VMware vCenter Update Manager Performance and Best Practices

VMware vCenter Update Manager 4.0

VMware vCenter Update Manager provides a patch management framework for VMware vSphere. IT administrators can use it to patch and upgrade VMware ESX/ESXi hosts, apply patches to Windows and certain versions of Linux in virtual machines, upgrade VMware Tools and virtual hardware for virtual machines, and patch and upgrade virtual appliances.

As data centers get bigger, performance implications become more important for patch management. This study covers the following topics:

Benchmarking Methodology

Experimental Setup

VMware Update manager 4.0, VMware vCenter 4.0, ESX 3.5, and ESX 4.0 were used for performance measurements. WANem 1.2 was used for simulating a high-latency network with packet drops. Microsoft Windows XP SP2 was used for powered-off virtual machine scan and remediation. Red Hat Enterprise Linux 32-bit was used for Linux virtual machine scan.

VMware Update Manager and vCenter Server

Host Computer: Dell PowerEdge 2970

CPUs: Two 2GHz AMD Opteron 2212 dual-core processors

RAM: 16GB

Hard drives: Eight 73GB SAS drives

Network: Broadcom NetXtreme II5708 1Gbps

Update Manager Server software: Update Manager 4.0

vCenter Server software: VMware vCenter 4.0

ESX System

Host Computer: Dell PowerEdge 2900

CPUs: Two 2.66GHz Intel Xeon 5355 quad-core processors

RAM: 32GB

Hard drives: Eight 73GB SAS drives

Network: Broadcom NetXtreme II5708 1Gbps

ESX software: VMware ESX 3.5 and ESX 4.0

Virtual Machine Operating Systems

Windows: Microsoft Windows XP SP2

Linux: Red Hat Enterprise Linux 32-bit (kernel: 2.6.18-8.el5)

Network Simulation Software

WANem 1.2

Network Configurations

The network configurations used in the experiments are shown in Figure 1 (basic network configuration) and Figure 2 (high-latency network configuration).

Figure 1. Basic Network Configuration

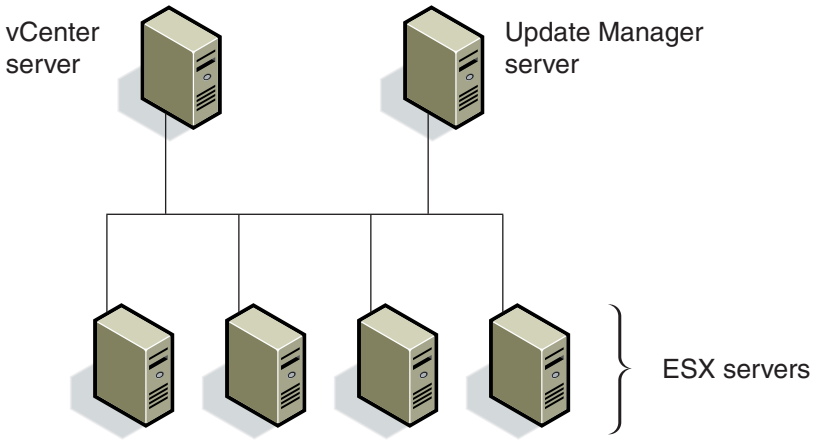

Figure 2. High Latency Network Configuration

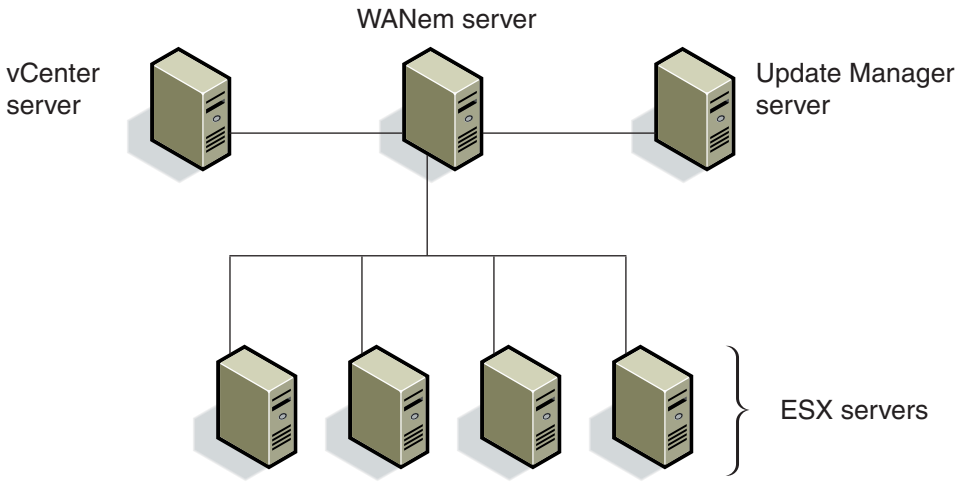

Update Manager Server Host Deployment

The three Update Manager server host deployment models, shown in Figure 3, are:

- Model 1 — The vCenter server and the Update Manager server share both a host and a database server.
- Model 2 — Recommended for data centers with more than 300 virtual machines or 30 ESX hosts. In this model, the vCenter server and the Update Manager server still share a host, but use separate database servers.
- Model 3 — Recommended for data centers with more than 1,000 virtual machines or 100 ESX hosts. In this model, the vCenter server and the Update Manager server run on different hosts, each with its own database server.

For both ESX and virtual machine patching the Update Manager server transfers patch files over the network. To avoid unnecessary disk I/O it is ideal if the Update Manager server host can cache patch files, some of which are several hundred megabytes, completely within the system cache. To this end it is desirable for the Update Manager server host to have at least 2GB of RAM.

It is also best to place the patch store and Update Manager database on separate physical disks. This arrangement distributes the Update Manager I/O and dramatically improves performance.

Figure 3. Update Manager Server Host Deployment Models

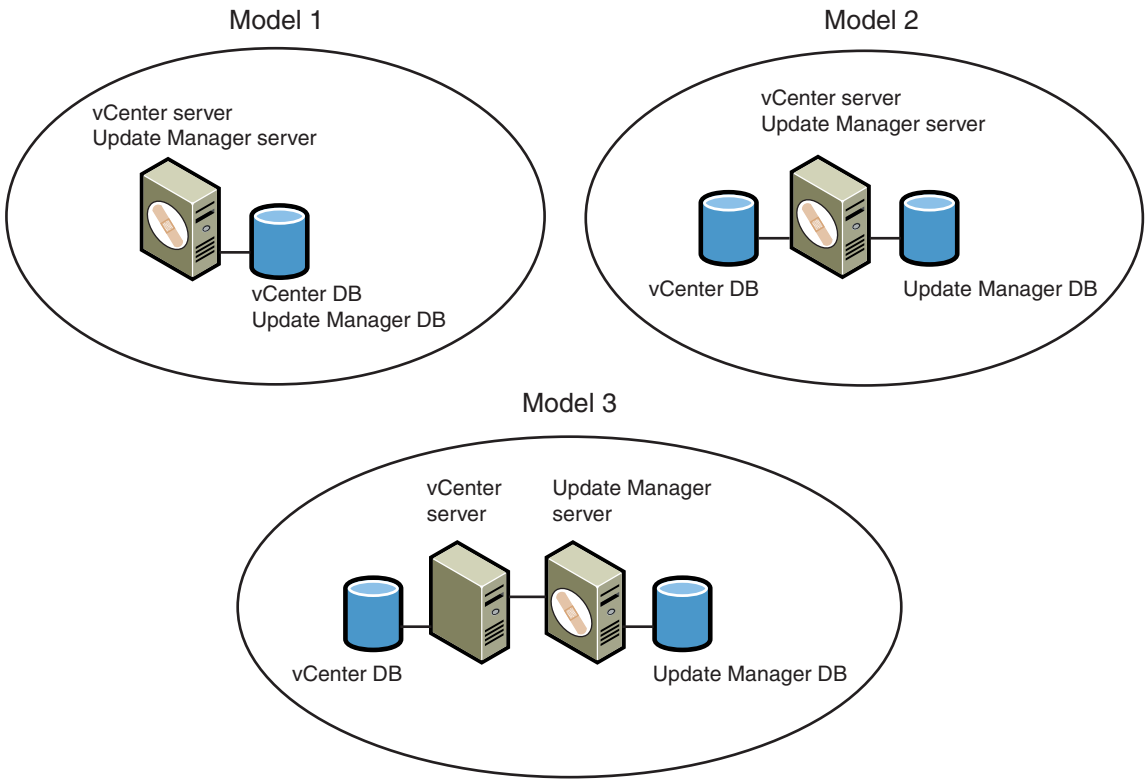

Performance Tips

- Separate the Update Manager database from the vCenter database when there are 300+ virtual machines or 30+ hosts.
- Separate both the Update Manager server and the Update Manager database from the vCenter server and the vCenter database when there are 1000+ virtual machines or 100+ hosts.
- Make sure the Update Manager server host has at least 2GB of RAM to cache patch files in memory.
- Allocate separate physical disks for the Update Manager patch store and the Update Manager database.

Latency Overview

Update Manager operation latency is an important metric. IT administrators need to finish applying patches within maintenance windows. Figure 4 and Figure 5 show latency results for guest agent download, powered-off Windows virtual machine scan, powered-on Windows virtual machine scan, ESX host scan, VMware Tools scan, virtual machine hardware scan, virtual machine hardware remediation, and VMware Tools remediation.

Guest agent download consists of multiple steps. The number presented in Figure 4 is the time required to send the file through the network from the Update Manager server to the virtual machine. For both the powered-on and powered-off virtual machine scan performance data presented in Figure 4 the guest agent is already installed. For the virtual machine hardware upgrade latency result in Figure 4, the virtual machine was originally powered off. For VMware Tools upgrade, the virtual machine was originally powered on. For compliance view of a folder with 500 virtual machines, the default critical virtual machine patch baseline and non-critical virtual machine patch baseline are attached to the folder and the latency is measured for fetching the compliance data for all attached baselines. All the numbers are averaged over multiple runs. The powered-on virtual machine scan showed higher deviation than did the powered-off virtual machine scan. Note that these latency numbers are only references. Actual latency varies significantly with different deployment setups.

Figure 4. Update Manager Operation Latency Overview (I)

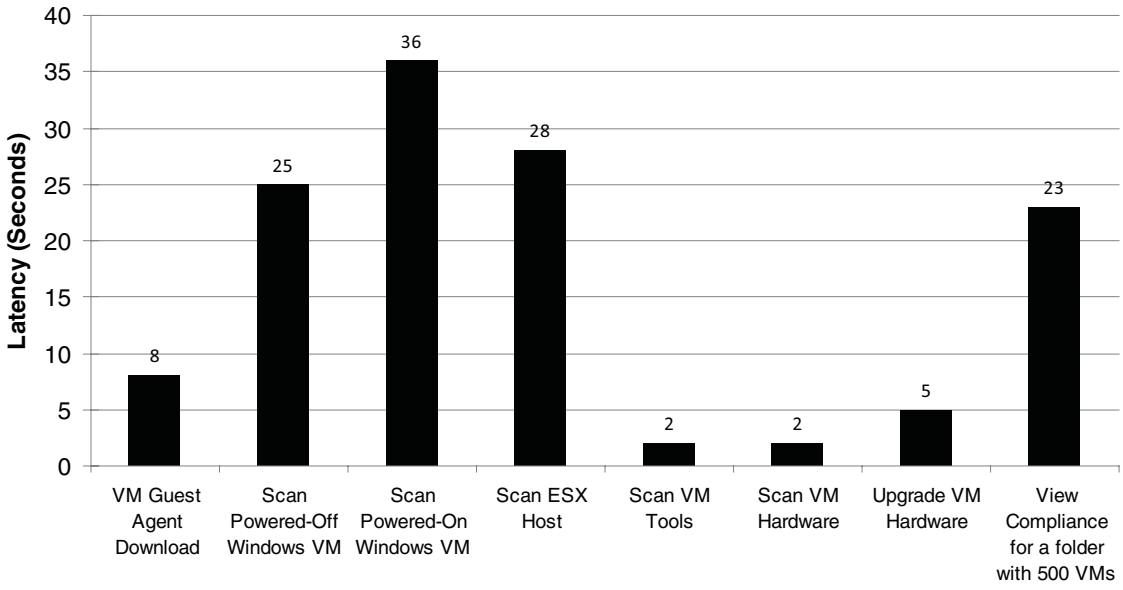

Figure 5. Update Manager Operation Latency Overview (II)

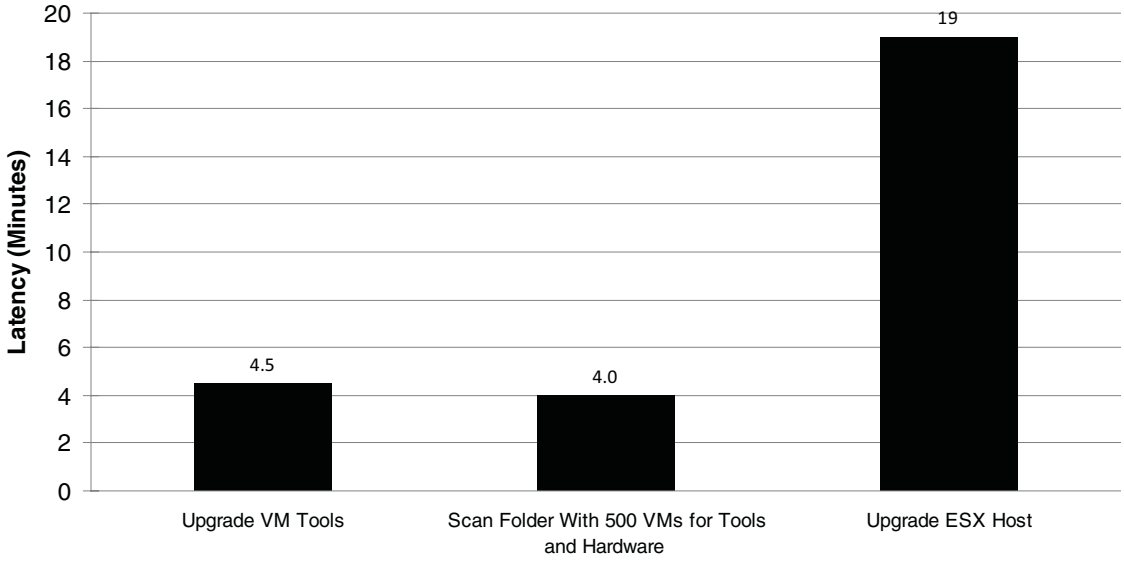

The results in Figure 4 and Figure 5 were measured on a low-latency local network setup. However network latency plays an important role for most Update Manager operations. For example, Figure 4 shows that a powered-off virtual machine scan finished faster than a powered-on virtual machine scan. In a high-latency network, however, a powered-on virtual machine scan can perform better than a powered-off virtual machine scan because the powered-on scan doesn't need to transfer as much data. For results obtained in a high-latency network see "Network Latencies" on page 10.

Performance Tips

- Because the Update Manager guest agent is installed in each Windows virtual machine the first time a powered-on scan is run on that machine, the first powered-on scan command can take longer than subsequent scans. It may therefore be desirable to run the first scan command when this additional time will not be an issue.
- Upgrading virtual machine hardware is faster if the virtual machine is already powered off. Otherwise, Update Manager will power off the virtual machine before upgrading the virtual hardware. This could increase the overall latency.
- Upgrading VMware Tools is faster if the virtual machine is already powered on. Otherwise, Update Manager will power on the virtual machine before the VMware Tools upgrade. This could increase the overall latency.
- For compliance view for all attached baselines, latency is increased linearly with the number of attached baselines. We recommend the removal of unused baselines, especially when the inventory size is relatively large.

Resource Consumption Matrix

Update Manager operations have different loads on the Update Manager server, vCenter server, and ESX host. Figure 6 and Figure 7 divide Update Manager operations into low, medium, and high resource consumption levels. For example, scanning a powered off Windows virtual machine results in high Update Manager server CPU usage, medium vCenter server CPU, network, disk, and database usage, and low ESX CPU usage. The powered-on Linux virtual machine scan latency measurement was done with a Linux virtual machine that did not have guest agent installed.

Figure 6. Resource Consumption for Update Manager Virtual Machine Patching Operations

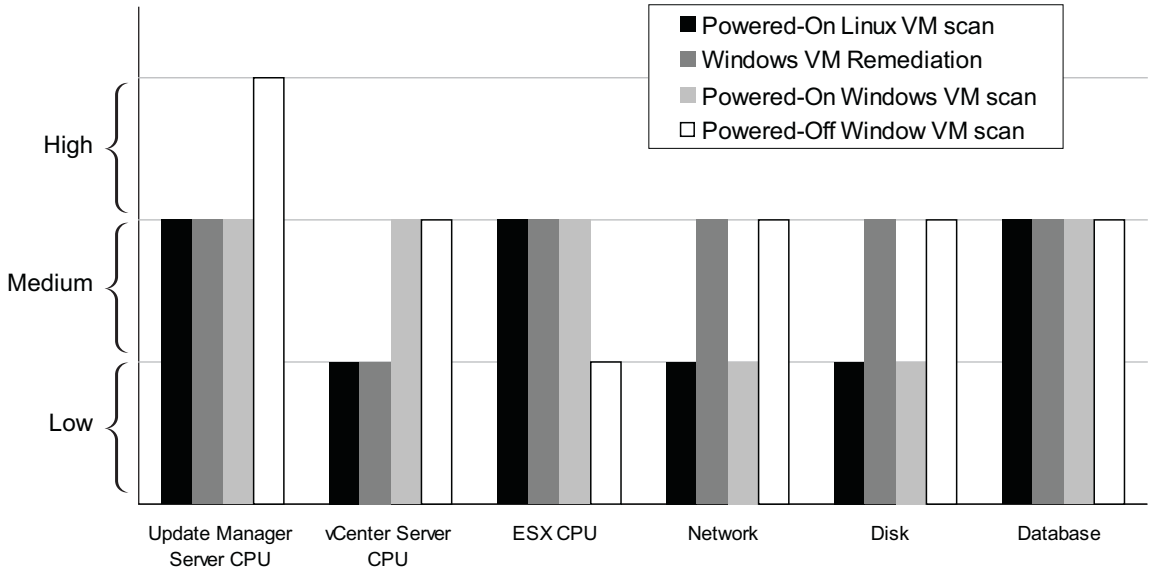

Figure 7. Resource Consumption for Update Manager 4.0 Only Operations

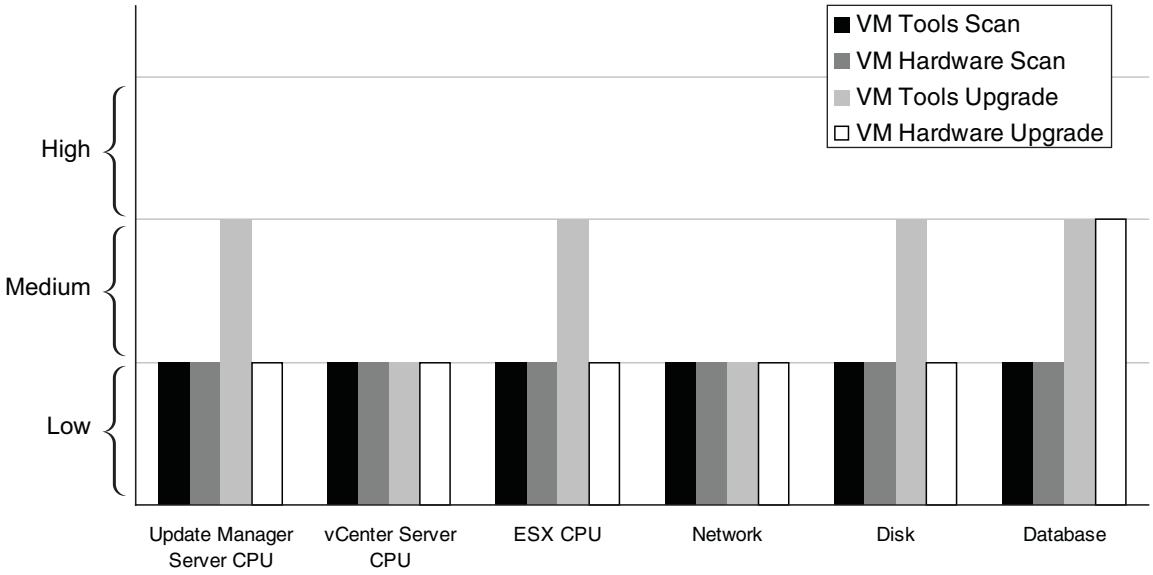

To alleviate resource pressure on Update Manager servers and ESX hosts, Update Manager performs job throttling by limiting the maximum number of concurrent operations. Table 1 gives the default performance limits.

Table 1. Job Throttling Default Limits for Each Update Manager Operation

Update Manager Operations	Maximum Tasks per ESX Host	Maximum Tasks per Update Manager Server
VM remediation	5	48
Powered-on Windows VM scan	6	72
Powered-off Windows VM scan	6	10
Powered-on Linux VM scan	6	72
ESX Host scan	1	72
ESX host remediation	1	48
VMware Tools scan	145	145
VM hardware scan	145	145
VMware Tools upgrade	145	145
VM hardware upgrade	145	145
ESX upgrade	1	48

The Update Manager server combines all operation types when calculating the limit. If a combination of operation types are being performed simultaneously, the formulas should be used with caution. For example, if powered-on Windows virtual machine scans and powered-off Windows virtual machine scans are being performed simultaneously, the maximum number of concurrent scans is *not* 12 (6+6) per ESX host, but rather 6 per ESX host.

Performance Tips

- For a large setup, powered-on virtual machine scans are preferred if Update Manager server resources are constrained or more concurrency is needed for scans.

Guest Operating System Tuning

The Update Manager guest agent installed on a virtual machine is single threaded for both powered-on virtual machine scans and virtual machine remediation. Multiple vCPUs thus do not reduce Update Manager operation latency.

The amount of RAM available to the guest operating system, however, can have a significant effect on the performance of virtual machine remediation operations. As shown in Figure 8, when a remediation command is issued from the vSphere client through the Update Manager plug-in, the Update Manager server first downloads the patches, creates an ISO patch file, and mounts it as a virtual CDROM to the guest operating system. The guest agent inside the guest operating system verifies the checksum of the patch file before the actual patch is applied. If the guest operating system RAM size is too small to cache the entire patch file into system cache, the patch file has to be read again. Because some Windows service packs can be as big as a few hundred megabytes in size, it is best to configure each virtual machine with at least 1GB of RAM so duplicate reads for the whole patch files can be avoided.

Figure 8. Guest Operating System Patching Flow

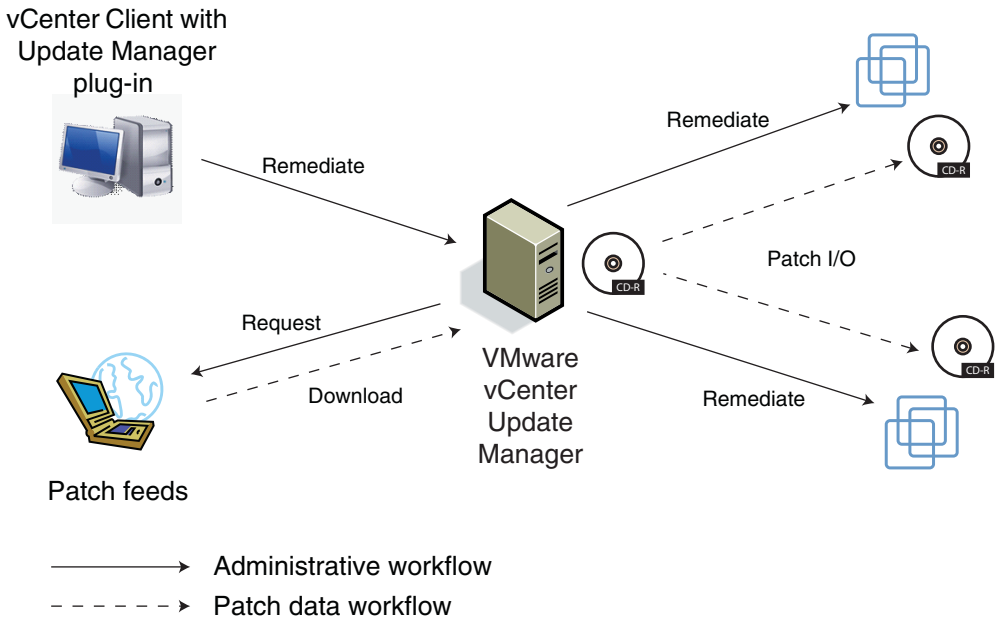

Performance Tips

- Multiple vCPUs do not help Update Manager operations as the Update Manager guest agent is single threaded.
- Configure each virtual machine with at least 1GB of RAM so large patch files can fit in the system cache.

Network Latencies

Update Manager server performance is strongly affected by network speeds. For powered-off virtual machine scans, the Update Manager server needs to read the registry information remotely from the ESX host. For virtual machine remediation, the Update Manager guest agent reads the patch ISO file through the network from the Update Manager host. We would thus expect the latency results for Update Manager operations to be impacted when the network latency increases. Figure 9 shows the latency results for powered-off virtual machine scans, powered-on virtual machine scans, virtual machine remediation, and guest agent download. As shown in Figure 2, WANem is used in this test to simulate a high latency network.

Figure 9. Operation Latency with a High Latency Network

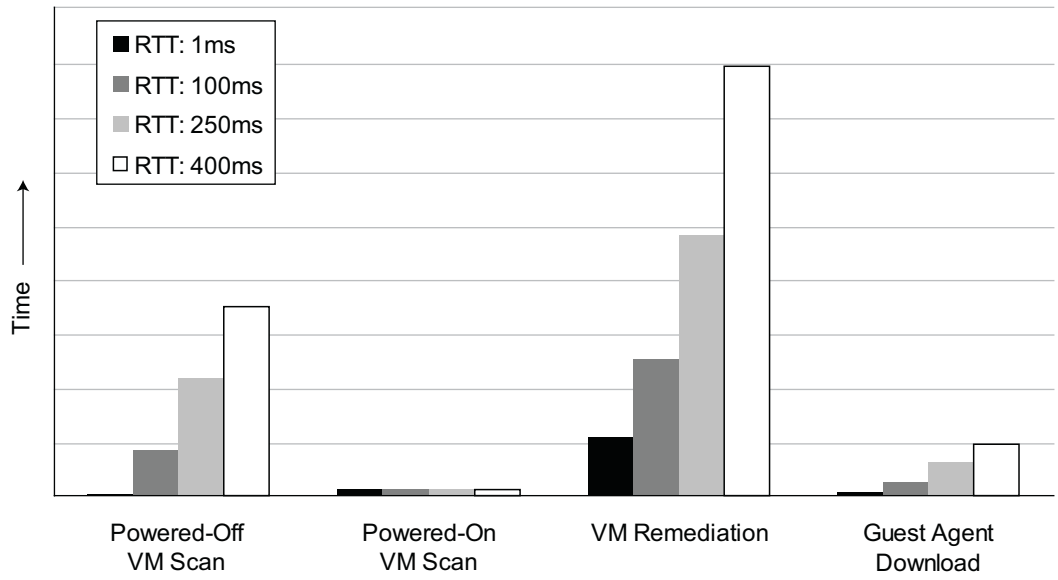

Network round trip times (RTTs) are typically 75 to 100 milliseconds for intra-U.S. networks, about 250 milliseconds for transatlantic networks, and 320 to 430 milliseconds for satellite networks.

In Figure 9, the latency of powered-off virtual machine scans, virtual machine remediation, and guest agent download operations increases linearly with the network speed. Powered-on virtual machine scans stay flat as most of the job is done inside the virtual machine and only the scan result info is transferred over the network.

For an unreliable network with packet drops, both powered-off virtual machine scans and virtual machine remediation are impacted significantly. The root cause is that the TCP retransmission timer varies from a few hundred milliseconds to a few seconds depending on the latency of the network. With a high-latency network, the TCP retransmission timer is much longer than with a low-latency network. It is therefore highly desirable to deploy the Update Manager server close to the ESX hosts to avoid network latency and packet drop.

Performance Tips

- Deploy the Update Manager server close to the ESX hosts if possible. This reduces network latency and packet drops.
- On a high-latency network, powered-on virtual machine scans are preferred as they are not sensitive to network latency.

On-Access Virus Scanning

As shown in Figure 10, for powered-off or suspended virtual machines a scan command is issued from the vSphere Client with the Update Manager plug-in and the Update Manager server mounts the virtual machine as a disk in the Update Manager server host. The Update Manager server then copies the hive files of the virtual machine and scans each file in the registry one by one. The file scan process has been optimized to read only necessary blocks of the scanned files to determine the version information. On average, there are only four reads per scanned file regardless of the actual size of each file. This is very important because the network between the Update Manager server and ESX host might be a high-latency network and reducing traffic on such a network can improve the performance dramatically.

Figure 10. Scan for Powered-Off or Suspended Virtual Machine

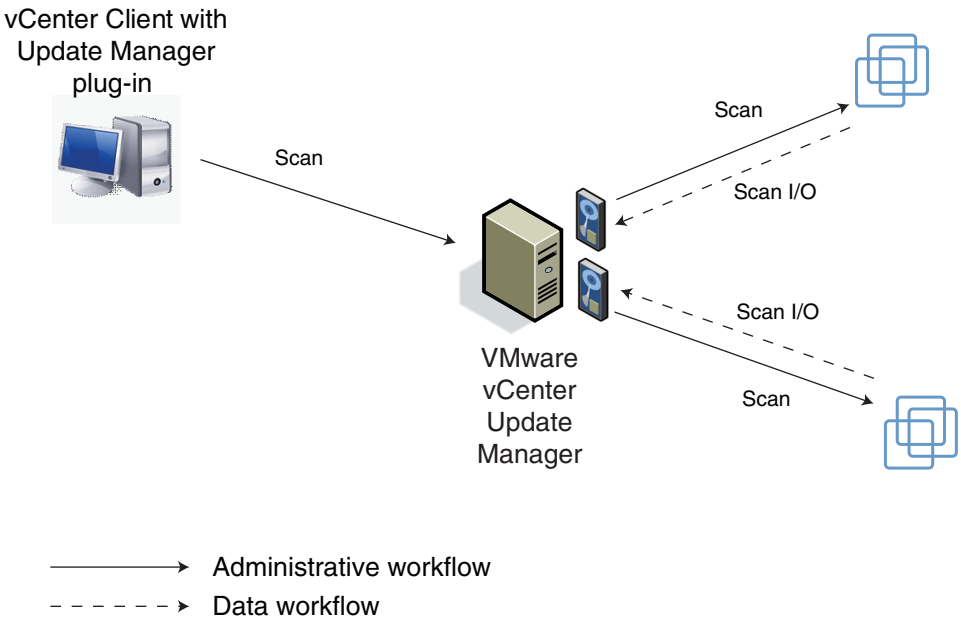

With on-access virus scanning, however, the optimization for reading only necessary blocks has no effect. When a file is opened, the on-access virus scanner will try to read the whole file to scan for viruses. This means all blocks of a scanned file will be fetched through the network. This is less of an issue on a low-latency network for powered-off scans as the file transfer speed is very high and thus the time to transfer the files can be ignored. Figure 11 show that on a low-latency network the total time to do powered-off scans is about the same with on-access virus scanning on or off. On a high-latency network, however, turning on-access virus scanning off can improve powered-off scan latency by almost 50%.

Figure 11. Powered-off Scan Latency With McAfee On-Access Anti-Virus

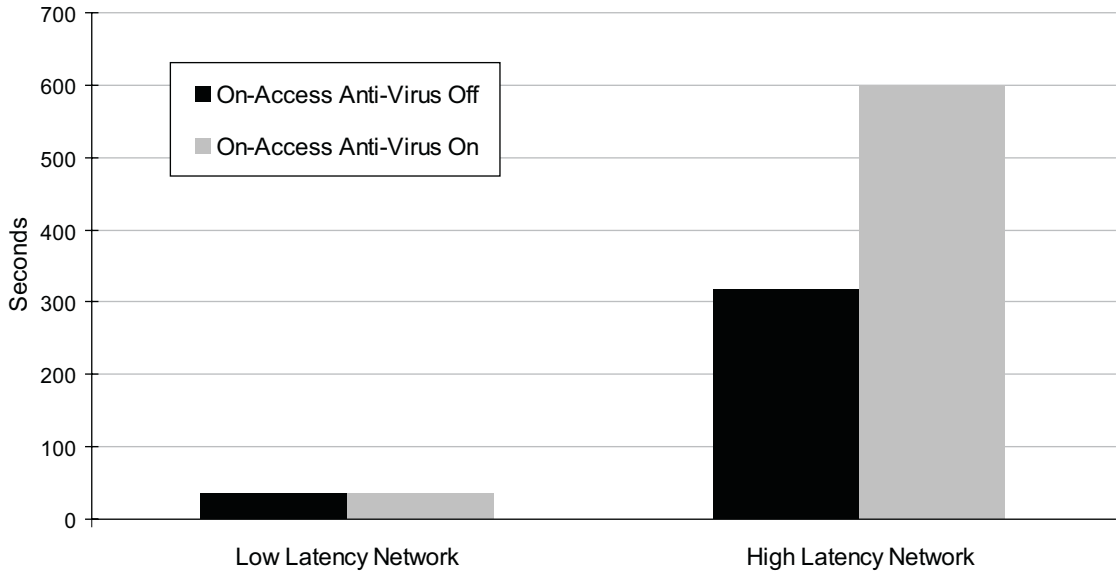

It is recommended that the mounted disk be excluded from on-access virus scanning to achieve optimal performance. Here is an example of how to exclude it for McAfee On-Access Anti-Virus.

1 Click the **All Processes** icon on the left of the **VirusScan On-Access Scan Properties** window (see Figure 12).

2 Select the **Detection** tab.

3 Click the **Exclusions...** button.

Figure 12. VirusScan On-Access Scan Properties Window

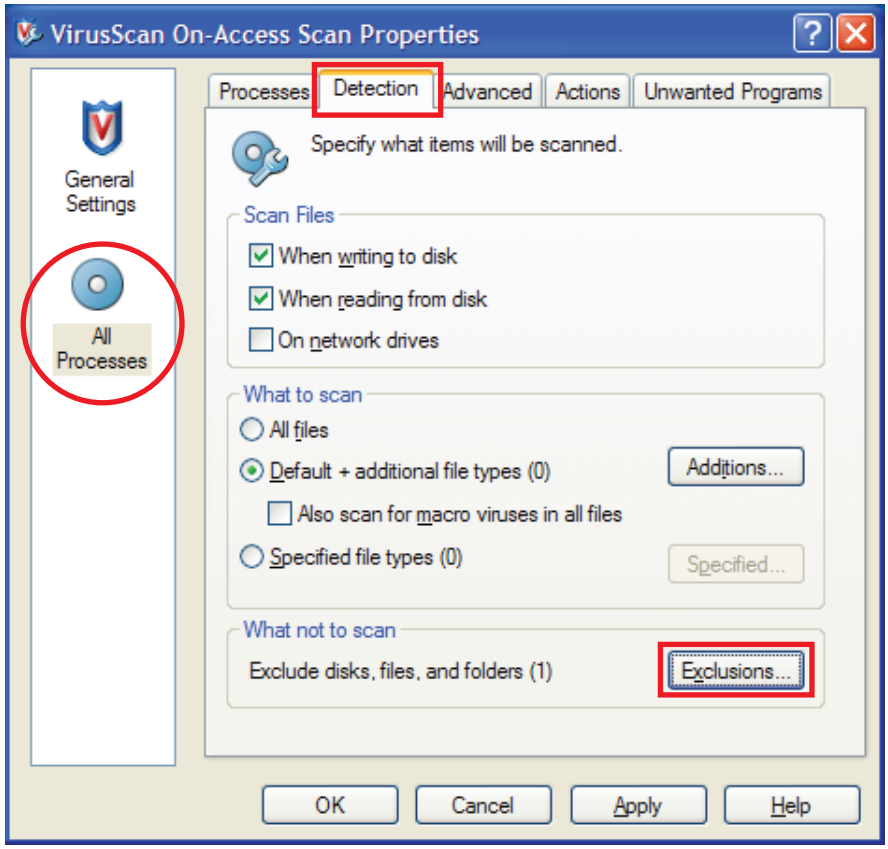

4 Click the **Add** button (see Figure 13).

Figure 13. VirusScan Set Exclusions Window

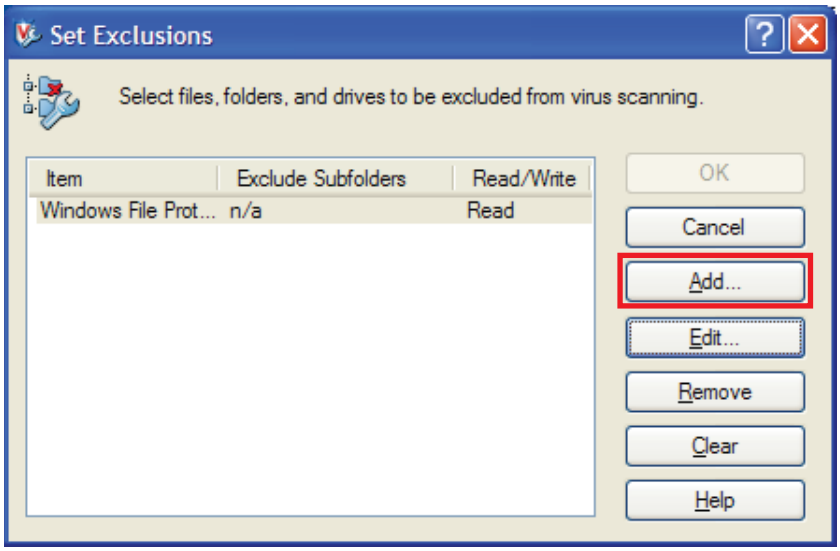

5 Select the **By name/location** radio button (see Figure 14).

6 Type:
\Device\vstor*

7 Select the **Also exclude subfolders**, **On read**, and **On write** checkboxes.

Figure 14. VirusScan Add Exclusion Item Window

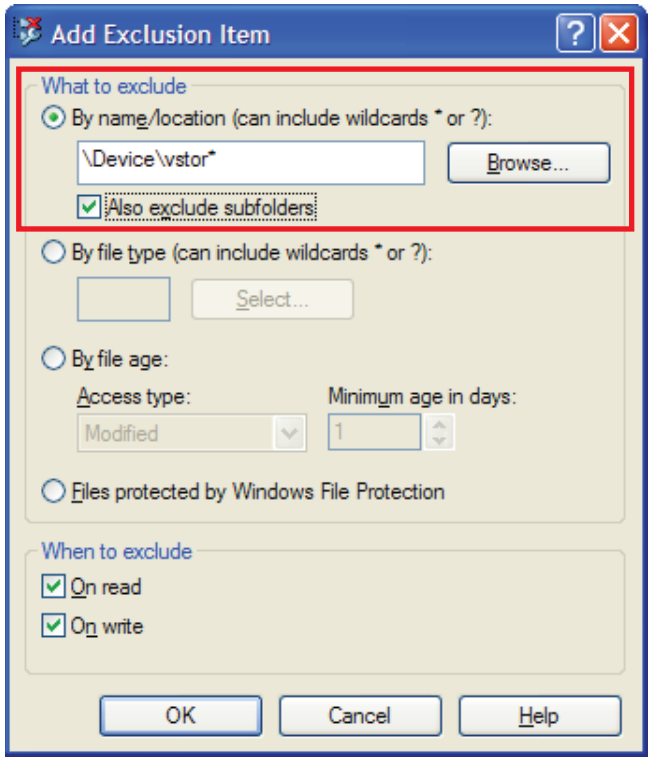

8 Click **OK** in the **Add Exclusion Item** window.

9 Click **OK** in the **Set Exclusions** window.

10 Click **OK** in the **VirusScan On-Access Scan Properties** window.

Performance Tips

- Check if on-access virus scanning software is running on the Update Manager server host. If it is, exclude the mounted disk, as described in this section.

Conclusion

VMware vCenter Update Manager delivers the most full-featured and robust patch management product for vSphere 4.0. In this white paper, the following performance recommendations have been made:

- Separate the Update Manager database from the vCenter database when there are 300+ virtual machines or 30+ hosts.
- Separate both the Update Manager server and the Update Manager database from the vCenter server and the vCenter database when there are 1000+ virtual machines or 100+ hosts.
- Make sure the Update Manager server host has at least 2GB of RAM to cache patch files in memory.
- Allocate separate physical disks for the Update Manager patch store and the Update Manager database.
- Because the Windows guest agent is installed in each virtual machine the first time a powered-on scan is run, the first powered-on scan command can take longer than subsequent scans. It may therefore be desirable to run the first scan command when this additional time will not be an issue.
- For a large setup, powered-on virtual machine scan is preferred if Update Manager server resources are constrained or more concurrency is needed for scans.
- Upgrading virtual machine hardware is faster if the virtual machine is already powered off. Otherwise, Update Manager will power off the virtual machine before upgrading the virtual hardware. This could increase the overall latency.
- Upgrading VMware Tools is faster if the virtual machine is already powered on. Otherwise, Update Manager will power on the virtual machine before the VMware Tools upgrade. This could increase the overall latency.
- For compliance view for all attached baselines, latency is increased linearly with the number of attached baselines. We recommend the removal of unused baselines, especially when the inventory size is relatively large.
- Multiple vCPUs do not help Update Manager operations as the Update Manager guest agent is single threaded.
- Configure each virtual machine with at least 1GB of RAM so large patch files can fit in the system cache.
- Deploy the Update Manager server close to the ESX hosts if possible. This reduces network latency and packet drops.
- On a high-latency network, powered-on virtual machine scans are preferred as they are not sensitive to network latency.
- Check if on-access virus scanning software is running on the Update Manager server host. If it is, exclude the mounted disk on a high-latency network.

References

- VMware vCenter Update Manager Administration Guide: http://www.vmware.com/pdf/vsp_vum_40_admin_guide.pdf
- VMware vCenter Update Manager Database Sizing Estimator: http://www.vmware.com/support/vsphere4/doc/vsp_vum_40_sizing_estimator.xls

About the Author

John Liang is a Staff Performance Engineer at VMware. Since 2007, John has been working as a technical lead for performance projects on VMware products such as VMware vCenter Update Manager, VMware vCenter Converter, VMware vCenter Site Recovery Manager (SRM), and VMware vCenter Lab Manager. Prior to VMware, John was a Principal Software Engineer at Openwave Systems Inc., where he specialized in large-scale directory server development and performance improvements. John received a Master of Science in Computer Science from Stony Brook University.

Acknowledgements

I want to thank Rajit Kambo for his valuable guidance and advice. His willingness to motivate me contributed tremendously to this white paper. I would also like to thank Sirish Raghuram, Jerome Purushotham, Roopak Parikh, Karthik Sreenivasa Murthy, Soam Vasani, Yufeng Zheng, Kevin Wang, and the Update Manager team for their tremendous support. Without their help this white paper would not have been possible. Last, I want to express gratitude to Monica Sharma for her great work communicating with Update Manager customers and providing useful feedback.

If you have comments about this documentation, submit your feedback to: docfeedback@vmware.com

VMware, Inc. 3401 Hillview Ave., Palo Alto, CA 94304 www.vmware.com

Copyright © 2009 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents. VMware, the VMware "boxes" logo and design, Virtual SMP, and VMotion are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Revision: 20090520; Item: EN-000157-01