



Hewlett Packard
Enterprise

Data Center Automation Suite 2016.10

Linux Patch Management
Performance Characterization

White Paper

Contents

Summary of performance results	1
Test case description	1
Performance results	2
Overall job throughput	2
Resource utilization on DCA Virtual Appliance.....	3
Tunable configuration parameters	5
Conclusions	5
Appendix: Test system configuration	6
DCA Server.....	6
Managed servers	6
Send documentation feedback	7

Summary of performance results

HPE tested the Linux Patch Management functionality of Data Center Automation Suite version 2016.10 in the HPE Performance lab. The aim was to validate the overall throughput and resource demand for a well-defined workload.

The DCA Suite 2016.10 is a virtual appliance that provisions and manages servers and data center operations more effectively.

In this study, the HPE Performance team exercised the capabilities of DCA Operations Portal, Operations Orchestration and Server Automation in a patching scenario for Linux-based managed servers. The throughput data was measured at DCA Operations Portal level. This white paper details the analysis of the DCA Linux Patching performance capabilities.

DCA Operations Portal offers the ability to perform patching operations on Linux-managed servers by leveraging the Software Policy remediation feature in Server Automation. Server Automation's Red Hat patching feature uses the Linux Yum utility for installing patches on managed servers.

Test case description

The HPE Performance team selected the following test based on typical customer usage of the feature **Patching RHEL servers with a monthly erratum set including its dependencies**.

The implemented RHEL Patch management test case:

- patches simultaneously up to 1000 RHEL 6.7 managed servers using the erratum set published in the first month after the General Availability release of the RHEL minor version. The erratum set is internally represented as a Software Policy.
- uses a Server Automation managed server device group with the Software Policy attached.
- uses a payload consisting of 29 errata. The errata contains 128 RPM patch packages out of which 32 (157 MB in size) are updated when the job runs.

All remediate jobs were submitted via the DCA UAPI.

For more information about the hardware and software configuration of the environment under test, see **Appendix: Test system configuration**.

Performance results

Job throughput is measured in number of managed servers processed per minute. This is computed by dividing the number of managed servers on which the job is run by the total execution time of the job.

Overall job throughput

The overall job throughput represents an average of two iterations for each load level: 1, 100, 200, 500 and 1000 managed servers.

The highest achieved throughput is 18.18 servers/minute at a workload of 1000 managed servers.




Figure 1: Linux patch remediation job throughput (managed servers per minute)

Resource utilization on DCA Virtual Appliance

Remediation in DCA is characterized by a balanced resource demand on the DCA Virtual Appliance, where the DCA Operations Portal, Operations Orchestration and Server Automation applications share hardware resources managed by the same operating system.

The following graphics show the CPU and network usage for a DCA server while processing a patch remediation job on 100 managed servers.

For a detailed analysis on the performance of a patch remediation job at Server Automation level, see the *HPE Server Automation 10.50 Linux Patch Management Performance Characterization* whitepaper.

CPU usage on a DCA server running a patch remediation job for 100 managed servers

As highlighted in Figure 2, a patch remediation flow in DCA starts with moderate CPU usage. This corresponds to the job preparation stage.

The first phases of the remediation job, analysis and the staging (download) phase are characterized by moderate to high CPU usage.

CPU usage decreases on the DCA server while the target managed servers process the workload. CPU usage spikes during the verification and compliance stages.




Figure 2: CPU usage for a DCA server while performing a patch remediation job on 100 managed servers

Network utilization on a DCA server running a patch remediation job for 100 managed servers

Figure 3 shows the network pattern during a patch remediation job. Utilization spikes correspond to the download phase of the remediation job, reaching up to 150 MB/s in tested environment.

Payload is downloaded in batches on target managed servers, ranging throughout the second quarter of the total duration of the job.




Figure 3: Network utilization for a DCA server while performing a patch remediation job on 100 managed servers

Tunable configuration parameters

The HPE Performance team adjusted the following system configuration parameters to facilitate stable operations:

Increase SA Device Remediation timeout at DCA level

DCA Administration Portal enables you to configure the SA Remediation timeout parameter¹.

In requests targeting a large number of managed servers, the remediation operation may reach the default timeout value of 3600 seconds. For this study, the HPE Performance team increased this value to 18000 seconds.

Increase Server Automation way.remediate.package_alarm_timeout

The remediate operation includes the action phase, which implies a workload on the managed server proportional to the payload being remediated.

When the number of managed servers in a job is large enough, this operation may reach the default timeout value of 3600 seconds. In this study, this value has been increased to 10200 seconds. You can adjust the value of this parameter from the SA Client².

Increase Server Automation 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 Client³.

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

Conclusions

For the given test configuration, DCA patch remediation performs concurrent operations achieving a steady state with maximum throughput of **18.18** servers per minute at a workload level of 1000 RHEL 6.7 64-bit managed servers.

This throughput is computed in the context of a well-defined payload: all RHEL-published errata during the first month after the official release of RHEL 6.7.

Different representative patch sets would exhibit correspondingly different characteristics.

Resource utilization on the DCA server is characterized by medium to high CPU and network usage throughout the remediation job.

¹ Designs > Sequenced > Resource Offerings > Managed Servers > Lifecycle > Provisioning Stages / Deploying > During > SA Device Remediation

² Administration View > System Configuration / Configuration Parameters > way.remediate.package_alarm_timeout

³ Administration View > System Configuration / Configuration Parameters > way.remediate.get_dicts_timeout

Appendix: Test system configuration

DCA Server

		<ul style="list-style-type: none"> • ESXi 5.1
ESXi host specifications		<ul style="list-style-type: none"> • HW: Model: HP ProLiant BL460c Gen9 • CPU: 16 CPUs x 2.6 GHz Intel Xeon E5-2640 • Memory: 256 GB
VM specifications		<ul style="list-style-type: none"> • Disk: 256 GB Linux ext4 • CPU: 8x vCPU @ 2.60GHz , Memory: 32 GB
Network configuration		Network: 10 GBPS LAN, dedicated VLAN
Software specifications		<ul style="list-style-type: none"> • OS: CentOS 6.5 64-bit • DCA Suite 2016.10

Managed servers

Managed servers		RHEL 6.7 VMware VMs
ESXi host specifications		<ul style="list-style-type: none"> • ESXi 5.1 • Hardware: <ul style="list-style-type: none"> ◦ Model: HP ProLiant BL460c Gen8 ◦ CPU: 16 CPUs x 2.6 GHz Intel Xeon E5-2670 ◦ Memory: 192 GB
VM specifications		<ul style="list-style-type: none"> • Local Disk: 20 GB Linux ext4 • CPU: 1 vCPU @ 2.60GHz • 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 Patching Performance White Paper (Data Center Automation Suite 2016.10)

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!