HPE Server Automation 10.50

Audit Performance

White Paper

Contents

Summary of performance results

HPE tested the Audit feature of Server Automation version 10.50 in the HPE Performance lab. The aim was to asses overall throughput, scalability and resource demand for a well-defined workload.
For the hardware configuration specified in **Appendix: Test system configuration**, the tested audit generated the following results:

Linux

- 25.50 servers/minute, on a one-slice SA Core with 1000 managed servers
- 52.56 servers/minute, on a two-slice SA Core with 1000 managed servers

Windows

- 17.78 servers/minute, on a one-slice SA Core with 1000 managed servers
- 33.37 servers/minute, on a two-slice SA Core with 1000 managed servers

Test case description

The test case used the Payment Card Industry Data Security Standard (PCI DSS) compliance test suite. The HPE Performance team ran this test flow separately on managed servers running on the following platforms:

- Linux: Red Hat Enterprise Linux 6.7
- Windows: Windows 2012 R2

The test ran over 190 audit/compliance checks (of files, data files, uninstall files, etc.) against a source server and a set of target servers.
Core configurations consisting of one and two slices were tested with managed server load levels of 1 to 1000.

Configuration

1. The following levels of audited managed servers were used: 1, 100, 200, 500 and 1000 target servers.
2. Audit job on Linux targets contained 192 separate checks for each managed server.
3. Audit job on Windows targets contained 193 separate checks for each managed server.
4. Audit tested in the HPE Performance Center of Excellence environment. This uses virtual machines for the managed servers that run the SA agents.
5. The audit jobs were submitted via the UAPI through Pytwist.

Performance results

Each performance test consisted of a job programmatically submitted to SA. The result of the audit shows the compliance status for all checks and target servers.

Load levels and throughput

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

Linux

For Linux targets, the highest throughput achieved was 54.10 servers/minute in a two-slice SA Core at a workload of 1000 managed servers.
On average, the throughput in a two-slice Core with 1000 managed servers was 52.56 servers/minute. This was significantly higher than the 25.50 servers/minute average throughput for a one-slice Core in the same configuration.

***Figure 1**: Linux PCI audit - server throughput*

Windows

For Windows targets, the highest throughput achieved was 33.41 servers per minute in a two-slice SA Core at a workload of 1000 managed servers. On average, the throughput in a two-slice Core with 1000 managed servers was 33.37 servers per minute.
This was significantly higher than the 17.78 servers per minute average throughput for a one-slice Core in the same configuration.

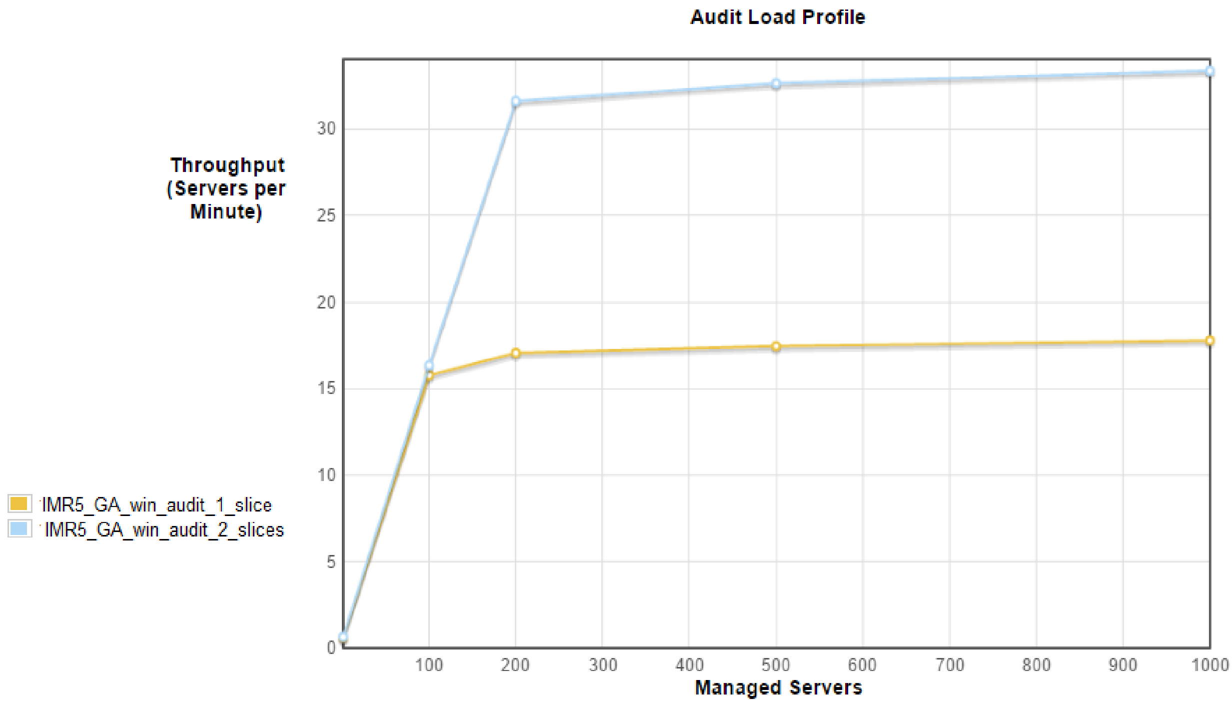

***Figure 2**: Windows PCI Audit - server throughput*

Scalability within SA Core

The SA Core automatically distributes audit workload across all SA slice servers within the Core. In this way, for large audit 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 tables give the horizontal scalability factor at the workload level of 1000 managed servers, as the number of slice servers is increased.

***Table 1**: Linux audit - 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	25.50	1
2	52.56	2.06

***Table 2**: Windows audit - 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	17.78	1
2	33.37	1.87

Workflow threads

The following graph shows operations on 20 managed servers being executed simultaneously in the system. In the initial job submission, the first 20 audit operations start concurrently across 20 managed servers. During later stages of the job execution, work thread processing shows a smooth distribution over time.

This shows that SA Core operation is predictable and stable at higher submission levels.

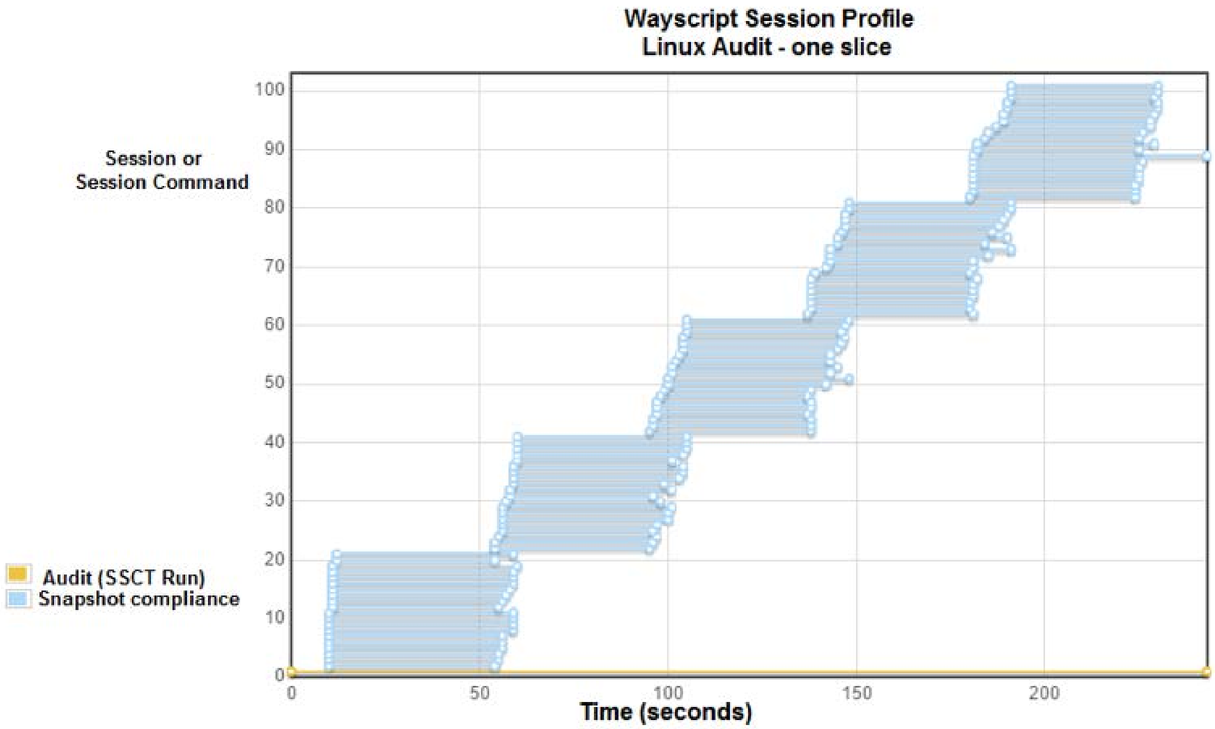

***Figure 3**: Work threads during audit on 100 audit managed server in a one-slice Core*

As additional slices are added, SA distributes the audit snapshot compliance tasks across the SA slices. Since, by default, the concurrency is limited to 20 operations per slice, the number of concurrent operations in progress increases with additional slices.

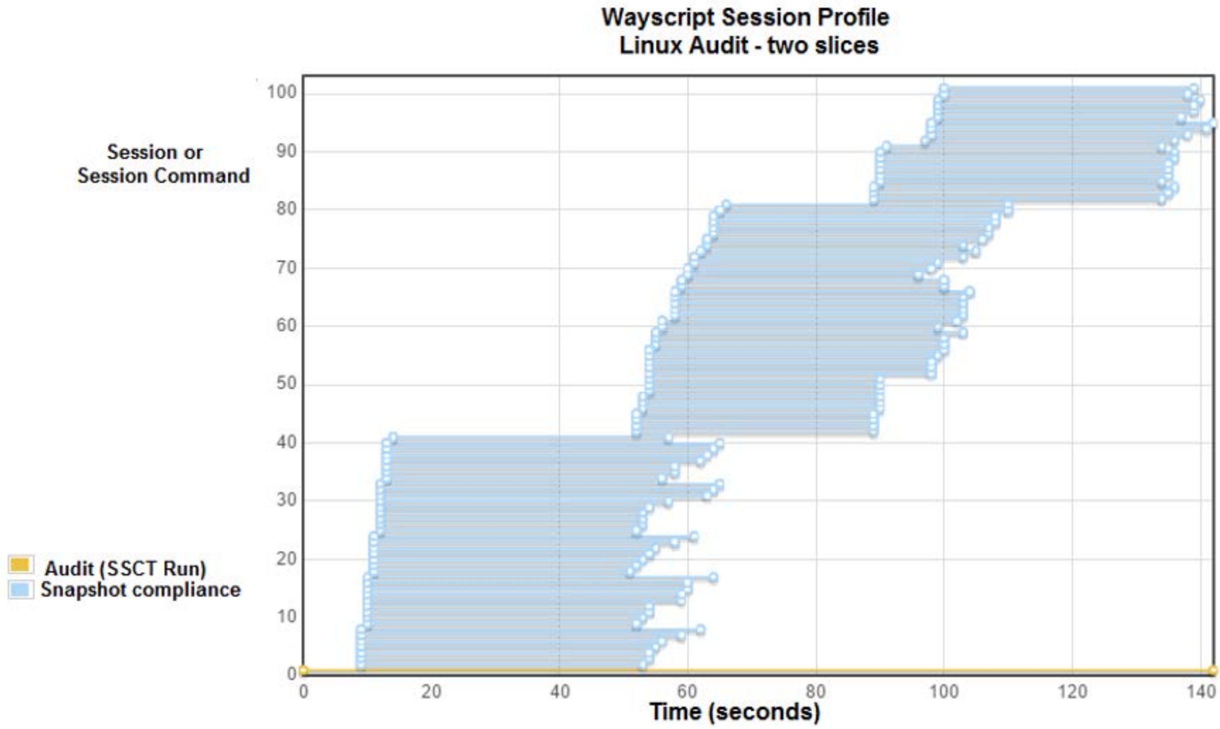

***Figure 4**: Work threads during audit on 100 managed servers for two-slice Core*

Workload characterization in SA Core

Auditing requires varying system resources over a job's lifetime.

The following series of graphs show the workflow and resource demands for an audit submitted on 100 managed server in a one-slice SA Core. This is the same job as shown in Figure 3.

SA Truth database system utilization

CPU load and network load on the Truth database server stays low during most of the auditing job.

As several of managed servers complete an audit at the same time, the CPU and network requirements spike moderately.

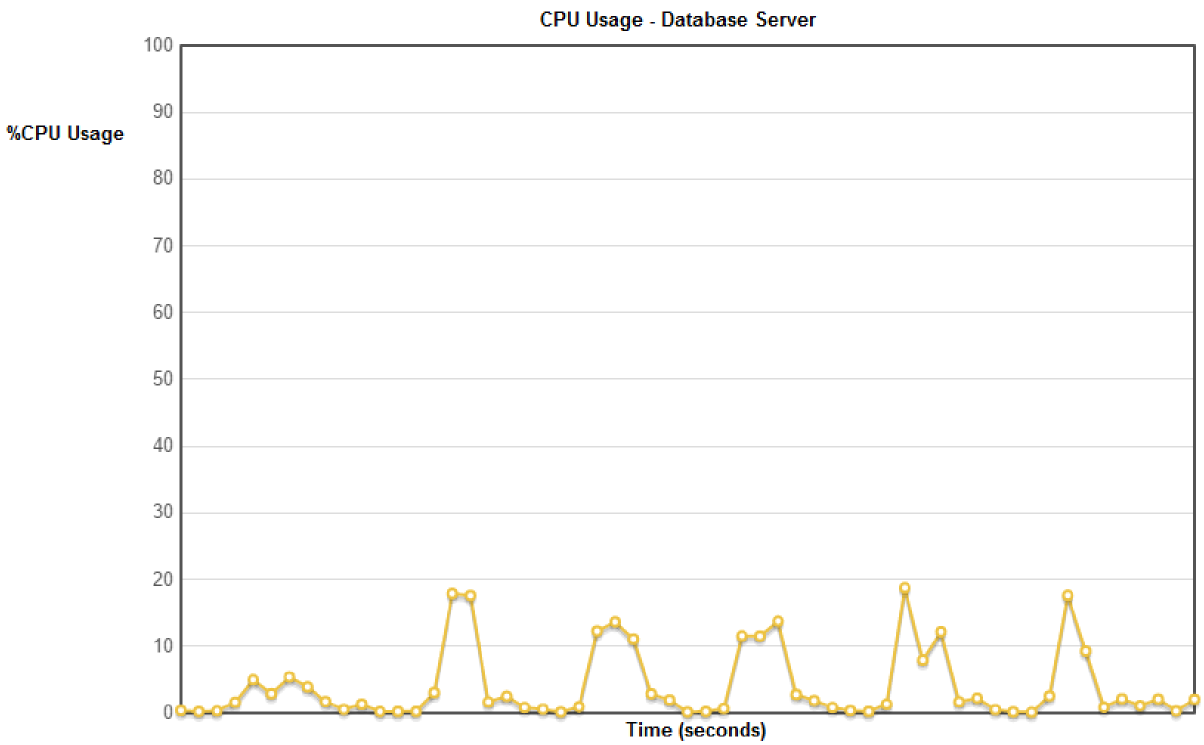

Figure 5: *Truth database CPU utilization during an audit job on 100 managed server in a one-slice Core*

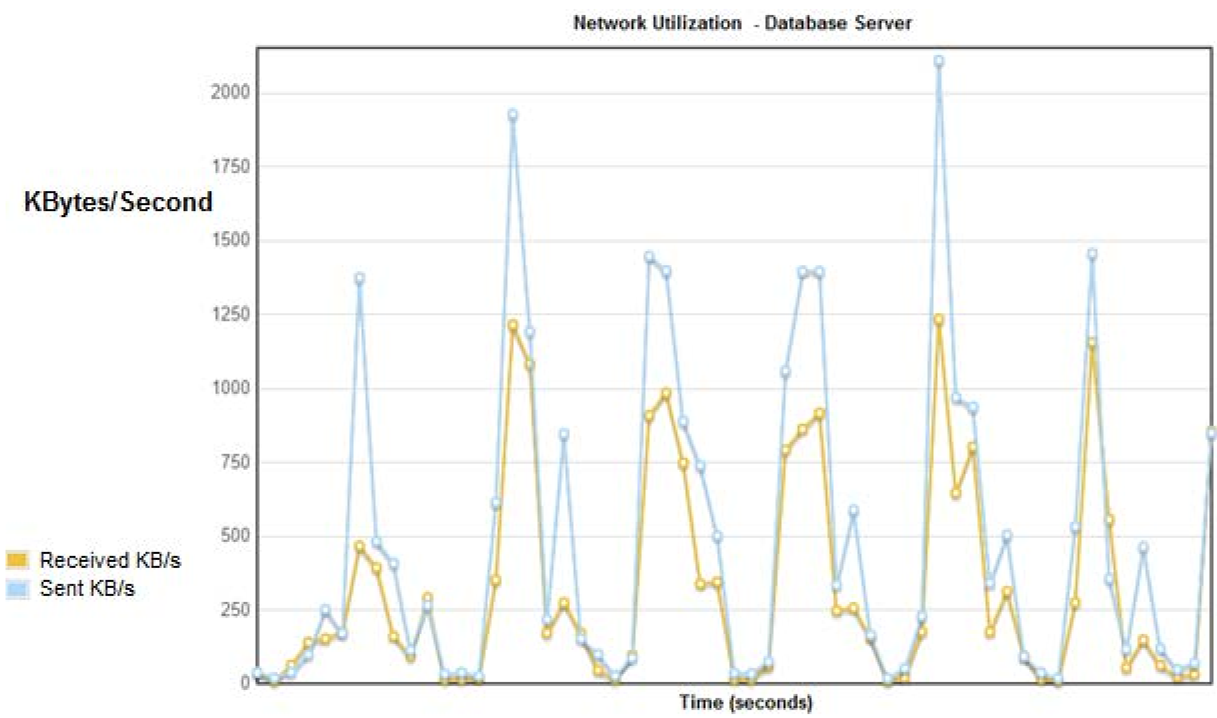

Figure 6: *Truth database network utilization during an audit job on 100 managed server in a one-slice Core*

SA slice system utilization

The server that hosts the SA slice shows periods of heavy CPU load during the snapshot compliance phase of the audit job.
The graph in **Figure 7** shows the total CPU load for all CPUs on the slice server belonging to a single-slice Core. The peaks correspond to the Job Setup phase, which happens on the SA slice. The load drops during the actual audit.

As highlighted in **Figure 9**, on 100 managed servers, the processing load for the same audit job is split between the available Core slices. When adding a second slice to a Core, the CPU demand reduces by 50%. As a result, the total job duration is reduced by 40% in a two-slice Core compared to a single-slice Core.

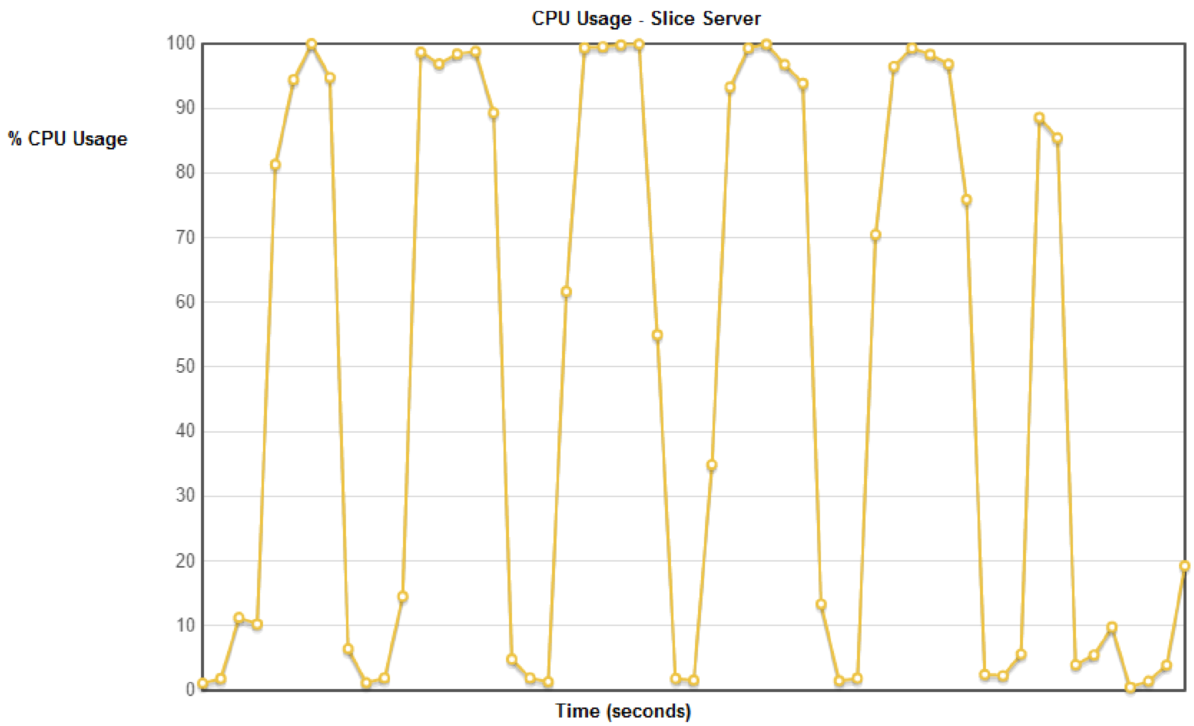

Figure 7: *Slice server CPU demand during an audit job on 100 managed server in a one-slice Core*

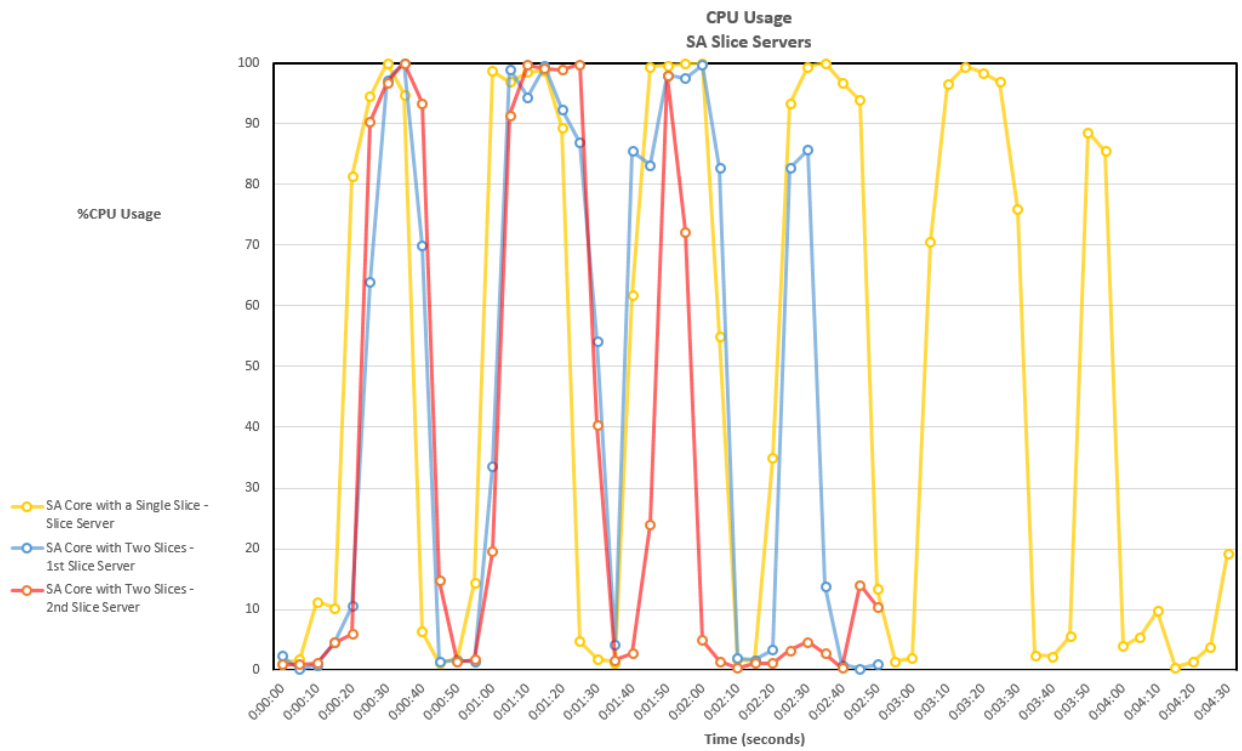

Figure 9: *Slice server CPU demand during an audit job on 100 managed server for one-slice and two-slice Cores*

The graph in **Figure 10** shows the network utilization on a slice server belonging to a Core configured with a single slice. The visible network load peaks correspond to the completion of the audit task for multiple server at the same time.

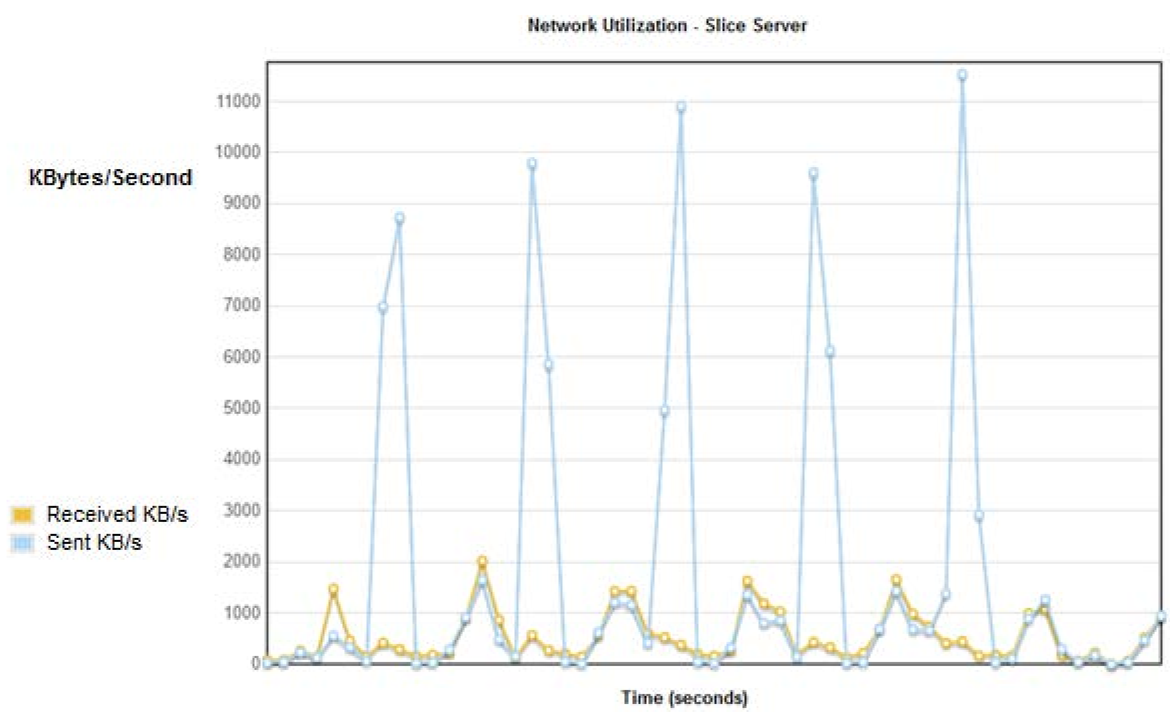

Figure 10: Slice server network utilization during an audit job on 100 managed server in a one-slice Core

Tuning parameters and work concurrency

The parameter **way.max_ssct_waiter** in SA can be changed to increase concurrency of an audit job. By default, this is set to 20 concurrent operations per slice. You can tune this in the SA Client. This parameter only affects concurrency of a single audit job. If you run multiple audits in parallel, they each have a maximum concurrency of 20 operations per slice. This means that the number of operations occurring on each slice will increase with the number of audit jobs you run in parallel.

Conclusions

For the tested configurations, the one-slice audit steady state throughput reached 17.78 servers/minute for Windows targets and 25.50 servers/minute for Linux targets. Adding additional slices increases the throughput levels through horizontal scalability.

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 configuration	• 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 specifications	• ESXi 5.1 • HW: Model: HP ProLiant BL460c Gen8 • CPU: 16 CPUs x 2.6 GHz Intel Xeon E5-2670 • Memory: 192 GB
VM specifications	• 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

Managed servers	Windows 2012 R2 VMware VMs
ESXi host specifications	• ESXi 5.1 • HW: Model: HP ProLiant BL460c Gen8 • CPU: 16 CPUs x 2.6 GHz Intel Xeon E5-2670 • Memory: 192 GB
VM specifications	• Local Disk: 40 GB • CPU: 2 vCPU @ 2.60 GHz , Memory: 4 GB
Network configuration	Network: 10 GBPS LAN, dedicated VLAN
Software specifications	OS: Windows 2012 R2 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 Audit Performance 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!

February 2017