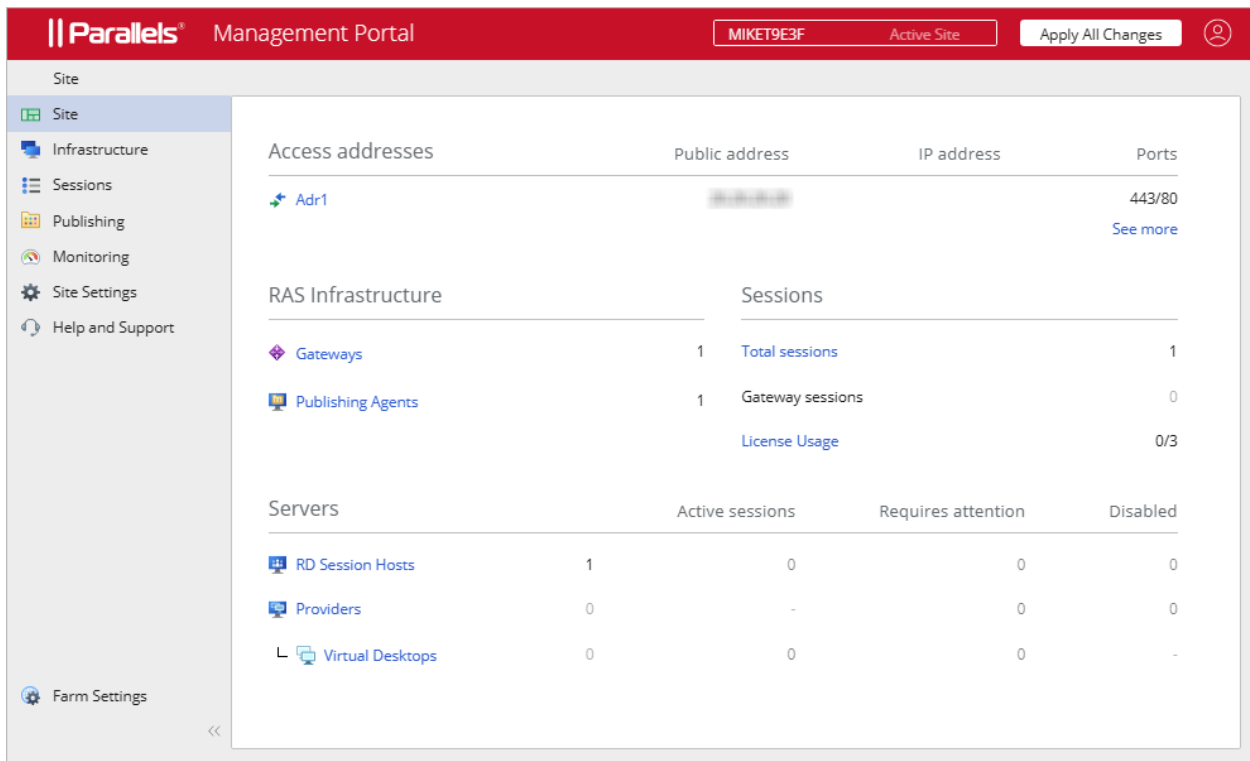
- 1 In RAS Management Portal, click the "User" icon in the upper right-hand corner and choose **Configure Management Portal**.
- 2 You will be asked to sign in again. Note that the RAS Web Administration Service must be running on the local server for this sign in to work. This is necessary to prevent users from remote servers to enter the RAS Web Administration Service configuration pages.
- 3 Enter the username and password of a member of local administrators or domain administrators and click **Sign in**.
- 4 The **RAS Management Portal Configuration** page opens.
- 5 In the **RAS Farm Address** field, specify the RAS Farm address that this RAS Management Portal will manage. This is the RAS Publishing Agent address installed in the Farm.
- 6 In the **Advanced Settings** section, specify the following:
 - **Certificate:** A certificate to use for this connection. Click **Upload** to select a certificate.
 - **Certificate Password:** The certificate password.
 - **Port:** The port number on which RAS Management Portal listens for connections. The default port is 20443. This port number is chosen not to conflict with RAS Secure Client Gateway ports. You can change it to 443 (if possible), in which case the port number doesn't need to be included in the connection URL. You can also change it to any custom port. For example, the default "URL": "https://*:20443" can be changed to "URL": "http://*:20080".

- **Admin Session Timeout:** The timeout after which the admin session will be disconnected.
- **Polling Interval:** The interval at which RAS Management Portal will update the information displayed in it. You can increase this number up to 30 seconds if you have a large number of admins working at the same time and/or if you have a large number of hosts, sessions, etc.

7 Click **Save** when done.

RAS Management Portal user interface

All navigations in the RAS Management Portal start from the sidebar on the left, which lists management categories. The **Site** category is selected by default.



Categories

The following table lists all available categories that can be managed in the RAS Management Portal. The Root Administrator can see and manage all categories. Administrators of other types (Power, Custom) may need permissions to see a particular category.

Category	Description
Site	Displays the current Site overview.
Infrastructure	RAS infrastructure management, including RD Sessions Hosts, VDI, Gateways, Publishing Agent, etc.

Sessions	Session management.
Publishing	Publishing and published resources management.
Monitoring	RAS Performance Monitor.
Site Settings	Connection, authentication, FSLogix, Universal printing and scanning.
Help and Support	Help and support.
Farm Settings	Displayed at the bottom of the sidebar on the left, this category manages global Farm settings, such as Administrators, Mailbox, Licensing.

Each category is described in detail later in this guide.

Admin permissions

Some categories and actions in the RAS Management Portal may not be viewed or allowed depending on the Admin permissions configured in the desktop RAS Console. For the information about how to configure administrator permissions, please refer to the **Parallels RAS Administrator's Guide**. In the guide, look for the **Administrator Account Permissions** topic. The guide is available on the Parallels website at <https://www.parallels.com/products/ras/resources/>.

Subcategories

Some categories have subcategories (namely **Infrastructure** and **Site Settings**). When you selected a category, the right side of the RAS Management Portal may include one or more additional panes where you can select a subcategory.

Navigation bar

Some components have their settings and information grouped by functionality (e.g. General, Properties, Sessions, etc.). When you view component properties, a navigation bar is displayed in the middle allowing you to browse these settings. When you select an item in the navigation bar, the settings are displayed in the right pane.

Breadcrumbs

As you select categories, subcategories, individual items, a breadcrumb trail is displayed at the top of the page to show where you are. To take one or more steps back, click a link in the trail.

Page header items

The page header includes the following items:

- The Farm and the current Site names. If you have more than one Site, you can select one from the drop-down list. The RAS Management Portal will switch to that site allowing you to manage the Site components.

- The "User" icon is a drop-down list with the following items: Current user name (e.g. **Administrator**); **About** (opens the About dialog); **Give feedback** (takes you to a web page where you can give feedback to Parallels); **Configure Management Portal** (p. 456), **Logout** (logs you out).
- **Apply All Changes**: This button applies changes that you've made in the RAS Management Portal to Farm components. When you create or modify components and objects, the changes are not applied to Farm components automatically and don't have any effect on the Site or Farm. When you click the **Apply All Changes** button, the changes are applied across the Farm or Site. Note that you shouldn't always click this button every time you make a change. If you are working on a task that requires multiple changes in different areas, complete all of them and then click the **Apply All Changes** button, so all changes are applied together.

Editing

When you open a view where you can modify some settings, the view is normally read-only. To enable editing, click the **Edit** button in the upper right-hand corner. The button name changes to **Save**. When done editing, click **Save**. To discard the changes, click **Cancel**.

Please note that an object that is opened for editing by an admin cannot be edited by another admin at the same time. If you try to enable editing for such an object, you will get an error with the name of the admin who has the object locked.

Edit toolbar

Some views (specifically lists) have a toolbar in the upper right-hand corner from which you can execute actions. To see a toolbar item name, hover over it with the mouse. The standard items (icons) on the toolbar are the following:

- **Show filter**: Specify a filter to show only the entries that match it.
- **Select columns**: Select table columns to display or hide.
- **Add**: Add a new entry. For example, add a new Gateway or RD Session Host, etc.
- **Refresh**: Refresh the view.
- **Ellipsis**: The ellipsis menu may have different items in different types of views. Some items have a corresponding toolbar items (e.g. **Add**, **Refresh**).

Other items may be present depending on the view you are in. For example, **Show running processes** and **Show sessions**.

Wizards

When you add a component to a Farm, a wizard usually opens which takes you through a series of pages where you specify component settings and options. A wizard has the usual **Next** and **Back** navigation buttons, and the **Cancel** button that closes the wizard and cancels the operation.

Modal dialogs

Clicking some menu and navigation bar items brings up a modal dialog. These are usually items that require you to confirm an action or enter additional information.

Object properties views

All objects (components) in the RAS Management Portal have properties. To view these properties, you select a category and a subcategory and click the object name in the list. This opens a view where object properties are displayed with its own navigation bar from which you can configure the object, perform actions, and view additional information.

Parallels RAS APIs

Parallels RAS comes with APIs to help you develop custom applications that integrate with it. This includes RAS PowerShell API and RAS REST API.

In addition, the RAS HTML5 Gateway API and Parallels Client URL scheme allow you to integrate with Parallels Client for Windows/macOS/Linux/iOS/Android and the RAS HTML5 Client.

In This Chapter

RAS PowerShell API.....	461
RAS REST API	463
RAS HTML5 Gateway API and Parallels Client URL scheme.....	467

RAS PowerShell API

RAS PowerShell API is intended for RAS administrators who would like to automate their RAS administration. The API includes commands to perform most of the RAS management tasks.

Parallels RAS requirements

The Parallels RAS PowerShell API version must match the version of the RAS Publishing Agent with which it communicates. Since the two components can be installed separately, you need to make sure that their versions match.

Microsoft Windows component requirements

The following components must be installed on the computer where you'll be executing Parallels RAS PowerShell cmdlets:

- Windows PowerShell 3.0 or higher
- Microsoft .NET Framework 4.5.2 or higher

Installation

To install Parallels RAS PowerShell, run the standard Parallels RAS installer, choose **Custom** installation, and then select to install the **Parallels RAS PowerShell** component. Follow the onscreen instructions to install the component.

RAS PowerShell API versions

The RAS PowerShell API has undergone changes in Parallels RAS 18 as follows:

- The RAS PowerShell module name has changed from PSAdmin to RASAdmin.
- Most of the commands now have the "RAS" prefix, such as RASGW or RASApply.
- API versions: Version 2.0 (latest) and version 1.0 are supported for backward compatibility.

Note that the API version 1.0 is still available in the current RAS PowerShell module. If you have existing scripts that use the older module and command names, you can use them with minimal changes. To do that, you need to load the API version 1.0 when you import the RAS PowerShell module. See below for more information about API versions.

Version 2.0

This version is the one loaded by default by the system or if the `-RequiredVersion` parameter is not specified when importing the module. See **RAS PowerShell API concepts** for examples.

Version 1.0

This version keeps backward compatible with the old PSAdmin module to allow administrators to use existing scripts with minor modification. This version includes:

- Cmdlet aliases
- Aliased parameters
- Returns old and new properties

RAS PowerShell API concepts

To quickly get started with RAS PowerShell, do the following:

- 1 Open the Windows PowerShell console.
- 2 Import the Parallels RAS PowerShell module using one of the following commands:
 - `Import-Module RASAdmin` – Loads the current API (version 2.0).
 - `Import-Module RASAdmin -RequiredVersion 1.0` – Loads the API version 1.0.
- 3 Create a Parallels RAS session by executing the `New-RASSession` cmdlet (see example below). Substitute the server name (in quotes) with the name or IP address of your Parallels RAS Licensing Server. Type your RAS administrator username and password when prompted:
`New-RASSession -Server "server.company.dom"`
- 4 Execute the following cmdlet to see the list of cmdlets included in the Parallels RAS PowerShell module:
`Get-Command -Module RASAdmin`

- 5 Execute other cmdlets. For example, try executing the `Get-RASGW` cmdlet to retrieve information about RAS Secure Client Gateway(s). The example below returns information about all RAS Secure Client Gateways available in the RAS Licensing Server Site:

```
Get-RASGW
```

- 6 To see help for a cmdlet, execute `Get-Help` passing a cmdlet name:

```
Get-Help Get-RASGW
```

- 7 To apply changes you've made to the Farm configuration, use the `Invoke-RASApply` cmdlet (this performs the same action as the **Apply** button in the RAS Console):

```
Invoke-RASApply
```

- 8 To activate a Parallels RAS license, use the `Invoke-RASLicenseActivate` cmdlet:

```
Invoke-RASLicenseActivate
```

When executing the cmdlet above, you'll be prompted to enter your Parallels account email address and password. You can include an optional `-key` parameter and specify a Parallels RAS license key. If omitted (as in the example above), Parallels RAS is activated as a trial.

Parallels RAS PowerShell Guide

To view and download the new **Parallels RAS PowerShell Guide** version 2.0, visit Parallels website at <https://www.parallels.com/products/ras/resources>.

RAS REST API

This section gives you an introduction to the RAS REST API. Read it to learn about system requirements, installation, configuration, and basic usage.

Installation

To enable RAS REST API in a RAS Farm, you need to install the RAS Web Administration Service. It can be installed on the RAS Publishing Agent server or any other server. If you install the service on a separate server, you will need to change its configuration (after the installation) to point to RAS Publishing Agent. By default, the configuration points to "localhost".

Note: If you've already configured and are using Parallels RAS Management Portal, you may skip this step because you should already have the RAS Web Administration Service installed.

To install RAS Web Administration Service:

- 1 Run the Parallels RAS installer on the RAS Publishing Agent or any other server.
- 2 On the **Select Installation Type** page, select **Custom**.

- 3 On the next page, select to install the **Parallels RAS Web Administration Service** component.
- 4 Click **Next** and follow the onscreen instructions.

Configure RAS Web Administration Service

If the RAS Web Administration Service was installed on a separate server, you need to modify its configuration and specify the RAS Publishing Agent server address. You may also want to change the port number and certificate information in the same configuration file. For details about configuring RAS Web Administration Service, please refer to KB article <https://kb.parallels.com/en/124701>.

When modifying the service configuration, please note the following:

- In the configuration JSON file, the RAS Publishing Agent address is specified using the "LicenseServer" parameter.
- The default HTTPS port number is set to 20443. This number is chosen not to conflict with RAS Secure Client Gateway ports. You can change it to 443 (if possible), so when opening the portal, you don't have to include the port in the URL.

Permissions

To access any of the RAS REST resources, the user executing a request must have sufficient rights to access a particular resource. These are basically the same rights a RAS administrator has in the Parallels RAS Console. For example, a root administrator can access any of the RAS REST resources. On the other hand, a power administrator who doesn't have rights to modify Site settings (as an example) will not be able to access a corresponding REST resource. Similarly, a custom administrator who, for instance, only has rights to view and modify RD Session Hosts will be able to access just that particular REST resource and no other.

Getting started

Applications communicate with Parallels RAS by sending HTTP or HTTPS requests. Parallels RAS answers with a JSON file in a response to every HTTP request.

All HTTP requests that you will use to retrieve and manage Parallels RAS resources have the following base structure:

```
https://<API-host>/api/<URI>
```

The parameters in the above URL are:

- <API-host> is the IP address or FQDN of the server on which the RAS Web Administration Service is installed.
- <URI> is a path to a REST resource that you would like to work with.

Logging in and sending requests

This section contains an example of RAS REST API usage that can help you quickly get started. The example demonstrates how to:

- 1 Login to Parallels RAS and obtain a session token.
- 2 Retrieve the information about all available RD Session Hosts.
- 3 Retrieve the information about a specific RD Session Host.
- 4 Modify RD Session Host properties.

Log in to Parallels RAS and obtain a session token

Before you can access any of the resources, you need to log in to Parallels RAS using administrator credentials and obtain a session token. This is accomplished by sending the following request:

```
POST https://<API-host>/api/session/logon
```

Request headers: The logon request must contain just the Content-Type request header. Subsequent requests must additionally contain the auth_token header, as you'll see in the examples that follow this one.

Content-Type: application/json; api-version=1.0

Request body: The request body must contain the RAS administrator user name and password.

```
{
  "username": "USER",
  "password": "PASSWORD"
}
```

Response: After sending the logon request, you will receive a reply containing the session token, which you will use in all subsequent requests:

```
{
  "authToken": "Lj+KddoJkANhzvbDRvB=K=DFCcroRjXJHeeWGbGlIRKaz-EXplbmhVWvWTiDVqtOq"
}
```

Retrieve information about RD Session Hosts

Now that we have the session token, we can send requests to access various resources. In this example we'll first obtain the information about all available RD Session Hosts. In the example that follows, we'll obtain the information about a specific RD Session Host.

To retrieve the RD Session Host info, send the following request:

```
GET https://<API-host>/api/RDS
```

Request headers: This time the auth_token request header must also be included and must contain the session token that we've obtained earlier.

Content-Type: application/json; api-version=1.0

auth_token: Lj+KddoJkANhzvbDRvB=K=DFCcroRjXJHeeWGbGIIrKaz-EXplbmhVWvWTiDVqtOq

Response: The response will look similar to the following (with multiple RD Session Hosts in the Farm each block of the result set will contain the information about an individual server).

```
[
  {
    "directAddress": "IP_ADDR",
    "rasTemplateId": 0,
    "inheritDefaultAgentSettings": true,
    "inheritDefaultPrinterSettings": true,
    "inheritDefaultUPDSSettings": true,
    "inheritDefaultDesktopAccessSettings": true,
    "port": 3389,
    ...
    "restrictDesktopAccess": false,
    "restrictedUsers": [],
    "server": "IP_ADDR",
    "enabled": true,
    "description": "",
    "siteId": 1,
    "id": 2
  }
]
```

Retrieve information about a specific RD Session Host

To retrieve the information about a specific server, we'll use the same request as above but will add the server ID at the end:

```
GET https://<API-host>/api/RDS/2/
```

The response will also be similar to the example above and will contain the information just for the specified server.

Modify RD Session Host properties

In this example we'll modify a property of the RD Session Host that we retrieved earlier. For simplicity let's modify the "description" field.

The request to modify properties of an RD Session Host has the following syntax:

```
PUT https://<API-host>/api/RDS/2/
```

Note the "2" at the end of the request, which specifies the ID of the RD Session Host that we want to modify.

Request headers:

- Content-Type: application/json; api-version=1.0
- auth_token: Lj+KddoJkANhzvbDRvB=K=DFCcroRjXJHeeWGbGIIrKaz-EXplbmhVWvWTiDVqtOq

Request body:

```
{
  "description": "description was updated!"
}
```

Response: If the PUT request succeeds, you will get an empty response with code "204: No Content". To verify that the "description" field was in fact modified, let's use the same GET request that we used earlier: GET https://<API-host>/api/RDS/2/

As we can see, the result now contains the updated "description" field:

```
[
  {
    "directAddress": "IP_ADDR",
    "rasTemplateId": 0,
    "inheritDefaultAgentSettings": true,
    ...
    "server": "IP_ADDR",
    "enabled": true,
    "description": "description was updated!",
    "siteId": 1,
    "id": 2
  }
]
```

More information

Parallels RAS REST API comes with the **Parallels RAS REST API Guide**. The guide contains more examples and the complete resource and schema reference. To view and download the guide, visit <https://www.parallels.com/products/ras/resources/>.

RAS HTML5 Gateway API and Parallels Client URL scheme

RAS HTML5 Gateway API and Parallels Client URL scheme allow you to integrate with Parallels clients.

Using the RAS HTML5 Gateway API or the URL scheme, you can implement an in-house solution, such as an application hub or web portal, for authenticating users and launching remote applications, desktops, and other published resources. Such an implementation is possible by integrating a custom solution with Parallels Clients, including Parallels Clients for supported platforms (Windows, macOS, Linux, iOS, Android) and RAS HTML5 Client.

The following is a quick summary of the API and the URL scheme:

- **RAS HTML5 Gateway API** — provides connection, user authentication, and resource launching methods called from a web browser via the RAS HTML5 Gateway.
- **Parallels Client URL scheme** — a custom URL scheme that allows you to perform actions in a Parallels Client installed on a user device. Actions include configuring a connection, authenticating a user, and launching published resources.

RAS HTML5 Gateway API and Parallels Client URL scheme are described in detail in the **Integrating with Parallels Clients** guide, which is available for download on the Parallels website at the following location: <https://www.parallels.com/products/ras/resources/>.

Appendix

In This Chapter

Microsoft license requirements in Parallels RAS	469
Port reference	473
RAS performance counters	483

Microsoft license requirements in Parallels RAS

This section is to be used as guidance to provide clarity on Microsoft license requirements in a Parallels RAS environment while not used as an exclusive list. It is recommended to refer to your Microsoft licensing partner for further information.

Microsoft license requirements include:

General

- Any Windows Server and Desktop Operating System (OS) to be used.
- Windows Server OS to be accessed must be covered by Microsoft Windows Server Client Access Licenses (CALs).

RD Session Hosts

If Windows Server is accessed remotely (for non-administrative work) then you need Remote Desktop Service (RDS) access license:

- RDS CALs are required for users or devices that want to utilize Remote Desktop Service functionality on Windows Server. The following types of RDS CAL are available:
 - a** RDS Device CAL: Permits one device (used by any user) to use Remote Desktop Services functionality on any of your servers.
 - b** RDS User CAL: Permits one user (using any device) to use Remote Desktop Services functionality on any of your servers.
 - c** RDS External Connector: Permits multiple external users to access a single Remote Desktop server. If you have multiple servers, you need multiple external connectors in addition to any required Windows Server External Connectors.

You may choose to combine RDS Device CALs and RDS User CALs simultaneously with the server software. Regular User or Device CALs are required in addition to the RDS User or RDS Device CALs.

- RDS SAL is a service that provides a Microsoft Remote Desktop Service Subscriber Access License (called an "RDS SAL") on Virtual Machines created in Compute Resource. This makes it possible for three or more users to connect to a remote desktop (RD Session Host) for a specific Virtual Machine in Compute Resource (for SPLA partners).

Read more:

- License your RDS deployment with client access licenses (CALs): <https://docs.microsoft.com/en-us/windows-server/remote/remote-desktop-services/rds-client-access-license>.
- RDS Licensing Data Sheet https://download.microsoft.com/download/6/B/A/6BA3215A-C8B5-4AD1-AA8E-6C93606A4CFB/Windows_Server_2012_R2_Remote_Desktop_Services_Licensing_Datasheet.pdf.
- RDS CAL overview and FAQ <https://download.microsoft.com/download/3/D/4/3D42BDC2-6725-4B29-B75A-A5B04179958B/Licensing-Windows-Server-2012-R2-RDS-and-Desktop-Apps-for-RDS.pdf>.
- Licensing of Microsoft Desktop Application Software for use with Windows Server RDS https://download.microsoft.com/download/3/d/4/3d42bdc2-6725-4b29-b75a-a5b04179958b/desktop_application_with_windows_server_remote_desktop_services.pdf.

Hypervisor and VDI

- 1 In case using Microsoft Hyper-V as a hypervisor, Microsoft Windows Server Operating System (OS) Licenses are required

Read more:

- Windows Server 2022 license datasheet <https://www.microsoft.com/en-us/windows-server/pricing>.
 - Windows Server 2019 license datasheet https://download.microsoft.com/download/7/C/E/7CED6910-C7B2-4196-8C55-208EE0B427E2/Windows_Server_2019_licensing_datasheet_EN_US.pdf.
 - Windows Server 2016 license datasheet <https://download.microsoft.com/download/7/2/9/7290EA05-DC56-4BED-9400-138C5701F174/WS2016LicensingDatasheet.pdf>.
- 2 In case using Virtual Desktop Infrastructure (VDI), Windows Software Assurance or Azure Virtual Desktop Access (VDA) licenses are required. Microsoft licenses Windows by access device:
 - Virtual desktop access rights are a benefit of Windows Client Software Assurance (SA). Customers who intend to use PCs covered under SA have access to their VDI desktops at no additional charge.

- Customers who want to use devices that do not qualify for Windows Client SA, such as thin clients, will need to license those devices with Azure Virtual Desktop Access (VDA) in order to access a Windows VDI desktop. Windows VDA is also applicable to third-party devices, such as contractor or employee-owned PCs.

Read more:

- Windows 11 licensing portal <https://www.microsoft.com/en-us/Licensing/product-licensing/windows>.
- Windows 10 licensing portal <https://www.microsoft.com/en-us/licensing/product-licensing/windows10?activetab=windows10-pivot:primaryr3>.
- Licensing Windows desktop operating system for use with virtual machines guide https://download.microsoft.com/download/9/8/d/98d6a56c-4d79-40f4-8462-da3ecba2dc2c/licensing_windows_desktop_os_for_virtual_machines.pdf.
- Licensing the Windows Desktop for VDI Environments <https://docs.microsoft.com/en-us/answers/storage/temp/12620-microsoft-vdi-and-vda-faq-v3-0.pdf>.

Microsoft Azure

Microsoft Online business services, such as Microsoft 365 or Microsoft Azure, require Azure AD for sign-in and to help with identity protection. If you subscribe to any Microsoft Online business service, you automatically get Azure AD with access to all the free features. To enhance your Azure AD implementation, you can also add paid capabilities by upgrading to Azure Active Directory Premium P1 or Premium P2 licenses.

Read more:

- Azure AD Implementations <https://docs.microsoft.com/en-us/azure/active-directory/fundamentals/active-directory-what-is>
- Azure hybrid benefits <https://azure.microsoft.com/en-us/pricing/hybrid-benefit/>

Azure Virtual Desktop

- Access to Windows 10 Enterprise multi-session, Windows 11 Enterprise multi-session, Windows 10 Enterprise and Windows 11 Enterprise desktops and apps is provided at no additional cost (excluding compute, storage and networking costs) if you have one of the following per user licenses:
 - a Microsoft 365 E3/E5
 - b Microsoft 365 A3/A5/Student Use Benefits
 - c Microsoft 365 F3
 - d Microsoft 365 Business Premium
 - e Windows 10 Enterprise E3/E5
 - f Windows 10 Education A3/A5
 - g Windows 10 VDA per user

- Access to desktops powered by Windows Server Remote Desktop Services running Windows Server 2012 R2 and newer is provided at no additional cost (excluding compute, storage and networking costs) if you have a per-user or per-device RDS CAL license with active Software Assurance (SA).

Read more:

- Azure Virtual Desktop pricing overview <https://azure.microsoft.com/en-us/pricing/details/virtual-desktop/>

FSLogix

You are eligible to access FSLogix Profile Container, Office 365 Container, Application Masking, and Java Redirection tools if you have one of the following licenses:

- Microsoft 365 E3/E5
- Microsoft 365 A3/A5/ Student Use Benefits
- Microsoft 365 F1/F3
- Microsoft 365 Business
- Windows 10 Enterprise E3/E5
- Windows 10 Education A3/A5
- Windows 10 VDA per user
- Remote Desktop Services (RDS) Client Access License (CAL)
- Remote Desktop Services (RDS) Subscriber Access License (SAL)

FSLogix solutions may be used in any public or private data center, as long as a user is properly licensed.

Read more:

- FSLogix Overview <https://docs.microsoft.com/en-us/fslogix/overview>.

Microsoft SQL Server

SQL Server is required if using Parallels RAS Reporting. SQL Server installation may be based on:

- SQL Express which is free but has a database size limit of 10 GB.
- SQL Server commercial edition Standard or Enterprise, using Core based licenses or Server + CAL based licenses.

Read more:

- SQL Server 2019 licensing guide <https://download.microsoft.com/download/6/6/0/66078040-86d8-4f6e-b0c5-e9919bbcb537/SQL%20Server%202019%20Licensing%20guide.pdf>

App-V

App-V is not licensed on its own, but included in other license agreements such as Microsoft Volume Licensing, Windows Software Assurance Microsoft, Remote Desktop Services (RDS) CAL, as part of a wider Microsoft licensing agreement. For instance, with an RDS CAL (either per-user or per-device), App-V client may be used on RD Session Host to deliver App-V applications.

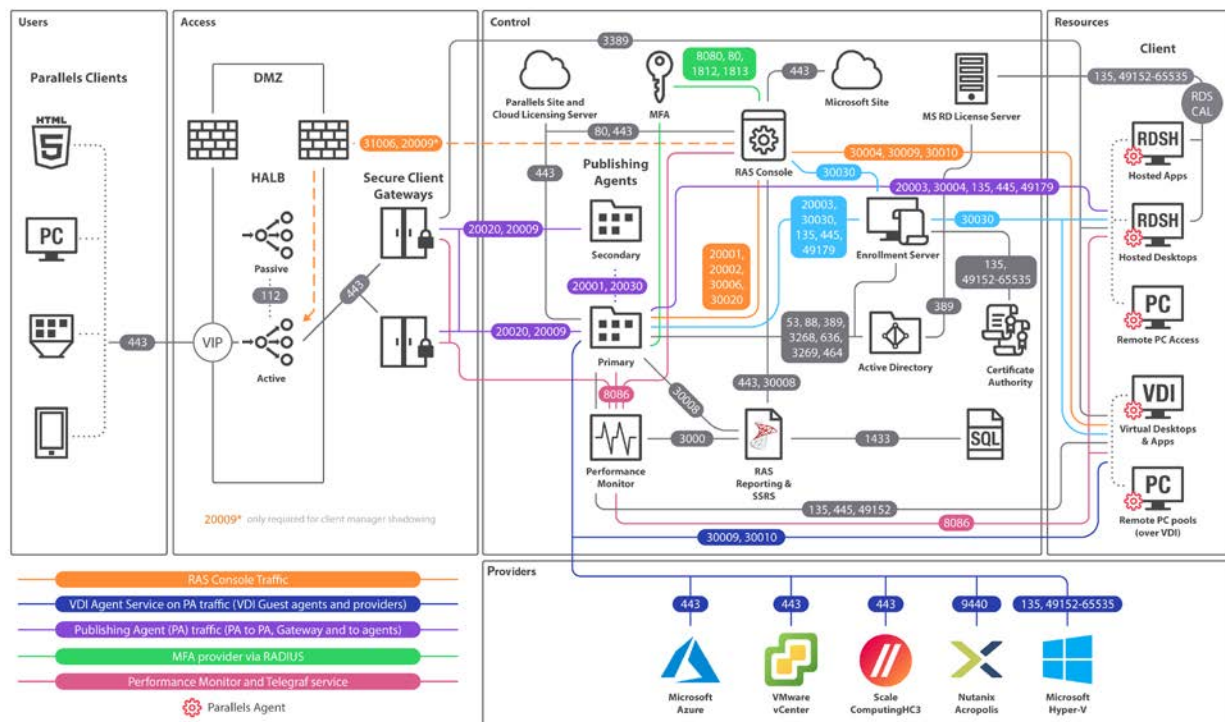
To license App-V correctly it is recommended you to engage with a Microsoft Partner (solution provider) knowledgeable on Microsoft Volume Licensing (list of Microsoft Partners: <https://pinpoint.microsoft.com/en-us/search?type=companies&competency=100010>).

Other References

For a detailed list of Microsoft Volume Licensing Product Terms please see <https://www.microsoftvolumelicensing.com/Downloader.aspx?documenttype=PT&lang=English>.

Port reference

The following diagram illustrates communication ports used in Parallels RAS.



The above diagram include SAML SSO components such as RAS Enrollment Server, however it does not include Tenant Broker.

Tip: If you are reading the PDF version of this guide, click the following link to view the full-sized diagram in a web browser: https://download.parallels.com/ras/v18/docs/en_US/Parallels-RAS-18-Administrators-Guide/index.htm#47092.

Parallels Client

Source	Destination	Protocols	Ports	Description
Parallels Client	HALB	TCP, UDP	80, 443	Management and user session connections.
		TCP, UDP	20009	Device Manager shadowing via Firewall (indirect network connection).
	RAS Secure Client Gateway Forwarding mode	TCP, UDP	80, 443	Management and user session connections.
		TCP, UDP	3389	Optional - Used for user session if RDP load balancing is enabled (Standard RDP).
	UDP	20000	Secure Client Gateway lookup broadcast.	
	RAS Secure Client Gateway Normal mode	TCP, UDP	80, 443,	Management and user session connections.
		TCP, UDP	3389	Optional - Used for user session if RDP load balancing is enabled (Standard RDP).
		TCP, UDP	20009	Device Manager shadowing via Firewall (indirect network connection)
	UDP	20000	Secure Client Gateway Lookup Broadcast	
Session host (VDI, RDS, RemotePC)	TCP, UDP	3389	Used for user session connections in Direct Mode only. RDP connection is always encrypted.	
Azure Virtual Desktop Services	TCP	443	Azure Virtual Desktop Gateway connection	
	UDP	3390	Used for user session connections in ShortPath mode only.	
Microsoft site	TCP	443	Download Microsoft Remote Desktop (MSRDC) client	
Parallels site	TCP	80, 443	Check for updates and download Parallels Client	

Web browsers

Source	Destination	Protocols	Ports	Description
Web browser (HTML5)	RAS Web Admin Service [RAS Management Portal]	TCP	20443	Admin access to HTML5 based Management Portal of RAS environment
	HALB	TCP	443	End-user access to Parallels RAS HTML5 Client (on Secure Client Gateway in Normal mode) through the HALB

	RAS Secure Client Gateway	TCP	443	End-user access to Parallels RAS HTML5 Client (on Secure Client Gateway in Normal mode)
--	---------------------------	-----	-----	---

HALB

Source	Destination	Protocols	Ports	Description
HALB	HALB	VRRP	112	HALB to HALB communication used for automatic assignment of VIP to active HALB.
	RAS Secure Client Gateway in Forwarding Mode	TCP, UDP	80, 443	Management and user session connections.
	RAS Secure Client Gateway in Normal Mode	TCP, UDP TCP, UDP	80, 443 20009	Management and user session connections. Device Manager shadowing via Firewall (indirect network connection).

RAS Secure Client Gateway

Source	Destination	Protocols	Ports	Description
RAS Secure Client Gateway in Forwarding mode	RAS Secure Client Gateway in Normal mode	TCP, UDP	80, 443	Management and user session connections.
		TCP, UDP	3389	Optional - Used for user session if RDP Load Balancing is enabled.
	RAS Performance Monitor	TCP	8086	Agent (Telegraf service) sends collected performance data to InfluxDB.
RAS Secure Client Gateway in Normal mode	Remote Desktop Services	TCP, UDP	3389	RDP Connections.
	RAS Publishing Agent	TCP	20002	RAS Publishing Agent service port - communications with RAS Secure Client Gateways and the RAS Console (in Normal mode only).
		TCP, UDP	20009	Device Manager shadowing via Firewall (indirect network connection) if RAS Console runs on RAS Publishing Agent
	RAS Performance Monitor	TCP	8086	Agent (Telegraf service) sends collected performance data to InfluxDB.
Localhost	TCP	20020	Communication with HTML5 Gateway web server (NodeJS).	

RAS Publishing Agent

Source	Destination	Protocols	Ports	Description
RAS Publishing Agent	AD DS controllers	TCP	389, 3268	LDAP
		TCP	636, 3269	LDAPS
		TCP,UDP	88	Kerberos
		UDP	53	DNS
	RAS Publishing Agent	TCP	20001 20030	Redundancy service. Communication between RAS Publishing Agents running in the same site.
	Parallels Licensing Server	TCP	443	RAS Publishing Agent (primary Publishing Agent in Licensing Site) communicates with Parallels Licensing Server (https://ras.parallels.com). Note: Not required for Tenant Broker RAS Publishing Agent (see the Tenant Broker section).
	RAS Performance Monitor	TCP	8086	Agent (Telegraf service) sends collected performance data to InfluxDB.
	RAS RD Session Host Agent	TCP, UDP	30004	Server for Publishing Agent requests.
	RAS VDI Agent	TCP, UDP	30006	VDI Agent communication port.
	RAS Remote PC Agent	TCP, UDP	30004	Remote PC Agent Communication Port (agent state, counters and session information)
	2FA Server(s)	TCP, UDP	8080, 80 1812, 1813	Deepnet/ Safenet Radius
	RAS Enrollment Server	TCP	30030	RAS Publishing Agent Sends RAS Enrollment Server connection Request
	RAS Reporting	TCP	30008	Master RAS Publishing Agent communicates with RAS Reporting (installed on the same host as SSRS).
	RAS Remote Installer Service	TCP	30020	Remote agent pushing
	RAS RD Session Host Agent RAS Guest Agent RAS Remote PC Agent RAS Publishing Agent RAS Secure Client Gateway RAS Enrollment Server	TCP	135, 445, 49179	Remote Install Push/Takeover of Software
SMTP	TCP	587	Notifdispatcher is the service which sends the emails using port specified in the Mailbox settings (+SSL/TLS)	

RAS Console

Source	Destination	Protocols	Ports	Description
RAS Console	RAS Reporting	TCP	30008	RAS Console is connected to primary RAS Publishing Agent which communicates with RAS Reporting (installed on the same host as SSRS). SSRS talks to SQL via TCP 1433 (or dynamic if 1433 is not established in the settings).
	SSRS	TCP	443	Reports retrieval.
	HALB	TCP, UDP	31006	Used for configuration.
	Parallels Client	TCP	50005	Shadowing from the RAS Console in case of direct network connection.
	RAS RD Session Host Agent	UDP, TCP	30004	Used for the "Check Agent" task. Used to manage components.
	RAS Guest Agent	TCP UDP	30010 30009	Used for the "Check Agent" task. Used to manage components.
	RAS Remote PC Agent	UDP, TCP	30004	Used for the "Check Agent" task. Used to manage components.
	RAS VDI Agent	UDP, TCP	30006	Used for the "Check Agent" task. Used to manage component.
	MFA Server(s)	TCP, UDP	8080, 80, 1812, 1813	Deepnet / Safenet / Radius
	Microsoft site	TCP	80, 443	Check for updates and download Parallels Client
	Parallels site	TCP	80	Check for updates and download Parallels Client
	RAS Performance Monitor	TCP	20002, 20001	Communication with Publishing Agent and redundancy.
	RAS Publishing Agent	TCP	20002, 20001	Communication with Publishing Agent and redundancy.
	RAS Enrollment Server	TCP, UDP	30030	Used for the "Check Agent" task. Used to manage components and for troubleshooting.
	Wyse Broker	UDP	1234 (outbound only) 68 (inbound only)	Wyse broker discovery request broadcast packet (V_WYSEBCAST). Wyse broker discovery reply packet (V_WYSETEST).
	SMTP	TCP	587	RAS Console can send test emails using port specified in the Mailbox settings (+SSL/TLS)

SSRS

Source	Destination	Protocols	Ports	Description
SSRS	Microsoft SQL Server	TCP	1433	RAS Console is connected to RAS Reporting

RAS Reporting

Source	Destination	Protocols	Ports	Description
RAS Reporting Service	MS SQL	TCP	1433	Store RAS activity information
	SSRS	TCP	8085, 443	Enumeration of reports (incl. custom reports)

RAS Web Administration Service (REST/Management Portal)

Source	Destination	Protocols	Ports	Description
RAS Web Administration Service	RAS RD Session Host Agent	TCP	30004	Log retrieval
	RAS Guest Agent	TCP	30010	Log retrieval
	RAS VDI Agent	TCP	30006	Log retrieval
	RAS Publishing Agent	TCP	20002, 20001 30020	Communication with PA and Redundancy Used during publishing to browse for installed applications or single file/folder browsing. 30020 - remote agent pushing (pre-RAS 18).
	RAS RD Session Host Agent RAS Guest Agent RAS Remote PC Agent RAS Publishing Agent RAS Secure Client Gateway RAS Enrollment Server	TCP	135, 445	Remote Install Push/Takeover of Software (pre-RAS 18).
	RAS Reporting Service	TCP	3000	Integration of RAS Reporting in Management Portal iFrame

RAS PowerShell

Source	Destination	Protocols	Ports	Description
RAS PowerShell	RAS RD Session Host Agent	TCP	30004	Log retrieval
	RAS Guest Agent	TCP	30010	Log retrieval
	RAS Remote PC Agent	TCP	30004	Log retrieval
	RAS VDI Agent	TCP	30006	Log retrieval
	RAS Publishing Agent	TCP	20002, 20001	Communication with PA and Redundancy Used during publishing to browse for installed applications or single file/folder browsing.

RAS VDI Agent

Source	Destination	Protocols	Ports	Description
RAS VDI Agent	RAS Publishing Agent	TCP	20003	Publishing Agent communication port.
	RAS Guest Agent	TCP	30010	TCP is used to send the commands.
		UDP	30009	UDP is used during the initial handshake.
	RAS Performance Monitor	TCP	8086	Agent (Telegraf service) sends collected performance data to InfluxDB - applicable to Hyper-V only.
	Hyper-V	TCP	135, 49152-65535	Used to check if the guest is powered on and send export, import, delete, shutdown, restart or suspend commands.
	Nutanix	TCP	9440	Used to check if the guest is powered on and sends clone, delete, shutdown, restart commands (RestAPI calls, PoSH, remote ncli).
	VMWare	TCP	443	Used to check if the guest is powered on and sends clone, delete, shutdown, restart and suspend commands.
	Microsoft Azure	TCP	443	Used to check if the guest is powered on and sends clone, shutdown, restart commands (via REST).
Scale	TCP	443	Used to check if the guest is powered on and sends clone, shutdown, restart commands (via REST).	

	Remote PC over VDI	TCP	135, 49152-65535	Used to check if the guest is powered on and sends shutdown, restart or suspend commands.
--	--------------------	-----	------------------	---

RAS Enrollment Server

Source	Destination	Protocols	Ports	Description
RAS Enrollment Server	AD DS controllers	TCP	389, 3268	LDAP
		TCP	636, 3269	LDAPS
		TCP,UDP	88	Kerberos
		UDP	53	DNS
	RAS Publishing Agent	TCP	20003	Settings synchronization and performance counters.
		UDP	20003	Deny Connection Request
Certificate Authority (CA)	TCP	135	DCOM/RPC ports	
	TCP	dynamic range 49152 - 65535		

RAS RD Session Host Agent

Source	Destination	Protocols	Ports	Description
RAS RD Session Host Agent	RAS Publishing Agent	TCP, UDP	20003	Used for communications with RAS Publishing Agents.
	Localhost	TCP	30005	For internal commands (memshell, printer redirector).
	FSlogix	TCP	443	Download FSlogix installer
	RAS Performance Monitor	TCP	8086	Agent (Telegraf service) sends collected performance data to InfluxDB.
	RAS Enrollment Server	TCP	30030	RAS RD Session Host Agent (PrIsSCDriver) connects to get logon credentials.

RAS Guest Agent

Source	Destination	Protocols	Ports	Description
RAS Guest Agent (used by Azure Virtual Desktop)	VDI Agent	TCP, UDP	30006	Communication with VDI Agent Subnet broadcast is sent to find VDI agent Regular UDP heartbeats

	Localhost	TCP	30005	For internal commands - memshell, printer redirector)
	RAS Performance Monitor	TCP	8086	Agent (Telegraf service) sends collected performance data to InfluxDB
	RAS Enrollment Server	TCP	30030	RAS Guest Agent (PrisSCDriver) connects to get logon credentials
	FSlogix	TCP	443	Download FSlogix installer

RAS Remote PC Agent

Source	Destination	Protocols	Ports	Description
RAS Remote PC Agent	RAS Publishing Agent	TCP, UDP	20003	Used for communications with RAS Publishing Agents
	Localhost	TCP	30005	For internal commands - memshell, printer redirector)
	RAS Performance Monitor	TCP	8086	Agent (Telegraf service) sends collected performance data to InfluxDB
	RAS Enrollment Server	TCP, UDP	30030	RAS Remote PC (PrisSCDriver) connects to get logon credentials
	FSlogix	TCP	443	Download FSlogix installer

Tenant Broker

Source	Destination	Protocols	Ports	Description
Tenant - RAS Publishing Agent	Tenant Broker - RAS Publishing Agent	TCP	20003	Tenant's RAS Publishing Agent communicates with Tenant Broker to join Tenant Broker, synchronize configuration and statuses

Active Directory and Domain Services ports

For Active Directory and Active Directory Domain Services port requirements, please see the following article: <https://technet.microsoft.com/en-us/library/dd772723%28v=ws.10%29.aspx>.

Azure Virtual Desktop

The Azure virtual machines you create for Azure Virtual Desktop must have access to the following URLs in the Azure commercial cloud:

Address	Outbound TCP port	Purpose	Service tag
*.wvd.microsoft.com	443	Service traffic	WindowsVirtualDesktop
gcs.prod.monitoring.core.windows.net	443	Agent traffic	AzureCloud
production.diagnostics.monitoring.core.windows.net	443	Agent traffic	AzureCloud
*xt.blob.core.windows.net	443	Agent traffic	AzureCloud
*eh.servicebus.windows.net	443	Agent traffic	AzureCloud
*xt.table.core.windows.net	443	Agent traffic	AzureCloud
*xt.queue.core.windows.net	443	Agent traffic	AzureCloud
catalogartifact.azureedge.net	443	Azure Marketplace	AzureCloud
kms.core.windows.net	1688	Windows activation	Internet
mrslobalsteus2prod.blob.core.windows.net	443	Agent and SXS stack updates	AzureCloud
wvdportalstorageblob.blob.core.windows.net	443	Azure portal support	AzureCloud
169.254.169.254	80	Azure Instance Metadata service endpoint	N/A
168.63.129.16	80	Session host health monitoring	N/A

The following table lists optional URLs that your Azure virtual machines can have access to:

Address	Outbound TCP port	Purpose	Azure Gov
*.microsoftonline.com	443	Authentication to Microsoft Online Services	login.microsoftonline.us
*.events.data.microsoft.com	443	Telemetry Service	None
www.msftconnecttest.com	443	Detects if the OS is connected to the internet	None
*.prod.do.dsp.mp.microsoft.com	443	Windows Update	None
login.windows.net	443	Sign in to Microsoft Online Services, Microsoft 365	login.microsoftonline.us
*.sfx.ms	443	Updates for OneDrive client software	oneclient.sfx.ms

*.digicert.com	443	Certificate revocation check	None
*.azure-dns.com	443	Azure DNS resolution	None
*.azure-dns.net	443	Azure DNS resolution	None

For up to date information, please also visit the Microsoft website at <https://docs.microsoft.com/en-us/azure/virtual-desktop/safe-url-list#required-url-check-tool>.

RAS performance counters

The following table lists performance counters available in Parallels RAS per component:

Parallels RAS Gateway (2XProxyGateway.exe)

ID	Name	Description
ras_gw_tot_conn	Total connections	The total number of Connections with the Gateway.
ras_gw_tot_threads	Total threads	The total number of threads running on the Gateway.
ras_gw_rpd_sess	RDP tunneled sessions	The number of tunneled RDP sessions.
ras_gw_rpd_sess_s	RDP SSL tunneled sessions	The number of tunneled RDP sessions over SSL.
ras_gw_html	HTTP connections	The number of tunneled HTTP sockets
ras_gw_html_s	HTTPS connections	The number of tunneled HTTPS sockets
ras_gw_html5	HTML5 connections	The number of tunneled HTTP5 sockets
ras_gw_html5_s	HTML5 SSL connections	The number of tunneled HTTP5 sockets over SSL
ras_gw_cm	Device Manager connections	The number of Parallels Device Manager connections
ras_gw_cm_s	Device Manager SSL connections	The number of Parallels Device Manager connections over SSL
ras_gw_wyse	Wyse connections	The number of Wyse connections
ras_gw_wyse_s	Wyse SSL connections	The number of Wyse connections over SSL
ras_gw_rdpudp	RDP UDP tunneled sessions	The number of RDP UDP connections
ras_gw_rdpudp_s	RDP UDP DTLS tunneled sessions	The number of RDP UDP connections over DTLS
ras_gw_cache_sock	Cached sockets	The number of cached sockets between Gateway and Publishing Agent
ras_gw_idle_threads	Idle threads	The number of idle threads on the Gateway
ras_gw_client	Client connections	The number of Parallels Client connections
ras_gw_client_s	Client SSL connections	The number of Parallels Client connections over SSL

Parallels RAS Publishing Agent (2XController.exe)

ID	Name	Description
ras_pa_avg_client_connection_time	Average time for client connection	The average client connection time.

ras_pa_avg_client_auth_time	Average time for user authentication	The average time taken to authenticate a user.
ras_pa_avg_client_policy_time	Average time to retrieve user policy	The average time taken to retrieve the user's policy.
ras_pa_avg_client_rep_time	Average time to send client telemetry	The average time taken to send client telemetry. Used by CEP.
ras_pa_avg_client_applist_time	Average time to retrieve user's published items	The average time taken to retrieve user's published items list.
ras_pa_avg_client_appicons_time	Average time to retrieve icons	The average time taken to retrieve published items icons.
ras_pa_avg_client_getidle_time	Average time to start up a request	The average time taken for the start up request.

Parallels RAS RDS Agent (2XAgent.exe)

ID	Name	Description
act_sess	Active RDS sessions	The number of active RDS Sessions.
disc_sess	Disconnected RDS sessions	The number of disconnected RDS Sessions.

Index

A

- About Parallels RAS - 15
- About Sites - 44
- About this guide - 16
- Active Directory and Domain Services ports - 481
- Active Directory user account configuration - 327
- Add a cloud VDI provider - 131
- Add a host pool - 195
- Add a hypervisor VDI provider - 129
- Add a new client policy - 387
- Add a VDI provider - 129
- Add a workspace - 194
- Add an RD Session Host - 30, 84
- Add Microsoft Azure as a VDI provider - 135
- Adding a HALB virtual server - 293
- Adding a RAS Secure Client Gateway - 66
- Adding a Remote PC to a Farm - 216
- Adding a Site to the Farm - 47
- Adding a VDI provider - 182
- Adding an administrator account - 50
- Adding and deleting pool members - 170
- Adding and deleting pools - 170
- Adding Remote PCs to a pool - 184
- Adding scanning applications - 372
- Additional properties - 154
- Admin-initiated Remote PC enrollment - 217
- Administrator account permissions - 51
- Advanced - 155
- Advanced settings - 400
- Agent settings - 91
- Allowing users to change domain password - 286
- Appendix - 469
- Architecture description - 300
- Assign a public domain address - 309
- Assigning a certificate to Gateways and HALB - 261
- Assigning a template to a pool or group - 163

- Attributes - 271
- Audio - 395
- Auditing certificates - 263
- Auto login - 359
- Automation - 272
- Azure Virtual Desktop - 41, 188, 482

B

- Branding - 348

C

- Change RD Session Host Site assignment - 90
- Change VDI provider Site assignment - 144
- Changing the HALB appliance password - 297
- Check an RD Session Host Agent status - 89
- Checking effective access - 253
- Checking the RAS Secure Client Gateway status - 67
- Checking the RAS VDI Agent status - 138
- Client Policies - 386
- Client policy backward compatibility - 407
- Client settings - 74
- Colors - 348
- Common Management Tasks - 429
- Communication ports - 317
- Computer management tools - 431
- Conclusion - 41
- Configure an SSL certificate - 310
- Configure certificate authority templates - 329
- Configure client policy options - 401
- Configure control settings - 405
- Configure CPU optimization - 290
- Configure file transfer for a client policy - 410
- Configure file transfer for a server - 409
- Configure file transfer for HTML5 Gateway - 410
- Configure Gateway redirection - 406
- Configure HTML5 Client - 345
- Configure HTML5 Gateway - 74

- Configure HTTP proxy settings - 440
- Configure logging - 81, 107, 142
- Configure managing existing profiles by Parallels RAS - 100
- Configure network - 309
- Configure RAS Console idle sessions - 55
- Configure RAS Performance Monitor Security - 426
- Configure RAS Web Administration Service - 456
- Configure session settings - 388
- Configure Themes - 346
- Configuring a RAS Secure Client Gateway - 67
- Configuring a Remote PC - 220
- Configuring a VDI provider - 183
- Configuring an RD Session Host - 89
- Configuring Azure MFA - 273
- Configuring DualShield 5.6+ Authentication Platform - 276
- Configuring Duo - 275
- Configuring MFA exceptions - 285
- Configuring notification handlers - 440
- Configuring notification scripts - 443
- Configuring notifications - 316
- Configuring Parallels RAS to use the DualShield Authentication Platform - 279
- Configuring preferred routing - 250
- Configuring RAS Publishing Agents - 57
- Configuring remote file transfer - 408
- Configuring SafeNet - 282
- Configuring SMTP server connection for event notifications - 446
- Connect to a RAS Farm - 281
- Connecting to a Parallels RAS Farm - 42
- Connection - 270, 389
- Connection and Authentication Settings - 265
- Create a smartcard logon certificate template - 333
- Create a template - 203
- Create an Enrollment Agent template - 329
- Create Microsoft Azure AD application - 132
- Creating a template - 150

D

- Delegating session management permissions - 352
- Deleting a Tenant object - 312

- Deploy Azure Virtual Desktop - 192
- Deploying a Parallels HALB appliance - 292
- Deploying a Tenant - 305
- Deploying Tenant Broker - 304
- Deploying Tenant Broker and Tenants - 303
- Desktop access - 95
- Direct App access - 360
- Display - 392
- Distribution - 154
- Drive redirection cache - 107

E

- Enable Azure Virtual Desktop and add a provider - 193
- Enable or disable a Gateway - 67
- Enable or disable HTML5 Client - 74
- Enabling Help Desk support - 375
- Enabling high availability for VDI - 142
- Error messages - 343
- Experience - 398
- Exporting a certificate - 261
- Exporting and importing Farm settings from the command line - 448

F

- Farm and Sites - 42
- Font management - 369
- FSLogix antivirus exclusions - 101
- FSLogix Profile Container - 97

G

- Gateway - 349
- Gateway mode and forwarding settings - 69
- Gateway network options - 69
- Gateway security - 77
- Gateway tunneling policies - 80
- GDPR compliance - 421
- General - 91
- General management tasks - 235
- General Theme settings - 346
- General Theme tasks - 351
- Generating a certificate signing request (CSR) - 260
- Generating a self-signed certificate - 259
- Getting additional device information - 377
- Getting started - 464
- Getting Started with Parallels RAS - 27

Grouping and cloning RD Session Hosts - 108

Guest VM naming - 159

H

HALB - 475

HALB connection and session information - 297

HALB Device status and version number - 296

HALB maintenance - 296

Hardware requirements - 19

Hiding toolbar items - 365

High availability load balancing (HALB) - 291

Host name resolution - 430

How guest VMs are created from a template - 161

HTML5 Client and Themes - 314

HTML5 Client Theme settings - 347

I

IdP side configuration - 324

Implementation overview - 301

Importing a certificate - 260

Input prompt - 349

Install Microsoft SQL Server - 413

Install Microsoft SQL Server 2016 or earlier - 413

Install Microsoft SQL Server 2017 or 2019 - 416

Install Parallels RAS - 23

Install Parallels RAS Reporting - 416

Install RAS Performance Monitor - 423

Installation - 455, 463

Installing Parallels RAS - 19

Installing RAS VDI Agent using the installer - 137

Installing the agent manually - 86

Introduction - 15, 188, 319

Introduction and prerequisites - 131

Invite users - 37

Inviting users to connect to Parallels RAS - 373

J

Join a Tenant to Tenant Broker - 305

Joining Customer Experience Program - 56

Joining with a secret key - 307

K

Keyboard - 396

L

Language bar - 349

Legal policies - 350

License keys - 158

Licensing - 438

Load Balancing and HALB - 288

Local devices and resources - 396

Log in and activate Parallels RAS - 24

Log in to RAS Management Portal - 456

Logging - 451

Logging in and sending requests - 465

M

Main menu options - 355

Maintaining RD Session Hosts based on a template - 115

Maintenance and backup - 448

Manage Azure Virtual Desktop - 197

Manage existing templates - 204

Manage folders - 243

Manage host pools - 198

Manage hosts - 205

Manage providers - 197

Manage published applications - 236

Manage published desktops - 240

Manage published documents - 241

Manage sessions - 207

Manage templates - 202

Manage workspaces - 198

Managing administrator accounts - 50, 53

Managing guest VMs - 172

Managing guest VMs in pools - 171

Managing Licensing Site - 49

Managing logons - 116

Managing multi-provider template distribution - 168

Managing RD Session Host sessions - 116

Managing Remote PCs in a pool - 185

Managing Secondary Publishing Agents - 61

Managing sessions - 231

Managing template-based guest VMs - 169

Managing Tenants - 311

Managing Universal Printing settings - 367

Managing Universal Scanning - 371

- Managing VDI sessions - 181
- Managing Windows devices - 379
- Manually adding a guest VM - 162
- Manually adding a RAS Secure Client Gateway - 66
- Mass configuring user devices - 374
- Messages - 349
- Microsoft Azure and templates - 136
- Microsoft license requirements - 23
- Microsoft license requirements in Parallels RAS - 469
- Modifying template properties - 161
- Modifying VDI Provider configuration - 138
- Monitoring devices - 376
- Monitoring settings - 231
- Monitoring Tenants - 315
- More information - 467
- Multi-factor authentication - 269
- Multi-provider template distribution - 148

N

- Native clipboard experience - 358
- Network - 399
- Network load balancers access - 76

O

- Open Parallels HTML5 Client - 353
- Opening a Tenant console - 312
- Optimization - 94, 102, 157
- Other useful features - 358
- Overview - 64, 216, 227, 234, 299, 422, 454

P

- Parallels Client - 474
- Parallels Client configuration - 341
- Parallels Client for Windows Theme settings - 350
- Parallels client policy configuration - 342
- Parallels HTML5 Client - 345
- Parallels RAS 18 release history - 15
- Parallels RAS APIs - 461
- Parallels RAS Management Portal - 454
- Parallels Test Template Wizard - 160
- Performance Monitor - 422
- Permissions - 464
- Permissions to manage certificates - 263
- Persistent guest VMs - 174
- Persistent Remote PCs - 186

- Planning for high availability - 87
- Policy information in Parallels Client - 407
- Port reference - 473
- Preparation - 156
- Prerequisites - 190, 323, 455
- Printing - 392
- Problem reporting and troubleshooting - 450
- Properties - 153
- Publish applications - 36
- Publish resources - 213
- Published Resources Management - 234
- Publishing a desktop - 118, 176, 223
- Publishing a document - 122, 180, 225
- Publishing a network folder - 121, 179, 225
- Publishing a web application - 121, 178, 224
- Publishing an application - 119, 177, 223
- Publishing App-V applications - 124
- Publishing containerized applications - 124
- Publishing from a guest VM - 176
- Publishing from a pool-based Remote PC - 187
- Publishing from a Remote PC - 223
- Publishing from an RD Session Host - 118

Q

- Quick keypad - 256

R

- RAS Console - 477
- RAS Enrollment Server - 480
- RAS Enrollment Server configuration - 338
- RAS Enrollment Server high availability - 340
- RAS Guest Agent - 480
- RAS Guest Agent installation options - 186
- RAS HTML5 Gateway API and Parallels Client URL scheme - 467
- RAS Management Portal user interface - 457
- RAS Multi-Tenant Architecture - 299
- RAS performance counters - 483
- RAS PowerShell - 479
- RAS PowerShell API - 461
- RAS Publishing Agent - 57, 476
- RAS Publishing Agent connection settings - 265
- RAS RD Session Host Agent - 480
- RAS Remote PC Agent - 481
- RAS Reporting - 478
- RAS REST API - 463

- RAS Secure Client Gateway - 64, 475
 - RAS session variables - 446
 - RAS VDI Agent - 479
 - RAS VDI Agent information - 127
 - RAS VDI Agent installation options - 128
 - RAS Web Administration Service (REST/Management Portal) - 478
 - RD Session Host drain mode examples - 115
 - RD Session Host templates - 148
 - RD Session Host types - 83
 - RD Session Hosts - 83
 - RDP printer - 95
 - Recovery - add a root administrator - 429
 - Remote PC pools - 181
 - Remote PCs - 216
 - Remote session settings - 266
 - Replicating Site settings - 48
 - Reporting - 411
 - Resource based & round robin load balancing - 288
 - Restricting access by Parallels Client type and build number - 268
 - Running Parallels RAS reports - 418
 - Running remote applications and desktops - 357
- S**
- SAML basics - 322
 - SAML configuration - 323
 - SAML integration examples and tips - 340
 - SAML SSO Authentication - 319
 - Scanning - 395
 - Scheduling Windows devices & groups power cycles - 385
 - Secondary Publishing Agents - 59
 - Security - 159
 - Security tip - 340
 - Self-service Remote PC enrollment - 218
 - Server authentication - 399
 - Session information - 228
 - Session Management - 227
 - Set IP addresses for client connections - 68
 - Set public address - 68
 - Set up a basic Parallels RAS Farm - 29
 - Set up routing for incoming traffic - 310
 - Settings - 158
 - Settings audit - 435
 - Shared Gateways - 312
 - Site defaults (Azure Virtual Desktop) - 208
 - Site defaults (Gateways) - 68
 - Site defaults (Publishing) - 245
 - Site defaults (VDI) - 144
 - Site defaults for multi-session hosts - 208
 - Site defaults for single-session hosts - 211
 - Site information - 433
 - Site settings - 433
 - Sites in the RAS Console - 45
 - Software requirements - 20
 - SP side configuration (RAS side) - 325
 - Specifying client settings - 254
 - SSL Certificate Management - 258
 - SSL server configuration - 73
 - SSL/TLS encryption - 70
 - SSRS - 478
 - Stage 1
 - Check and install the Agent - 151
 - Stage 2
 - Configure the template - 152
 - Suggest a feature - 453
 - Summary - 159
 - Supported tokens - 276
 - Supported VDI providers - 126
 - System event notifications - 440
 - System requirements - 19, 321, 411
- T**
- Template maintenance - 163
 - Template status - 166
 - Template types - 147
 - Templates - 147
 - Tenant Broker - 481
 - Tenant Broker compatibility and updates - 315
 - Tenant configuration - 311
 - Terms and abbreviations used in this guide - 17
 - Test the SAML SSO deployment - 342
 - The Parallels RAS Console - 27
 - The Resources tab - 233
 - Third-party network load balancers - 313
- U**
- Understanding session prelaunch - 252
 - Universal Printing - 367
 - Universal Printing drivers - 368
 - Universal Scanning - 371

- Unjoining from Tenant Broker - 310
- Updating RAS Performance Monitor - 427
- Upgrading from an older RAS version - 264, 316
- Upgrading RAS agents - 437
- URLs - 347
- User account attributes - 340
- User authentication - 310
- User connection flow - 302
- User Device Management and Client Policies - 373
- User profile - 94, 96, 152
- User Profile Disks - 97
- Using a VDI provider in multiple farms - 175
- Using a wildcard to filter VMs - 171
- Using computer management tools - 63, 82, 117, 176, 222
- Using Deepnet DualShield - 275
- Using default settings - 91
- Using drag and drop functionality - 357
- Using filtering rules - 247
- Using Google Authenticator - 283
- Using instant messaging - 55
- Using Parallels Client with Azure Virtual Desktop - 214
- Using Parallels RAS Performance Monitor - 423
- Using RADIUS - 269
- Using SafeNet - 282
- Using scheduler - 112, 207
- Using Site defaults - 74
- Using the remote clipboard - 364
- Using the toolbar - 361
- Using the toolbar on desktop computers - 361
- Using the Toolbar on Mobile Devices - 363

V

- VDI pool management - 170
- Verify join status - 308
- Verify the deployment - 213
- View and modify RD Session Host properties - 90
- Viewing Gateway summary and metrics - 81
- Viewing guest VMs on a VDI provider - 146
- Viewing published resources - 125
- Viewing RD Session Hosts - 87
- Viewing Remote PC summary - 222

- Viewing VDI provider summary - 181
- Virtual Desktop Infrastructure (VDI) - 126
- Virtual desktop templates - 147

W

- Web browsers - 474
- Web request load balancing - 78
- Windows desktop replacement - 383
- Windows device groups - 377
- Wyse ThinOS support - 76