HPE Server Automation 10.50

Linux Remediation Performance Characterization

White Paper

Contents

Summary of performance results

HPE tested the Remediate feature of Server Automation version 10.50 in the HPE Performance lab. The aim was to validate throughput and resource demand for a well-defined workload.
For the hardware configuration specified in **Appendix: Test system configuration**, the tested remediation job achieved the following results:

- 48.60 servers/minute, on a one-slice SA Core with 1000 RHEL 6.7 x86_64 managed servers
- 87.40 servers/minute, on a two-slice SA Core with 1000 RHEL 6.7 x86_64 managed servers

Test case description

The HPE Performance team used a software policy containing one .zip file of 100 MB to remediate a number of managed servers. The team measured the results for different Core configurations and load levels. All remediate jobs were submitted via the UAPI through Pytwist.

Configuration

- Core configurations of one and two slices
- Managed servers running RHEL 6.7 x86_64
- 1, 100, 200, 500 and 1000 levels of managed servers
- A software policy with one zip file of 100 MB for each tested level

Performance results

Load levels and throughput

Throughput is measured in number of managed servers processed per minute. This is computed by dividing the number of targets by the time required to complete the job.

Figure 1 shows the remediation throughput as the average of two iterations for each of the 1, 100, 200, 500 and 1000 servers load levels, on a one-slice and two-slice Core.

On average, the throughput for a two-slice Core with 1000 managed servers was 87.40 servers/minute. This was significantly higher than the 48.60 servers/minute average throughput for a one-slice Core in the same configuration.

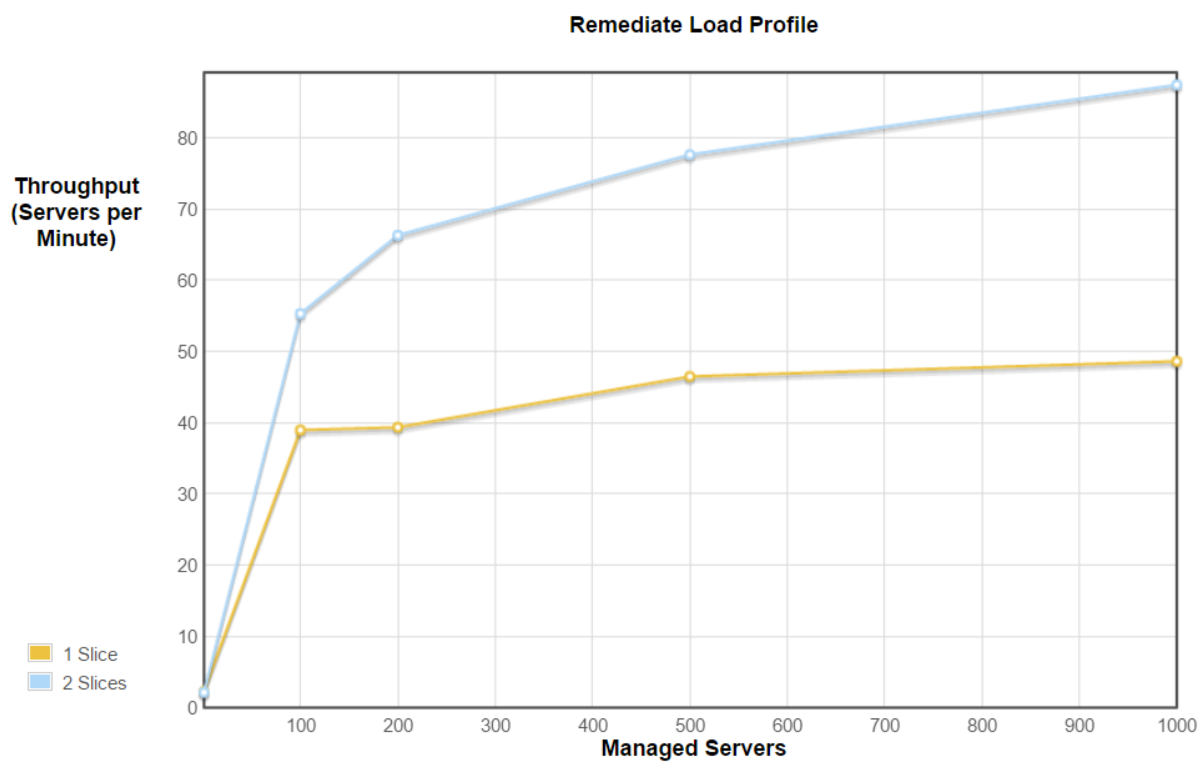

***Figure 1**: Overall job throughput, one-slice vs two-slice SA Core*

Scalability within the SA Core

The SA Core automatically distributes remediate workload across all SA slice servers within the Core. In this way, for large remediate job submissions, throughput can benefit from the horizontal scalability of the slice server. Overall throughput increases as the number of slice servers is increased.

The following table gives the horizontal scalability factor at the workload level of 1000 managed servers, as the number of slice servers is increased.

Table 1*: Linux remediate - effect of number of slices in SA Core on average throughput (servers per minute) and scalability factor at 1000 managed server load level*

# of slices	Throughput	Scalability factor
1	48.60	1
2	87.40	1.79

Workflow characterization in SA Core

Executing a remediation in SA requires varying resources. The following graphs show the workflow and resource demands on a one-slice vs. a two-slice SA Core processing 100 managed servers.

Workflow threads

The following graph shows the remediation operations across 100 managed servers on a one-slice Core. Each multicolored horizontal line represents a single managed server going through the remediation steps.
The X axis is the elapsed time, in seconds.
The remediation sessions on the managed servers are distributed across the two slices of the SA Core, as defined by the tunable parameters of the SA Command Engine.

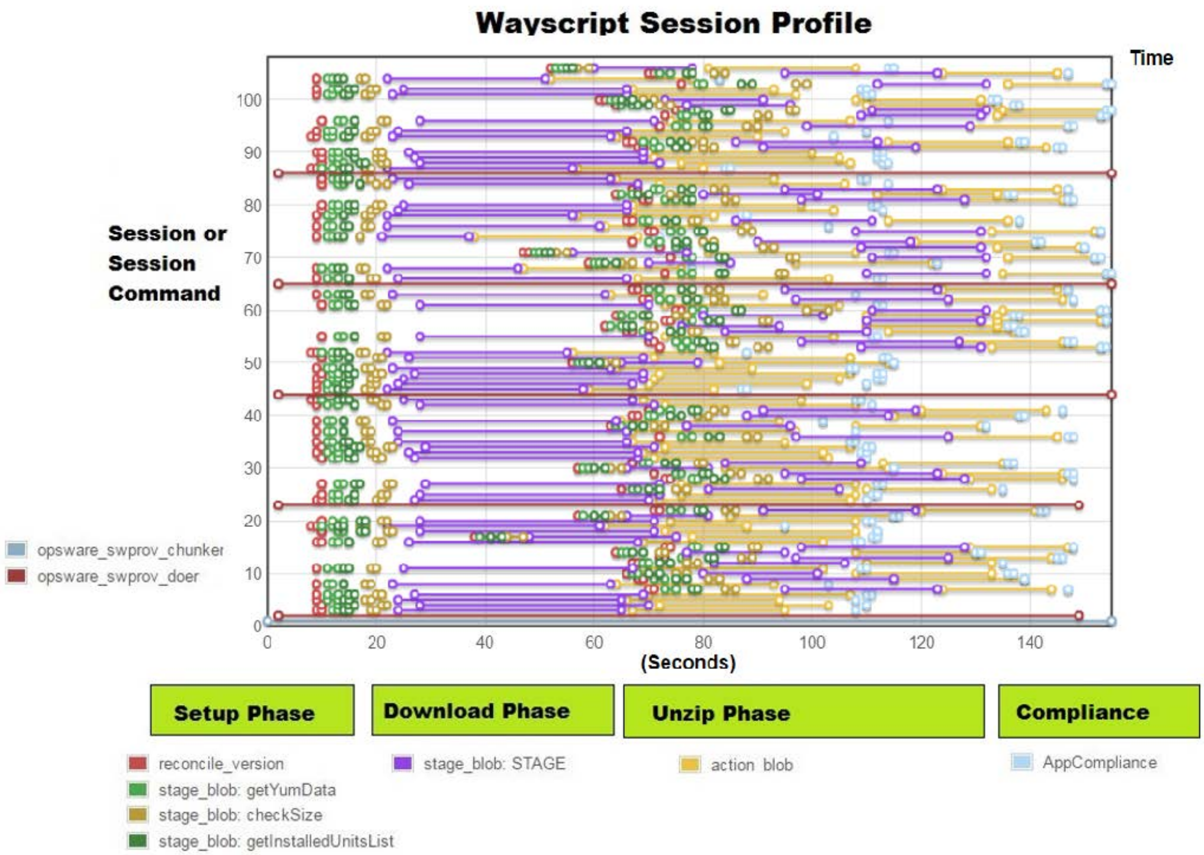

***Figure 2**: Work threads in 100 managed servers across a one-slice SA Core*

A remediate job consists of a single *chunker* session which spawns one or more *doer* sessions. Each doer operates on one or more managed servers. Several tunable parameters specify how the chunker session distributes and limits the load of the job across these doer sessions and across the SA Core as a whole.

In Figure 2 above, five doers operate on around 20 managed servers each. Job sessions with more managed servers would have more doer sessions to handle them. The job shows that each target server is managed by the first doer session below it, and that the chunker session at the bottom manages all doers.

If an SA Core has more than one slice, then the doers are distributed across the slices. Figure 3 shows how increasing the number of concurrent operations in progress for a two-slice Core results in a shorter job duration.

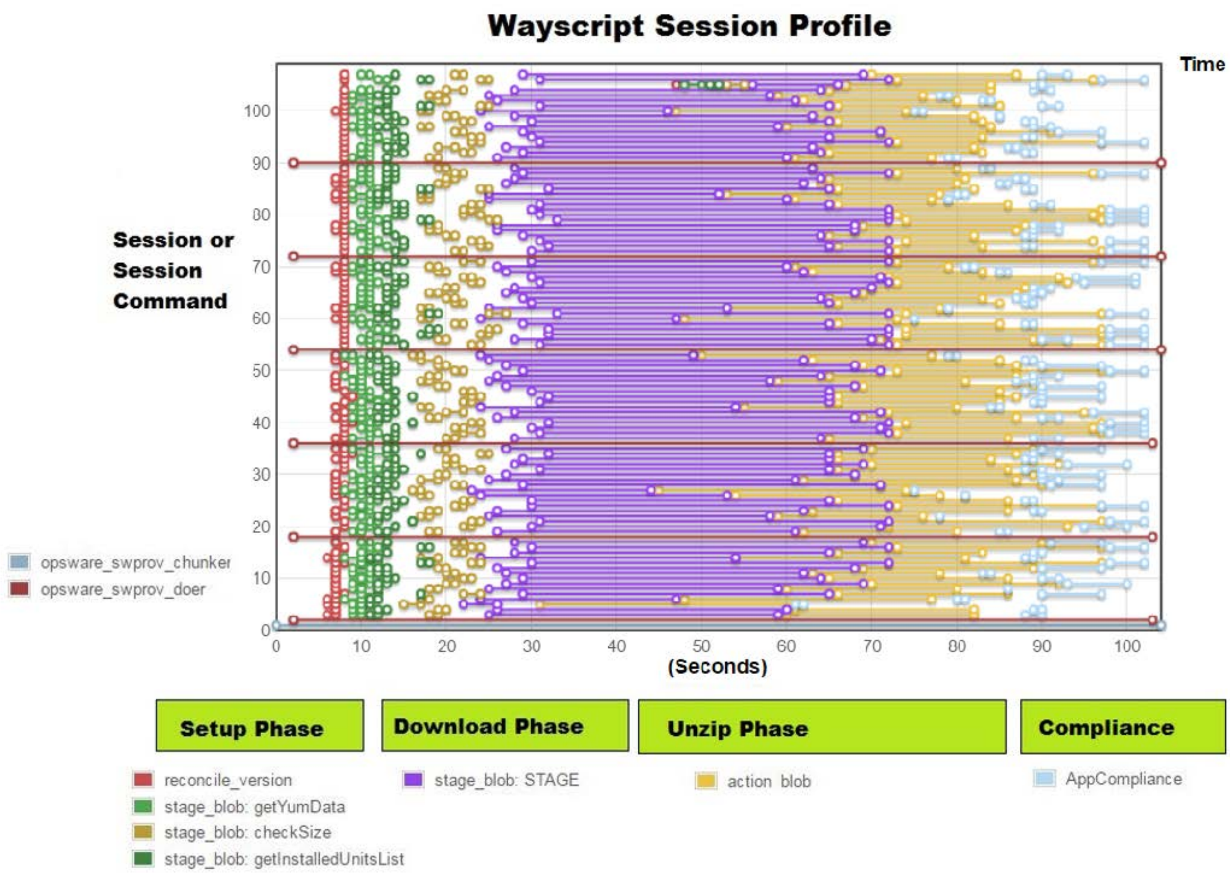

***Figure 3**: Work threads in 100 managed servers across a two-slice SA Core*

Resource utilization of SA slice servers

The following graphs show how much CPU and network bandwidth are required by the slice throughout the different steps of the remediation job.

CPU usage for an SA slice server in a job with 100 managed servers

Remediation exercises the slice's CPU and network resources. The Staging (Download) phase of the managed servers requires the most CPU and network usage.

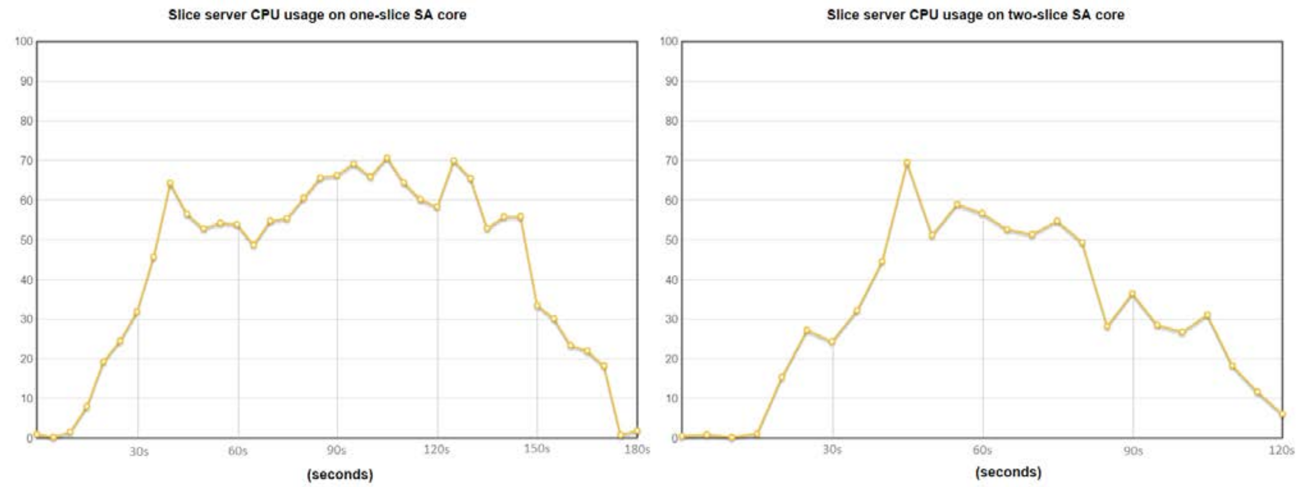

Figure 4*: CPU usage by a slice server when performing remediation with 100 managed servers on a one-slice vs two-slice SA configuration*

The first spike in the CPU usage graph below is related to the initial SA Core setup for each managed server. The initial setup involves:

1. Checking reconcile version.
2. Getting the list of installed software from the managed server. This avoids transferring an already existing item.
3. Checking the size of the files to be transferred to the managed server.

Figure 4 shows the CPU usage in a two-slice SA Core is lower than the one recorded in a one-slice configuration. The duration with peak CPU load is also considerably shorter for a two-slice SA Core.

CPU usage for an SA DB server in a job with 100 managed servers

On the database server of the SA Core, the CPU usage graph shows another peak towards the end of the job. This peak is aligned with the compliance phase of the job where the compliance information for the managed servers is updated in the DB.

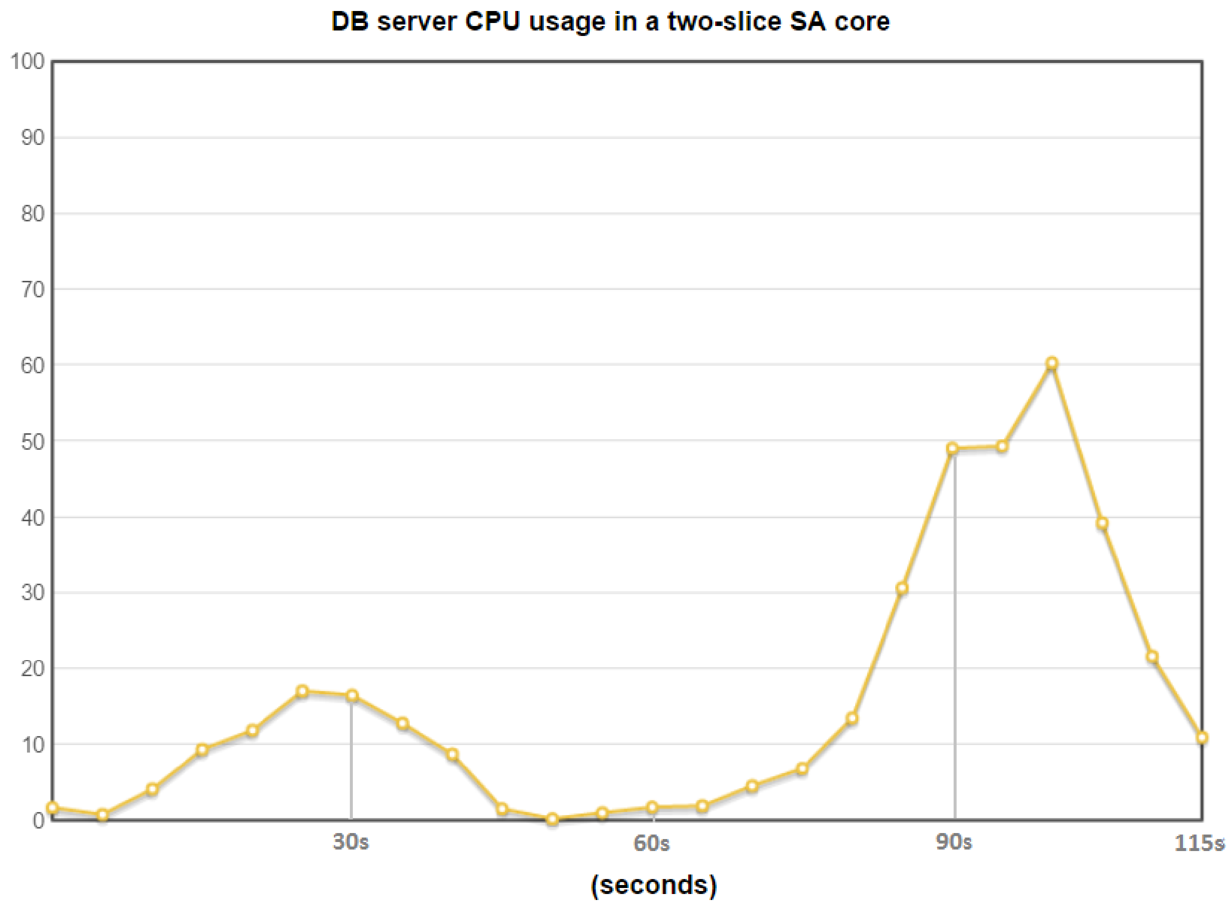

__Figure 5__: CPU usage DB server with 100 managed servers in a two-slice SA Core

Network usage for an SA slice server with 100 managed servers

Figure 6 shows the network usage on the slice server. The significant spike corresponds to the download phase of the remediation jobs.

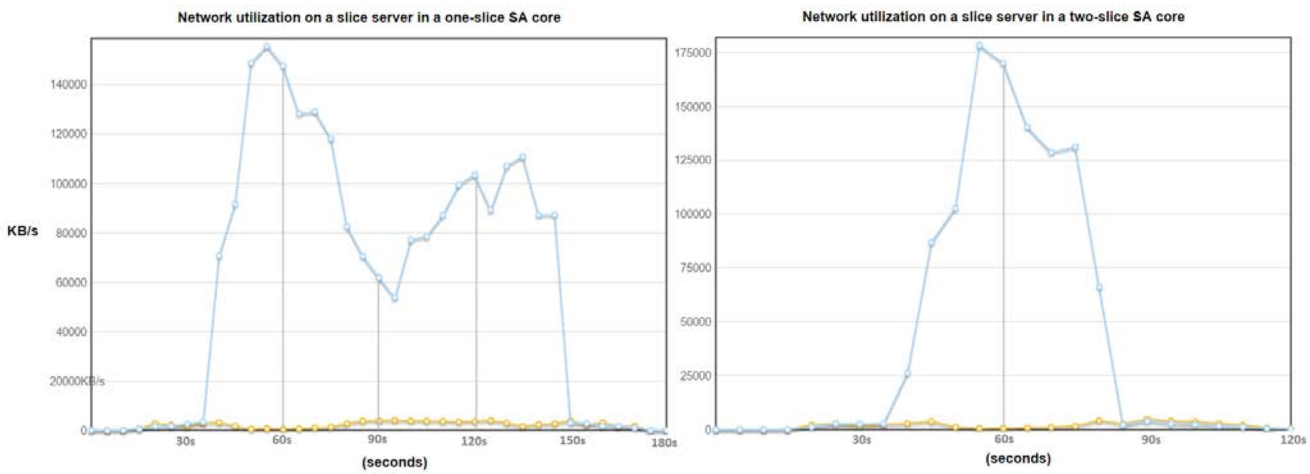

Figure 6*: Network usage by a slice server when performing remediation with 100 managed servers on a one-slice vs two-slice SA configuration*

Figure 6 also shows the network usage in a two-slice SA Core is slightly higher when compared to one-slice configuration, reaching a peak of 175MB/s, but of considerable shorter duration.

Tunable SA configuration parameters

The following system configuration parameters have been adjusted to facilitate stable operations:

Increase way.remediate.package_alarm_timeout

The remediate operation includes the action phase, which implies a workload on the managed server proportional with the payload being remediated. When the number of managed servers in a job is large enough, this operation may reach the default timeout value set to 3600 seconds. In this study, this value has been increased to 10200 seconds. The value of this parameter can be tuned from the SA Client[1].

Increase way.remediate.get_dicts_timeout

The *get_dicts_timeout* tuning parameter in SA is similar to *package_alarm_timeout*. It limits the number of seconds allowed for getting a list of installed software in the remediation action phase.

By default, this parameter is set to 1800 seconds but has been increased to 7200 seconds to work around timeout issues encountered at heavy workload levels. You can change the value of this parameter from the SA Client[2].

For this test, all other configuration parameters use the default values.

Conclusions

For concurrent operations on a number of managed servers, remediation achieves a steady state with a maximum throughput of **48.60 servers per minute** on a one-slice SA Core at a workload level of 1000 RHEL 6.7 64-bit managed servers. Adding an extra slice will greatly increase the throughput to a maximum of **87.40 servers per minute**.
This throughput is computed in the context of a well-defined 100 MB payload.

The resource utilization on both the slice and DB servers looking at CPU and network usage is moderate throughout the remediation job for the tested scenario.

[1] Administration View -> System Configuration / Configuration Parameters -> way.remediate.package_alarm_timeout
[2] Administration View -> System Configuration / Configuration Parameters -> way.remediate.get_dicts_timeout

Appendix: Test system configuration

SA Core servers

SA Core infrastructure	• Infrastructure & Slice services • Model Repository Multimaster Component (vault) • Data Access Engine (Spin - primary) • Gateways (mgw) • Media Repository (Word storage on NFS, SMB) • Model Repository Database (Truth)
ESXI host specifications	• ESXi 5.1 • HW: Model: HP ProLiant BL460c Gen9 • CPU: 16 CPUs x 2.6 GHz Intel Xeon E5-2640 • Memory: 256 GB
VM specifications	• Disk: 150 GB Linux ext4 • CPU: 8x vCPU @ 2.60 GHz , Memory: 32 GB
Network configuration	Network: 10 GBPS LAN, dedicated VLAN
Software specifications	• OS: RHEL6.7 64-bit • SA 10.50 (Build 65.0.70496.0)
SA Core slice #1 and #2	• “Slice” scalable services • Command Engine (Way) • Secondary Spin • Web service API (Twist) • Opsware Global File System (Hub) • Word • Tsunami • Gateways (cgw, agw)
ESXI host specifications	• ESXi 5.1 • HW: Model: HP ProLiant BL460c Gen9 • CPU: 16 CPUs x 2.6 GHz Intel Xeon E5-2640 • Memory: 256 GB
VM specifications	• Local Disk: 150 GB Linux ext4 • CPU: 8x vCPU @ 2.60 GHz , Memory: 32 GB
Network configuration	Network: 10 GBPS LAN, dedicated VLAN
Software specifications	• OS: RHEL 6.7 64-bit • SA 10.50 (Build 65.0.70496.0)

SA database	Oracle Database
ESXI host specifications	• ESXi 5.1 • HW: Model: HP ProLiant BL460c Gen9 • CPU: 16 CPUs x 2.6 GHz Intel Xeon E5-2640 • Memory: 256 GB
VM specifications	• Local Disk: 150 GB Linux ext4 • CPU: 8x vCPU @ 2.60 GHz , Memory: 32 GB
Network specifications	• Network: 10 GBPS LAN, dedicated VLAN
Software specifications	• OS: RHEL6.7 64-bit • Oracle Database 12c Standard Edition Release 12.1.0.2.0 – 64bit Production • SA 10.50 (Build 65.0.70496.0)

Managed servers

Managed servers	6.7 VMware VMs
ESXI host specs	• ESXi 5.1 • HW: Model: HP ProLiant BL460c Gen8 • CPU: 16 CPUs x 2.6 GHz Intel Xeon E5-2670 • Memory: 192 GB
VM specs	• Local Disk: 20 GB Linux ext4 • CPU: 1 vCPU @ 2.60 GHz , Memory: 2 GB
Network configuration	Network: 10 GBPS LAN, dedicated VLAN
Software specifications	OS: RHEL 6.7 64-bit
Additional notes	VMs are evenly distributed across 28 VMware ESXi hosts

Send documentation feedback

If you have comments about this document, you can contact the documentation team by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

Feedback on Linux Remediation Performance Characterization White Paper (Server Automation 10.50)

Just add your feedback to the email and click send.
If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to hpe_sa_docs@hpe.com.

We appreciate your feedback!

© Copyright 2017 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

This document contains confidential and/or legally privileged information. It is intended for Hewlett Packard Enterprise and Channel Partner Internal Use only. If you are not an intended recipient as identified on the front cover of this document, you are strictly prohibited from reviewing, redistributing, disseminating, or in any other way using or relying on the contents of this document.

February 2017

Hewlett Packard Enterprise