



Parallels Remote Application Server

Scalability Testing with Login VSI

Contents

Introduction	3
Scalability	4
Testing the Scalability of Parallels RAS.....	4
Configurations for Scalability Testing.....	4
Login VSI Workload.....	5
Testing Process	5
Task Worker Workload Findings	6
Conclusion	12

CHAPTER 1

Introduction

Parallels RAS is a comprehensive virtual application and desktop delivery solution that allows your employees to use applications and data from any device. Seamless and easy to deploy, configure, and maintain, Parallels RAS supports the delivery of applications and desktops via Microsoft RDS and major hypervisors.

This document presents an analysis of the scalability testing of Parallels RAS v15.5 using Login VSI.

CHAPTER 2

Scalability

In This Chapter

Testing the Scalability of Parallels RAS	4
Configurations for Scalability Testing	4
Login VSI Workload	5
Testing Process	5
Task Worker Workload Findings	6

Testing the Scalability of Parallels RAS

To validate Parallels RAS configurations, Parallels engineers conducted a series of performance tests. The goal was to analyze the scalability of Parallels RAS sessions running on Hyper-V virtual machines. As a part of this testing, Login VSI was used in testing to generate user connections to RD Session Host servers, simulating typical user workloads.

In a Parallels RAS deployment, users connect through Parallels Client to access remote applications and desktops. Login VSI clients simulate user connections, while RAS Publishing Agents distribute them and set up service connections between end-users and RD Session Host servers.

Configurations for Scalability Testing

For the Parallels RAS scalability testing, Intel S4TR1SY3Q “Brickland” IVT-EX 4U-server consisting of the following hardware components was used:

CPU	4x Xeon E7-4890 v2 (D1 stepping) 2.8GHz 15 cores, 37.5MB L3, 155W TDP
RAM	256 GB, 32x 8 GB Micron DDR-4-2100 at 1600MHz
HDD	2x SEAGATE ST9300653SS 300GB/15000RPM SAS in RAID
SSD	2x Intel DC S3710 SSDSC2BA400G401 2.5" 400GB SATA III MLC

Parallels RAS was deployed on Hyper-V Server 2012 R2 as follows:

Parallels RAS Component	Total VMs	vCPU in each VM	RAM in each VM
RAS Publishing Agent	2	2	4 GB
RAS Secure Client Gateway	2	2	4 GB

High Availability Load Balancing	1	1	2 GB
RD Session Host server	9	4	16 GB

Login VSI Workload


Login VSI was configured as follows to simulate the workload on Parallels RAS:

- 250 Task Worker users.
- Launch 1 session every 10 seconds.
- Steady state: 2500 sec.

More info about the Login VSI Task Worker workload can be found at the following links:

- https://www.loginvsi.com/documentation/index.php?title=Changes_old_and_new_workloads
- <https://www.loginvsi.com/documents/documentation/Login-VSI-41-Workloads.pdf>

All virtual machines comprising the testing environment are siloed on the same virtual network.



Testing Process

In the scalability testing, Login VSI 4.1.12.8 was used to run a user load on Parallels RAS. Login VSI helps to gauge the maximum number of users that a desktop environment can support. Login VSI categorizes workloads as Task Worker, Knowledge Worker, Power Worker, and Office Worker.

The Task Worker workload was selected for the testing. The workload includes segments with Microsoft Office 2013 Outlook, Excel, Internet Explorer, Adobe Acrobat and PDF Writer. While being diverse and not focused on one or two applications, the Task Worker workload does not place a very severe demand on the environment and represents users that do not overload the system with heavy tasks.


It is important to note that while scalability testing is a key factor in understanding how a platform and overall solution perform, it should not be inferred as an exact measurement for real world production workloads. Customers looking to better assess how applications will perform should conduct their own Login VSI scale testing using custom workload scripts.

Since the goal was to capture a baseline reflecting the densities possible, the Login VSI client launchers were configured to go through Secure Gateway in proxy SSL mode. Performance metrics were captured during user logon and virtual desktop acquisition (ramp-up), user workload execution (steady state), and user logoff. To achieve consistent measurements that reflected when components were appropriately cached, each workload ran for 48 minutes before Login VSI performance metrics were recorded. VSI tests were repeated 3 times on each VM instance to get an average number of users that successfully ran the test.

Task Worker Workload Findings

Following are test results for the Task Worker workload. VSI_{max} v4 (which indicates the maximum user density under a specific workload) is determined from the VSI Baseline and VSI Threshold metrics. VSI Baseline represents a pre-test Login VSI baseline response time measurement that is determined before the normal Login VSI sessions are sampled.


A VSImax v4 density of 250 users running the Task Worker workload was demonstrated. In out tests, VSImax was not reached with 240 sessions.



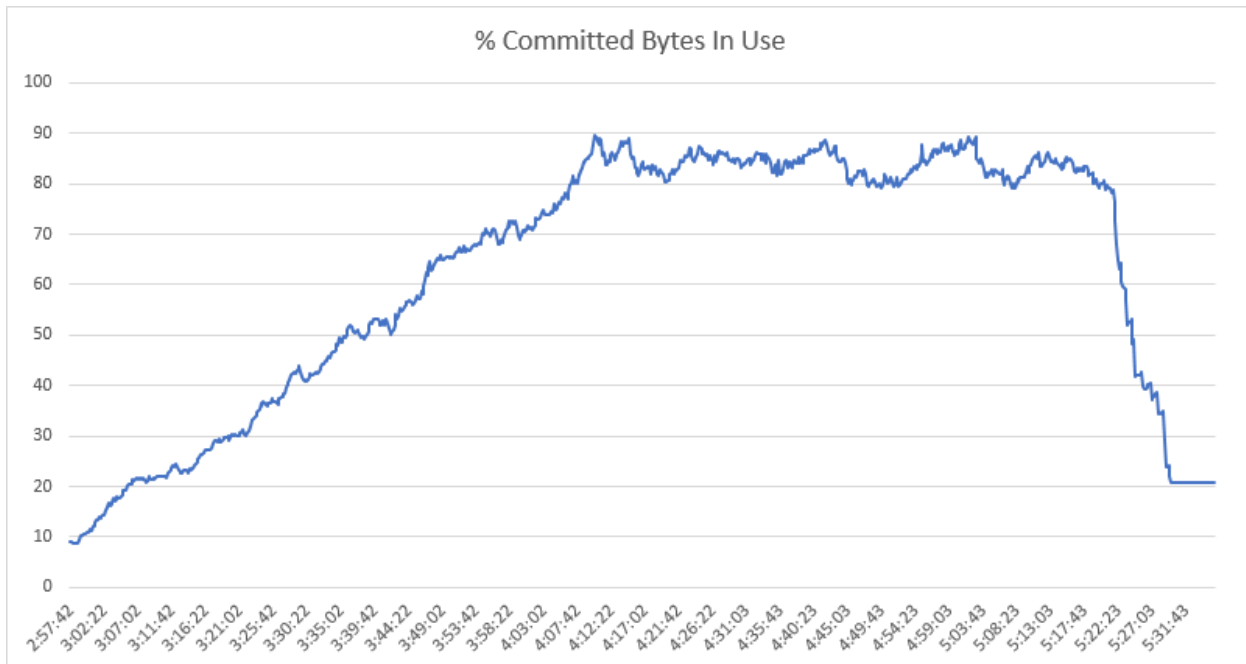
The following test results for CPU and memory consumption and disk I/O response times are helpful in evaluating performance under the test workload. Each chart below shows data collected from a single RD Session Host server. Since there were nine RD Session Host servers and 240 simulated users, a single RD Session Host served about 27 users.

In the following two charts, as user load increases, CPU and memory use peak where the number of users approaches VSImax v4.


In the following chart, the horizontal axis represents the time (h:mm:ss), while the vertical axis represents percentage of processor time. The orange line depicts the moving average.



In this chart, **% Committed Bytes In Use** means the amount of page file used.




The following chart shows average disk read and write response time. The write I/O response time averaged about 0.001287 sec. Read I/O response times averaged to 0.001395 sec.




The next two charts show resource consumption for network transfers and disk I/O.

The first chart shows networking transfer rates for data going out. For the Task Worker workload, the average outbound bandwidth at steady state is about 191 kB/s for our test group of 27 users (240 total users divided by nine RD Session Host servers). Therefore the outgoing transfer rate per user is about 6.82 kB/s.



The chart below shows disk transfer metrics. For the Task Worker workload, disk transfers during steady state averaged about 35 IOPS for the test group of 27 users, or about 1.25 IOPS per user. The peak transfer rate was 804 IOPS for 27 users, 29 IOPS per user. User profile data is recorded at logoff, generating disk transfer activity.



CHAPTER 3

Conclusion

The Parallels RAS scalability results presented here confirm that 240 Login VSI sessions using the Task Worker workload can be successfully launched using the given hardware configuration. Specifically, a total of nine RD Session Host servers with 4 vCPU and 16 GB of RAM each were sufficient to serve these sessions.

The results presented here should be used only as guidelines in configuring your Parallels RAS solution. Before making final sizing and deployment decisions, it is suggested that you run proof-of-concept tests using your own workloads.