McAfee VirusScan Enterprise for Storage 1.0

Sizing Guide for NetApp Filer on Data ONTAP 7.x

Contents

Overview

McAfee® VirusScan® Enterprise for Storage 1.0 (VSES) remotely scans NetApp filers and Internet Content Adaptation Protocol (ICAP) storage appliances, by expanding the capabilities of McAfee VirusScan Enterprise 8.7i.

McAfee VirusScan Enterprise for Storage 1.0 is an add-on to McAfee VirusScan Enterprise 8.7i. You can configure VirusScan Enterprise for Storage from the VirusScan Console or use ePolicy Orchestrator to centrally manage and enforce policies, then use queries and dashboards to track activity and detections.

The VirusScan Enterprise for Storage installation adds the following:

- Console — **Network Appliance Filer AV Scanner** and **ICAP AV Scanner**
- Service — **McAfee VirusScan Enterprise for Storage**, which manages scan requests received from the filer and uses the McAfee Engine Service to perform the actual scan.

When a file is accessed by the CIFS (Common Internet File System) user, the Data ONTAP sends the file's UNC path to VirusScan Enterprise for Storage for scanning. VirusScan Enterprise for Storage then scans the file and takes the necessary action, then sends the result back to the filer.

When a user tries to open, create or modify the files, the NetApp filer will send a scan request to VirusScan Enterprise for Storage when the following criteria is met:

- The file extension is included in the to-be-scanned file types list.
- The file has not already been marked as previously scanned, and no changes have occurred to the file.

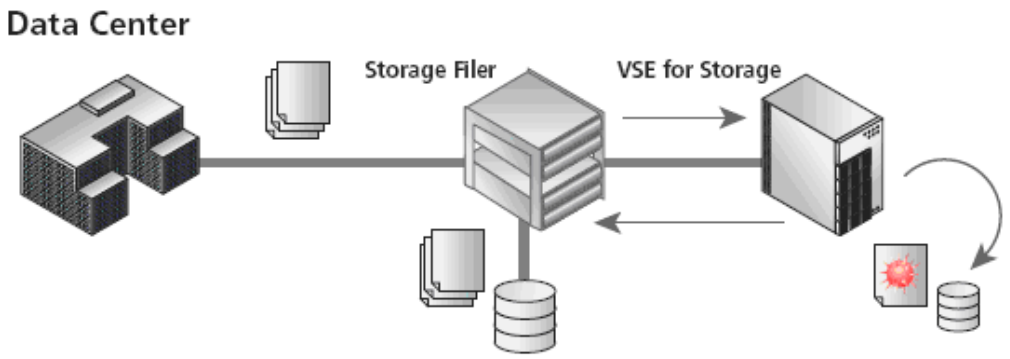

Sizing Overview

The objective of this sizing exercise is to measure the performance characteristics of the scanner and the NetApp Filer on various configurations. This helps an end-user to decide the number of scanners required for a given user load. McAfee has conducted these tests on the following filer configurations:

- Single NetApp filer with single Scan Server
- Single NetApp filer connected to two Scan Servers

All the tests are done with 100% write operation, which means each file sends a scan request to the scan server. The default Scanner settings and default NetApp filer "**vscan**" settings were used for these sizing tests.

A freeware tool having the capability of multi-threading feature was used for this exercise. The windows "**perfmon**" is used to capture scan server performance statistics. On the NetApp filer, "**systat**" command is used to capture the filer's performance statistics.

Contents

Samples used in the tests

Different file types with extensions of doc, pdf, xls, txt, zip, rar, ppt, eml, com, docx, dcom and tmp were included in the sample set. The sample set consisted of 37 files.

Sample size	Average = 3 MB Minimum = 1 KB Maximum = 30 MB
Clean vs. Infected ratio	60% to 40%

Hardware details

This section describes the hardware requirements, to measure the performance characteristics of the scanner and the NetApp Filer on various configurations for this sizing exercise.

Table 1: Sizing excercise — Hardware requirements

Hardware	Component
Filer	**Filer Model:** FAS 2020 **DataOntap Version:** 7.3.1.1 **Processor:** One, 2198MHz **Memory:** 896MB **NVMem:** 128MB **Network Adapter:** 1Gbps speed
Scan Server — Single Proc	**Model:** Dell Optiplex GX620 **RAM:** 1GB **Processor:** One Pentium4, 2.80GHz **Network Adapter:** 1Gbps speed **Operating System:** Microsoft Windows 2003 Enterprise Server (x64bit)
Scan Server — Dual Proc	**Model:** Dell PowerEdge 2900 **RAM:** 4GB **Processor:** 2x Intel Xeon CPU 5120 with Dual Core, 1.86GHz **Network Adapter:** 1Gbps speed **Operating System:** Microsoft Windows 2003 Enterprise Server (x64bit)
Domain Controller	**Model:** Dell Optiplex GX620 **RAM:** 1GB **Processor:** One Pentium4, 2.80GHz **Network Adapter:** 1Gbps speed **Operating System:** Microsoft Windows 2003 Enterprise Server (x64bit)

Testbed Configuration — Single Filer with single Scan Server

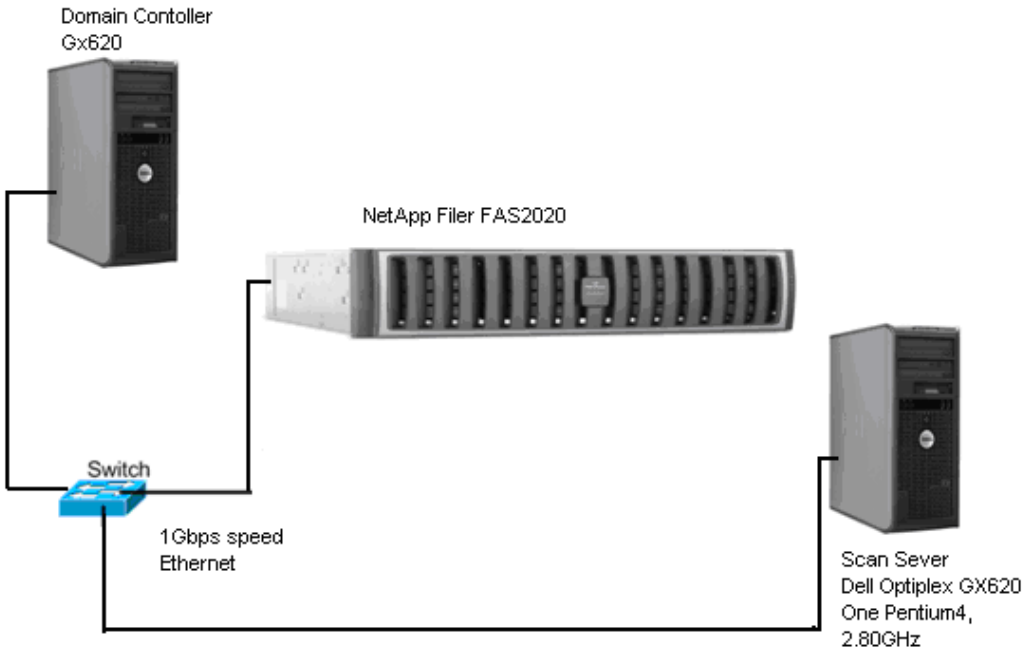

Testbed Configuration — Single Filer connected to two Scan Servers

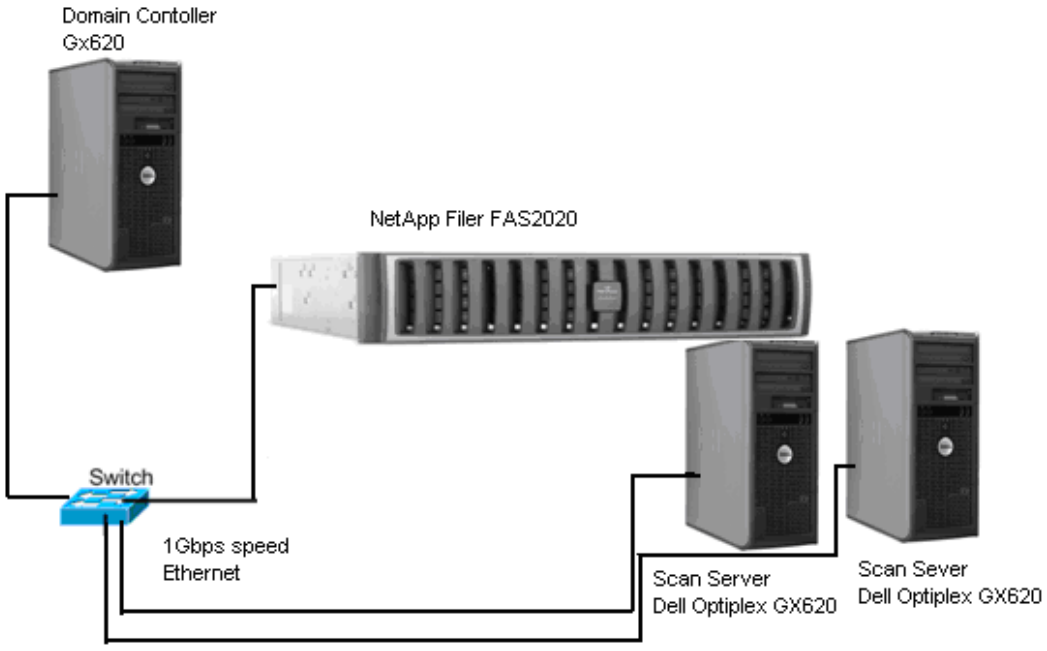

Performance Comparison between Single Proc and Dual Proc Scan Servers

Tests were conducted separately on a single proc server and dual proc server. Below are the comparison graphs of CPU utilization, response time and throughput for the same load and the same settings.

Contents

CPU utilization

The following figure shows the CPU utilization of Single Proc Scan Server and Dual Proc Scan Server:

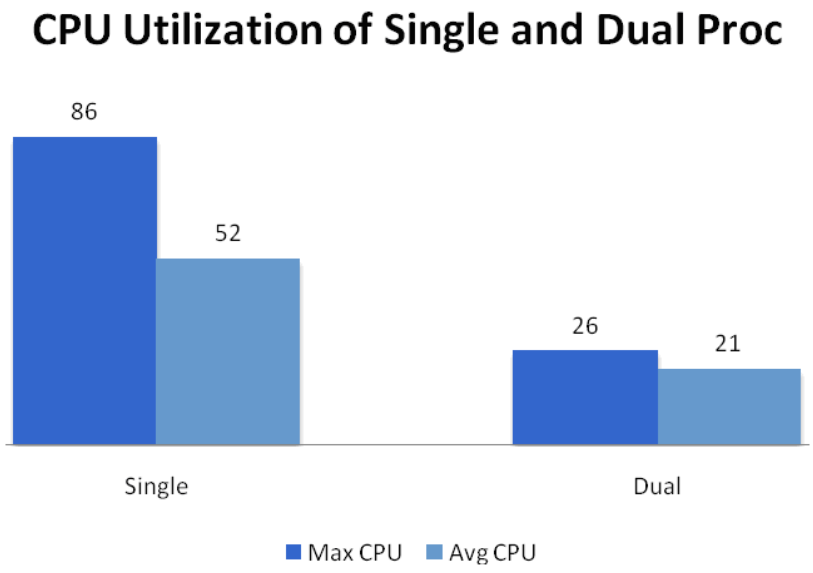

Tests were conducted on a single proc server and dual proc server seperately with the same load.

- **Load Profile:** 30 Users (30 clients accessing CIFS share)
- **Test Duration:** 1hour (Average of 3 runs)
- **VSES Settings:** Default settings

- **VSCAN Settings:** Default settings

Average Scan Response time

The following figure shows the average scan response time (in milliseconds) of Single Proc Scan Server and Dual Proc Scan Server:

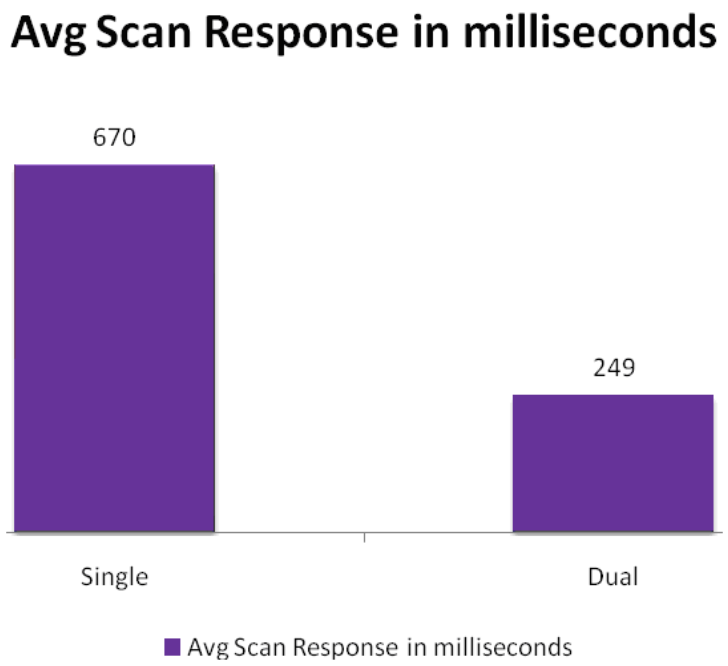

Tests were conducted on a single proc server and dual proc server seperately with the same load.

- **Load Profile:** 30 Users (30 clients accessing CIFS share)
- **Test Duration:** 1hour (Average of 3 runs)
- **VSES Settings:** Default settings.
- **VSCAN Settings:** Default settings

Average Throughput

When determining the sizing requirement, this average throughtput measurement is important. The following figure shows the average throughput in KB/s of Single Proc Scan Server and Dual Proc Scan Server.

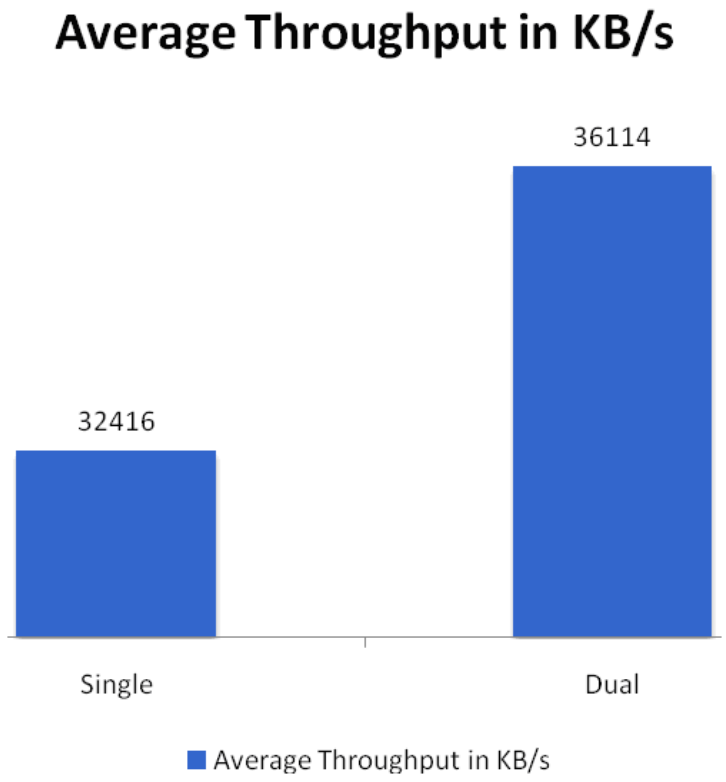

Tests were conducted on a single proc server and dual proc server seperately with the same load.

- **Load Profile:** 30 Users (30 clients accessing CIFS share)
- **Test Duration:** 1hour (Average of 3 runs)
- **VSES Settings:** Default settings
- **VSCAN Settings:** Default settings

Performance Statistics

The test was done on a Single filer with a **Dual Proc Scan Server**. The tests were done by gradually increasing the number of users writing to filer simultaneously with default settings. In this test filer's (FAS2020), the CPU utilization reached 100% but the Scan Server CPU averaged to 50% for 300 users.

Scan Server Performance Statistics

- **Average CPU utilization:** 50%
- **Max scan threads:** 53
- **Engine Server Throughput (KB/s):** 28850
- **Average response Time (milliseconds):** 3369

Filer Performance Statistics

- **Peak CPU utilization:** 100%
- **Average CPU utilization:** 90%
- **CIFS Throughput:** 64766 KB per second (network in+out)
- **KB/user/s:** 215

The following figure concludes that the average CPU and Memory (Working Set) utilization increases by increasing number of users accessing the filer:

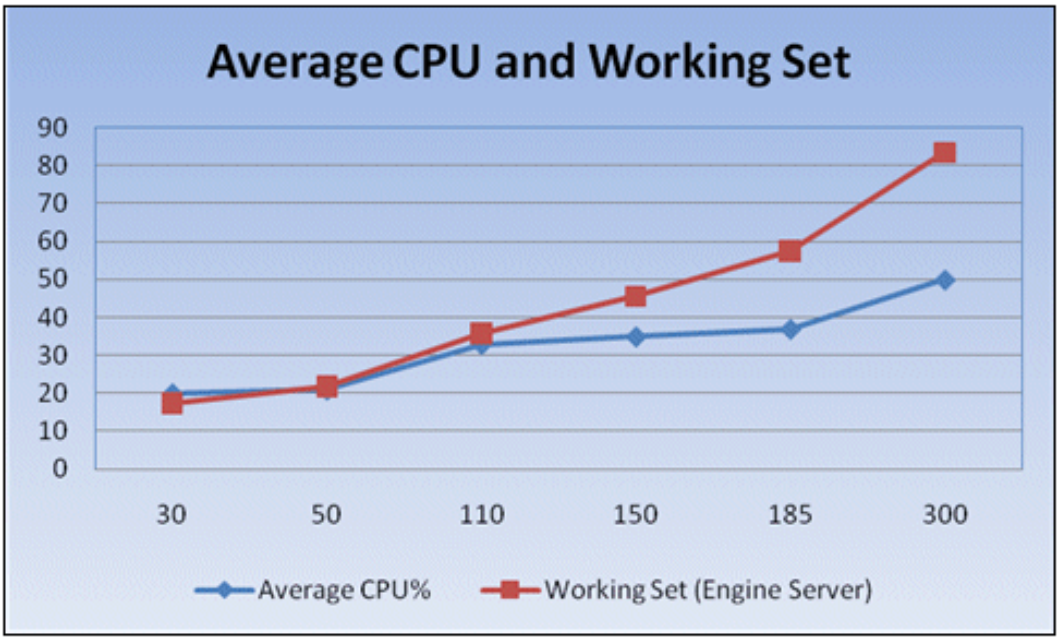

The following figure concludes that the average throughput increases, as the number of users writing to filers increase till the scanner/filer is able to handle. In this scenario, the throughput

actually decreases beyond 200 users accessing the filer, as this filer is not able to scale beyond this.

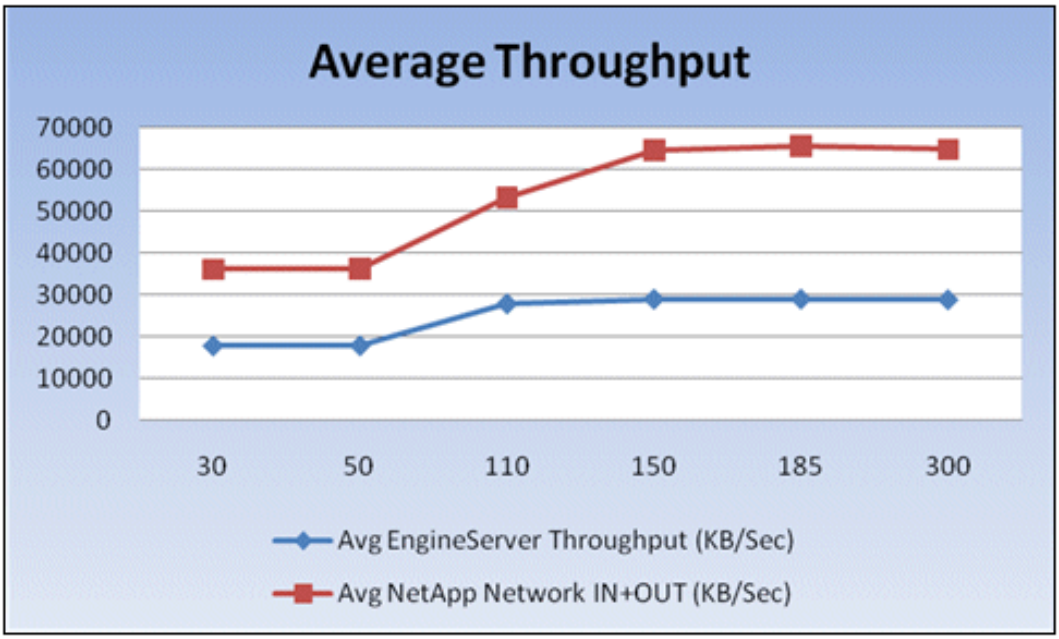

Performance Improvement by adding additional Scan Servers

By adding an additional server, there is a significant improvement in average scan response time and average data throughput. The scan requests will be sent to the scanners in a round-robin fashion, to distribute the load by NetApp filer.

The following table shows the significant improvement in the average scan response time:

Table 2: Average Response Time (in ms)

# of Users	One Scan Server	Two Scan Servers
15	546	305
30	631	307
50	670	320

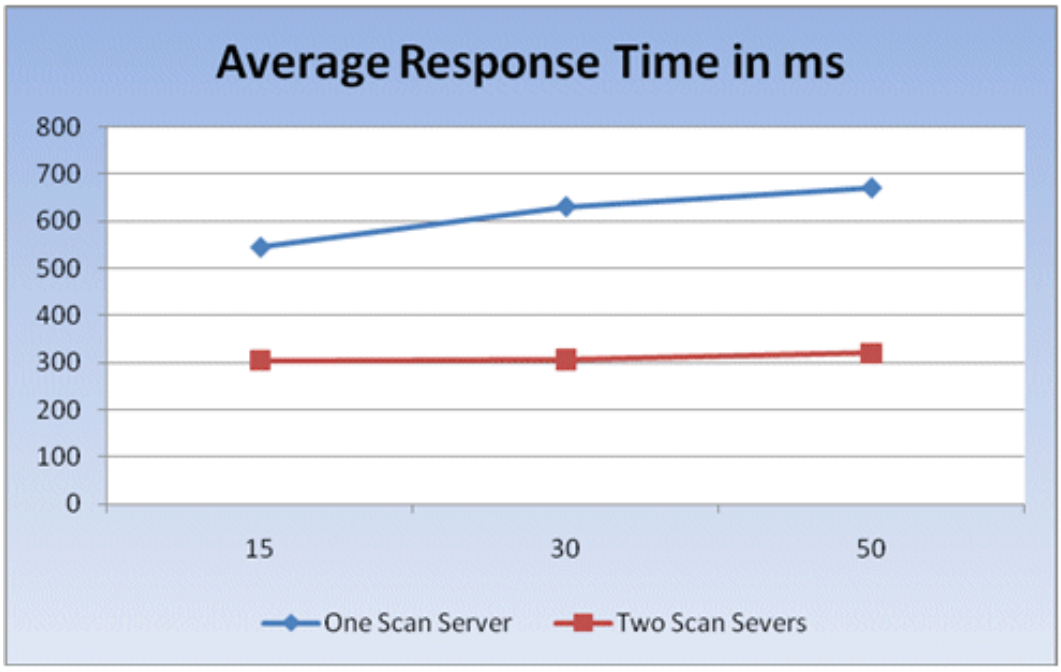

The following table shows the significant improvement in the average throughput:

Table 3: Average Throughput (in KB/s)

# of Users	One Scan Server	Two Scan Servers
15	32110	34346
30	32146	34678

# of Users	One Scan Server	Two Scan Servers
50	33112	34986

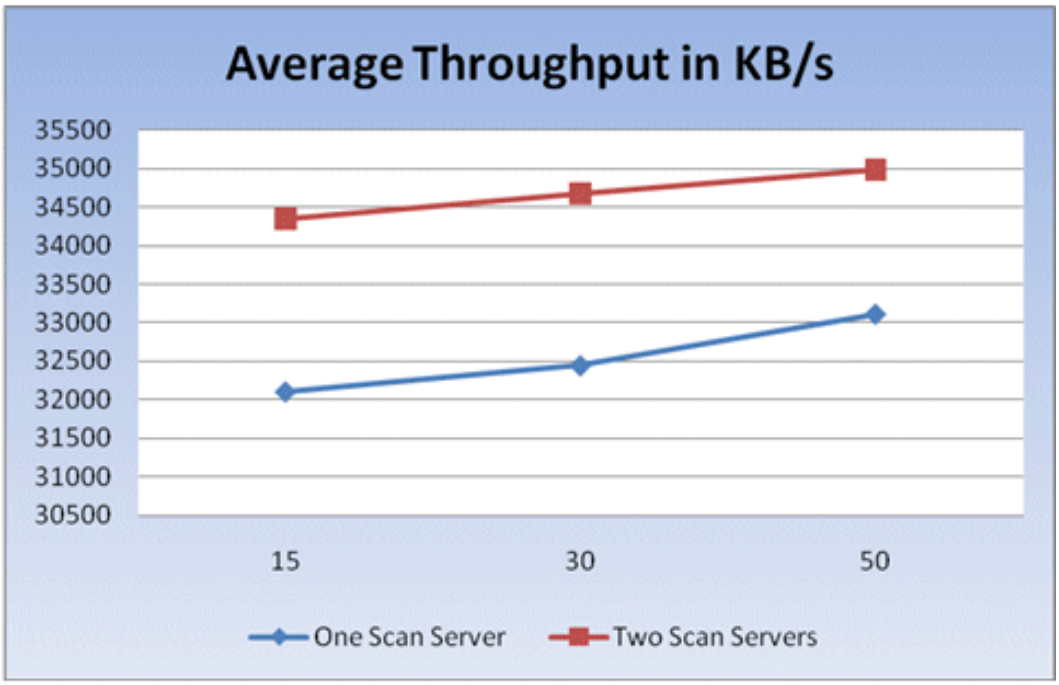

Performance Considerations

Network connectivity

Gigabit Ethernet (1000MB) networking is the minimum recommendation between the NetApp storage device and the Scan Server.

The NetApp filer and Scan Server can be connected directly or connected using a private network to achieve a clean network. If connecting through an Ethernet switch, use a dedicated switch or configure a virtual LAN (VLAN) to separate traffic traveling through the switch.

Scan server hardware

The limiting performance factors with AV scanners are the CPU frequency and memory bus performance. The amount of "**vscan**" requests the AV scanners will be able to process, will be greatly dependent on their CPU speed and the memory bus bandwidth.

McAfee recommends the following system requirements for each AV server:

- Dual Processor 2.6GHz (or greater)
- Network Interface Card 1000MB/s Ethernet

Average File Size and Type of files

AV Server performance will be affected by average size of the files and file types it scans (such as .zip, .cab, etc).

Deploying VSES on Virtual Environments

The limiting performance factors with AV scanners are the CPU frequency and memory bus performance. The VMware/ESX infrastructure should be capable of handling the CPU and Memory requirements of AV scanner, as given below:

- Dual Processor 2.6GHz (or greater)
- Network Interface Card 1000MB/s Ethernet

Filer Storage Capacity

The filer storage capacity does not affect the scanner's performance. The performance actually depends on the number of concurrent users accessing the filer and the average throughput of the filer.

Test Summary

This section summarizes the performance characteristics of the scanner and the NetApp Filer on various configurations, which helps the end-user to decide the number of scanners required for a given user load.

The following table summarizes the sizing tests for various configurations:

# of CPU	# of Scanners	# of Users	Avg CPU %	Avg Engine Server Thoughput (KB/s)	Avg NetApp n/w In+Out (KB/s)	KB/User/s	Avg Working Set Engine Sever(MB)	Response Time (ms)	Processor Queue Length	# of files/s
1	1	1	22	8484	17201	17201	8.26	202	0.81	2.19
1	1	15	52	15731	32110	2141	24.26	546	2.64	4
1	1	30	53	15946	32416	1081	31.5	632	3.30	4
1	1	50	76	16874	33112	662	36.69	692	3.98	4
1	2	15	26	17318	34346	2290	19.91	307	1.07	5
1	2	30	27	18452	34678	1156	22.12	312	1.09	10
2	1	30	20	17745	36114	1204	17.32	256	0.14	5.36
2	1	50	21	17797	36167	724	21.7	259	0.15	5.52
2	1	110	33	27861	53297	485	35.87	574	0.64	7.8
2	1	150	35	28962	64564	430	45.54	807	0.73	8.4
2	1	185	37	28979	65513	354	57.4	949	0.99	8.8
2	1	300	50	28850	64766	215	83.5	3369*	1.47	12

*The filer CPU Utilization reached 100%

The Standard CIFS User Classification from NetApp is classified as:

- **Heavy User** — A heavy user transfers around 20 KB/s of network data (in and out) from client to CIFS server.
- **Medium User** — A medium user transfers around 7 to 13 KB/s of network data (in and out) from client to CIFS server.
- **Low User** — A low user transfers around 3 to 6 KB/s of network data (in and out) from client to CIFS server.
- **Light User** — A light user transfers around 0 to 2 KB/s of network data (in and out) from client to CIFS server.

Based on the tests conducted on various scenarios, the following recommendations are suggested for the sample load as mentioned earlier.

Table 4: Recommendations for the sample load

Heavy users accessing the filer simultaneously	Recommended # of Scan Server(s)
Up to 30	One Single Proc Scan Server
Up to 75	Two Single Proc Scan Servers
Up to 300	One Dual Proc Scan Server

NOTE: For a failover, it is highly recommended to add an additional scan server as secondary scan server.

The following illustration shows the sample scanning pod:

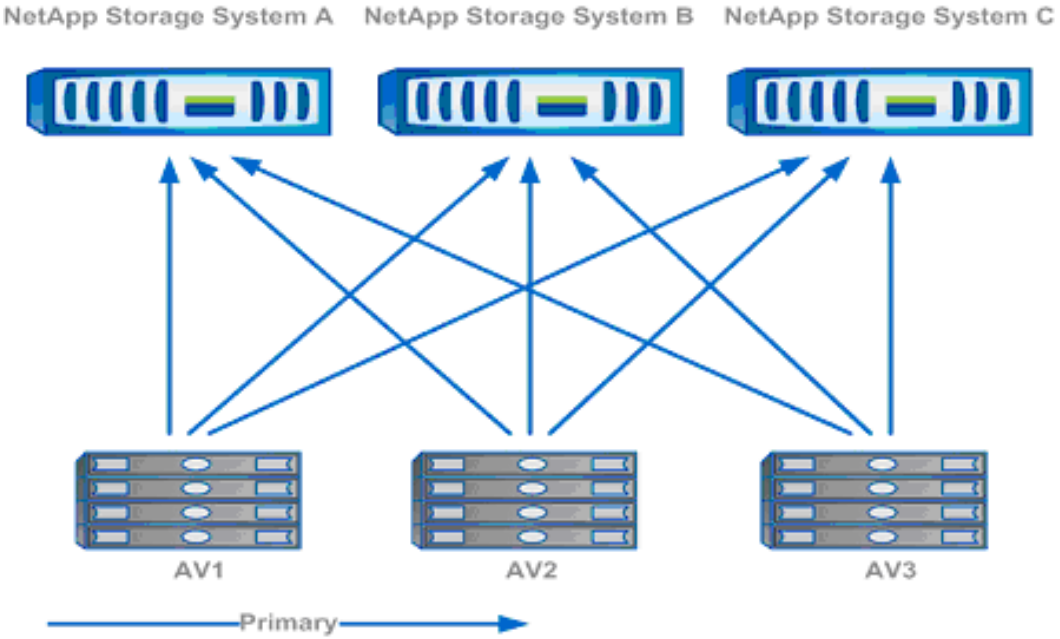

Figure 1: All the scanners as primary inside a data center

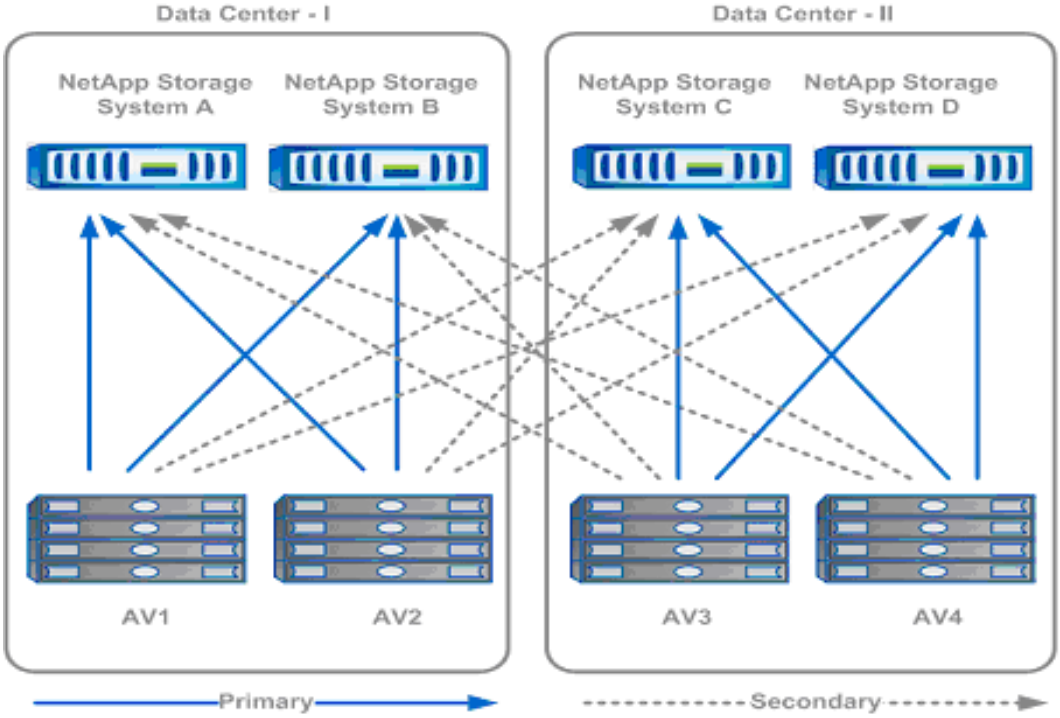

Figure 2: AV scanning pod between two data centers