

InterScan Web Security Virtual Appliance

Sizing Guide for version 6.0

July 2013

TREND MICRO INC.

10101 N. De Anza Blvd. Cupertino, CA 95014 • www.trendmicro.com

• **Toll free:** +1 800.228.5651 • **Fax:** +1 408.257.2003 • **Phone:** +1 408.257-1500

Contents

Executive Summary	1
Assumptions and Recommendations	1
General Sizing Guidelines.....	2
Sizing at a Glance – Software Appliance (Bare Metal)	2
Sizing at a Glance – Virtual Appliance (VMware ESXi 5.0).....	3
Sizing at a Glance – Virtual Appliance (Hyper-V 3.0).....	4
Sizing at a Glance – Log query impact.....	5
Sizing at a Glance – Log entries V.S. Disk size	5
Sizing at a Glance – Reporting	6
Calculating the Number of IWSVA Servers Required	6
Step 1: Obtain the Required Data for the Customer Environment	7
Step 2: Determine the Number of Required Servers.....	8
Step 3: Determine the Storage	11
Step 4: Determine the Standalone log server	12
Appendix A	13
How Tests Were Conducted	13
What Configuration Changes do to Sizing	14
Reporting	14
Caching	14
Performance Criteria for Tests	14
Scalability and Accuracy	14
Hardware Tested	15
About Trend Micro Incorporated	16

Copyright© 2013 by Trend Micro Incorporated. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the prior written consent of Trend Micro Incorporated. Trend Micro, the t-ball logo, and InterScan are trademarks or registered trademarks of Trend Micro Incorporated. All other company and/or product names may be trademarks or registered trademarks of their owners.

Information contained in this document is provided "as-is" and subject to change without notice. This report is for informational purposes only and is not part of the documentation supporting Trend Micro products.

TREND MICRO MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS REPORT.
[ETM Part No: SG_SVA_201206US]

This document is a product of Trend Micro Technical Sales Solutions.

Executive Summary

Trend Micro conducted tests on Trend Micro™ InterScan™ Web Security Virtual Appliance v6.0 (IWSVA) to obtain performance-sizing data for customer deployment.

Several factors impacts the sizing results for an IWSVA deployment, including:

- CPU
- Memory
- Which scanning features are enabled and what percentage of traffic will be scanned by those features

For the latest information about InterScan Web Security Virtual Appliance, including product documentation, server hardware support, and the latest software builds, visit the Trend Micro Web site at:

<http://www.trendmicro.com/us/enterprise/network-security/interscan-web-security/index.html>

<http://downloadcenter.trendmicro.com/>

Assumptions and Recommendations

Assumptions:

- The *user population* is the total number of potential Web users within an organization.
- The number of *active users* is the total number of client workstations that are simultaneously requesting HTTP content at any point in time. This is assumed to be 20% of the *user population* by default.
- Gigabit network interface cards (NIC) are used throughout the LAN.
- HTTPS scanning is disabled by default.
- When measuring the impact of Application Control policies on performance, we assumed 35% of the network traffic to be non-HTTP (e.g. UDP, Skype, bitTorrent, etc.)
- Each active user has in average 3.5 opening HTTP connections to the Internet at any one time.
- Each Internet access generates in average 5 log events
- There are three feature set configurations tested:
 - A basic feature set, including virus scanning, advanced threat scanning, Web reputation, URL filtering, and Anti-Botnet
 - An advanced feature set, which adds Application Control to the Basic feature set
 - The advanced feature set with Data Loss Prevention enabled
- The impact of HTTPS scanning and LDAP authentication is discussed after the performance data from the three configurations are presented.

Recommendations:

- Hardware should meet the minimum requirements specified in the Installation Guide. IWSVA requires having adequate memory to establish TCP connections. If there is not enough memory, system performance will be restricted. In IWSVA v6.0, we recommend the ratio of Memory to the number of CPU Threads shall be larger than the following calculation result:

$$\text{Memory} = \text{CPU_Threads_Num} \times 0.5 + 3$$

Note: Minimum memory requirement is 4GB.

General Sizing Guidelines

Administrators can obtain the general sizing on a per-server basis if the specification of the available hardware is known. The sizing below shows the three configurations of enabled features in both Forward Proxy mode and Transparent Bridge mode. Content caching was not used in these tests.

IWSVA performance is CPU-bound, and the number of CPU threads will affect performance. For multi-core or hyper-threading systems, each CPU thread is considered for the purpose of this sizing guide.

Note: A linear increase in allocated CPU GHz does not equate to a linear performance increase.

The sizing information in this guide is the **absolute maximum** the specified hardware configuration can support before the additional latency introduced to average page download times exceeds 2 seconds. The results in the tables below are accurate to within ± 5 percent.

Sizing at a Glance – Software Appliance (Bare Metal)

The features including virus scanning, advanced threat scanning, Web reputation, Anti-Botnet and URL filtering are active on an IWSVA 6.0 server in Forward Proxy and Transparent Bridge mode deployment; Tables 1, 2, and 3 provide recommendations for bare-metal software appliance general sizing respectively for the three feature set configurations described in Assumptions and Recommendations: 1. Basic feature set, 2. Basic feature set plus App control, 3. Basic feature set plus App control and DLP.

Table 1 IWSVA 6.0 Sizing at a Glance with All Scanning Policies Active and disable Application Control (Only basic feature set)

Server Type	Memory Size	Concurrent Connections	HTTP Transactions per Second	Throughput (Mbps)	Maximum Total User Population per device
Dell R210II 1 x Intel® Xeon® E3-1240 (4 Cores, 8 Threads)	8 GB	7,500	4,800	600 Mbps	10,000
Dell R420 2 x Intel® Xeon® E5-2420 (6 Cores, 12 Threads)	16 GB	10,000	7,000	850 Mbps	15,000
Dell R720 2 x Intel® Xeon® E5-2660 (8 Cores, 16 Threads)	24 GB	11,500	8,000	~1,000 Mbps	16,500

Table 2 IWSVA 6.0 Sizing at a Glance with All Scanning Policies Active and enable Application Control (Basic feature set + App control)

Server Type	Memory Size	Concurrent Connections	HTTP Transactions per Second	HTTP Throughput (Mbps)	Maximum Total User Population per device
Dell R210II 1 x Intel® Xeon® E3-1240 (4 Cores, 8 Threads)	8 GB	6,000	3,800	500 Mbps	8,500
Dell R420 2 x Intel® Xeon® E5-2420 (6 Cores, 12 Threads)	16 GB	9,000	5,800	750 Mbps	13,000

Dell R720 2 x Intel® Xeon® E5-2660 (8 Cores, 16 Threads)	24 GB	11,000	7,200	950 Mbps	15,000
---	-------	--------	-------	----------	--------

Table 3 *IWSVA 6.0 Sizing at a Glance with All Scanning Policies Active and enable Application Control and DLP (Basic feature set + App control + DLP)*

Server Type	Memory Size	Concurrent Connections	HTTP Transactions per Second	Throughput (Mbps)	Maximum Total User Population per device
Dell R210II 1 x Intel® Xeon® E3-1240 (4 Cores, 8 Threads)	8 GB	5,500	3,600	450 Mbps	5,200
Dell R420 2 x Intel® Xeon® E5-2420 (6 Cores, 12 Threads)	16 GB	8,500	5,600	700 Mbps	8,000
Dell R720 2 x Intel® Xeon® E5-2660 (8 Cores, 16 Threads)	24 GB	10,500	7,000	900 Mbps	13,500

Please refer to Appendix A for the specific hardware configuration of each server type used above.

Note: Features and Modes Impact on Performance:

- The feature with the largest use of system resources is virus scanning, whilst the second is application control.
 - The use of Web reputation, URL filtering, and Anti-Botnet, only modestly lowers system performance and does not significantly decrease the supported user population per server.
 - Transparent bridge mode performance was tested no obvious difference than forward proxy mode performance.
 - For DLP feature, we configured a policy with only a single template (HIPAA compliance). The sizing and capacity numbers in table 1 were reduced by approximately 10%.
 - After enabling LDAP authentication, it will cause nearly 14% drop of sizing capacity number in general.
 - Enabling HTTPS scanning in a network that contains 15% of all traffics, as HTTPS will reduce the performance numbers by approximately 26%. This is normal since HTTPS key negotiation; decryption and re-encryption are CPU-bound activities. Environments with HTTPs scanning requirements should consider to higher level CPU configuration;
-

Sizing at a Glance – Virtual Appliance (VMware ESXi 5.0)

Table 4 provides the general sizing for virtual appliance deployment in VMware ESXi environment.

Table 4 *IWSVA 6.0 Sizing at a Glance with All Scanning Policies Active and enable Application Control and DLP (Basic feature set + App control + DLP)*

Server Type	Memory Size	Concurrent Connections	HTTP Transactions per Second	Throughput (Mbps)	Maximum Total User Population per device
VM (2 core)	4 GB	1,500	950	120 Mbps	1,370
VM (4 core)	6 GB	3,250	2,050	260 Mbps	2,970
VM (16 core)	16 GB	6,990	4,400	560 Mbps	6,300

The performance difference of installing IWSVA on a VMware virtual appliance verses installing IWSVA on a bare metal software appliance server is nearly 10%. The performance degradation under VMware is normal and can be attributable to the VMware OS overhead that is required to manage the Virtual Machines (VMs) and the resources being shared.

Sizing at a Glance – Virtual Appliance (Hyper-V 3.0)

Table 5 provides the general sizing for virtual appliance deployment in Microsoft Windows Server 2012 Hyper-V environment.

Table 5 *IWSVA 6.0 Sizing at a Glance with All Scanning Policies Active and enable Application Control and DLP (Basic feature set + App control + DLP)*

Server Type	Memory size	Concurrent Connections	HTTP Transactions per Second	Throughput (Mbps)	Maximum Total User Population per device
VM (2 core)	4 GB	1,500	980	120 Mbps	1,390
VM (4 core)	6 GB	3,200	2,100	260 Mbps	2,990

The performance difference of installing IWSVA on a Hyper-V virtual appliance verses installing IWSVA on a bare metal software appliance server (both using 2/4 CPUs and the same amount of memory and disk) is approximately 10%. The performance degradation under Hyper-V is normal and can be attributable to the Hyper-V OS overhead that is required to manage the Virtual Machines (VMs) and the resources being shared.

Sizing at a Glance – Log query impact

Log query may need to read large amount of historical data. The data amount decides the system resource usage for the log query.

In this log-query testing we tested with the total 2.2-billion log entries and keeping query 1.1-billion log entries to identify the maximum performance impacts.

Table 6 IWSVA 6.0 Sizing at a Glance with Log query impact

Server Type	Memory Size	Concurrent Connections	HTTP Transactions per Second	Throughput (Mbps)	Maximum Total User Population per device
Dell R210II 1 x Intel® Xeon® E3-1240 (4 Cores, 8 Threads)	8 GB	5500	3628	462.60	5183
Dell R210II 1 x Intel® Xeon® E3-1240 (4 Cores, 8 Threads) <i>With Log query</i>	8 GB	4580	2988	355.49	4269

From the Table 6, large amount of log query will make system performance drop nearly 20%.

Note: Log query impact on performance:

- Querying logs requires the system CPU/Memory resources. The larger the amount of logs, the higher CPU usage is.
 - The log query does not require much physical memory. It uses Linux cache memory.
 - In average IWSVA completes one-time log queries on 1.1-billion logs in 5 seconds; after completion the CPU/physical Memory used by log query will be freed up.
 - The maximum system performance impact for the log query is 20%.
-

Sizing at a Glance – Log entries V.S. Disk size

Table 7 provides recommendations for required minimum disk space V.S. Log entries

Table 7 IWSVA 6.0 Sizing at a Glance with disk space V.S. Log entries

Log Entries	Disk Space (GB)	Average Disk Space per million log entries (MB)
3,000,000	0.6 GB	206 MB
110,000,000	11 GB	103 MB
2,200,000,000	134 GB	63 MB
4,400,000,000	256 GB	61 MB

IWSVA6.0 log storage mechanism makes non-linear relation between log entries and disk space. In general, the average disk space required will be reduced with log entries increasing.

Note: For simple disk size calculation with respect to log entries, the average disk size is **140MB** per million log entries.

Sizing at a Glance – Reporting

If the total log entries is less than 100 million, generating each single report type takes 10 seconds or less. Even selecting all report types the report can be generated in 3 minutes.

If the total log entries are greater than 1.1 billion, generating the most of single report type takes 1 minute or less. For the two “Top N sites visited” and “Top N sites by bandwidth” report types each takes about 50 minutes or less.

The screenshot displays the 'Report By' and 'Report Types' sections of the IWSVA reporting interface. Under 'Report By', 'All Users' is selected. Under 'Report Types', 'All Reports' is expanded, and 'Internet Security' is checked. A list of 14 report types follows, each with a checked checkbox, a 'Top N' value of 99, a description, and a bar chart icon. The reports are: Top N Malware Spyware Detection, Top N Botnet Detection, Top N Document Exploit APT Detection, Top N Custom Defense APT Blocking, C&C Contact Alert Count by Date, Top N C&C Address, Top N Users Hosts Detected by C&C Contact Alert, Top N Groups Detected by C&C Contact Alert, Top N Malicious Sites Blocked, Top N Users Blocked by Malware Spyware, Top N Users Blocked by Malicious Sites, and Top N Groups Blocked by Malware Spyware.

Report Type	Top N	Icon
Top N Malware Spyware Detection	99	Bar Chart
Top N Botnet Detection	99	Bar Chart
Top N Document Exploit APT Detection	99	Bar Chart
Top N Custom Defense APT Blocking	99	Bar Chart
C&C Contact Alert Count by Date		Bar Chart
Top N C&C Address	99	Bar Chart
Top N Users Hosts Detected by C&C Contact Alert	99	Bar Chart
Top N Groups Detected by C&C Contact Alert	99	Bar Chart
Top N Malicious Sites Blocked	99	Bar Chart
Top N Users Blocked by Malware Spyware	99	Bar Chart
Top N Users Blocked by Malicious Sites	99	Bar Chart
Top N Groups Blocked by Malware Spyware	99	Bar Chart

We recommend scheduling big reports in system idle hours, such as midnight, to avoid the resource conflict with content scanning functions.

Calculating the Number of IWSVA Servers Required

This procedure allows administrators to calculate the number of IWSVA servers required for the deployment based on the number of total users in the organization.

Step 1: Obtain the Required Data for the Customer Environment

At a minimum, the following information is required to size a customer environment:

- User Population Number

The following items are recommended to collect either for more accurate sizing:

- Peak number of concurrent user
- Peak throughput (Mbps)
- Caching percentage (if present)

Sizing variables are defined in Table 14. Obtain from the customer environment as many of the variables listed in Table 14 as practical and write them down on a worksheet to be used in the calculations or in the Sizing Calculator.

Note: To ensure proper sizing, Trend Micro recommends that customers use peak loads (the highest number of active users and peak throughput) when calculating the number of required IWSVA servers.

All the calculations are based on performance data from IWSVA 6.0.

Table 8 *Environment Variables for IWSVA Sizing*

Name	Variable	Description
Number of Users with Internet Access	USER_POPULATION	The total number of users with Internet access that will be supported by the IWSVA deployment.
Number of Users Accessing the Internet Simultaneously	%_CONCURRENCY	The percentage of the USER_POPULATION who are actively making an internet request (clicking a link in a web browser) If unknown, it is common to choose 20% (0.20) concurrency when sizing. For environments with more active users, increase the concurrency percentage to a value such as 0.33 or higher. (%_CONCURRENCY=0.20)
Peak Bandwidth Required	THROUGHPUT_TOTAL	The amount of HTTP traffic passing through the gateway in Megabits per second (Mbps). If unknown, it is common to choose a value of 75% available Internet bandwidth. For example, an organization with a T3 connection would have 44.74 Mbits per second. For this organization THROUGHPUT_TOTAL = 0.75 x 44.74 = 33.6.
Connections per user	CONNECTIONS_USER	The number of HTTP connections per active user. If unknown, the default value is 3.5. For environments with users accessing many Web 2.0 sites, increase the number of connections per user to 5 or higher. (CONNECTIONS_USER=3.5).

Name	Variable	Description
Caching percentage	%_CACHE	<p>If a caching solution is present, the percentage of caching occurring for the environment is required.</p> <p>If unknown yet caching is present, it is common to choose 25% caching (%_CACHE = 0.25) for the environment.</p> <p>If no caching is present, choose %_CACHE = 0.</p>
Maximum log entries per month	MAX_LOG_MONTH	The maximum log number one month in customer's environment
Total maximum log entries	MAX_LOG_TOTAL	The maximum total log number in customer's environment
Number of days to keep logs	DAYS_OF_LOG	<p>The number of days that log can kept on IWSVA</p> <p><i>(IWSVA default use 30 day as default value)</i></p>
Number of Internet accesses per user per day	NUM_OF_ACCESS	<p>An estimated number of one user's accesses to the internet in one day.</p> <p><i>(If you don't know it, we recommend 1300)</i></p>

Step 2: Determine the Number of Required Servers

Identical hardware configuration is assumed for all servers in a multi-server deployment. After calculation, the customer can determine which server type fits their needs the best and use the recommended number of servers for that specific hardware configuration.

Note: It is important to note that only one type of server is used for sizing. Do not add the results from the options below; simply choose one of the results for the sizing.

Option 1: Number of 2 x 4-core servers (3.3 GHz Intel E3-1240, 8 GB RAM)

Using the variables obtained or assumed from Table 8 and LDAP/HTTPS' impact ratio, calculate the number of servers required:

$$\text{Number_Servers} = \frac{(\text{CONNECTIO NS_USER} \times \%_CONCURRE NT \times \text{USER_POPUL ATION})}{3,100} \times (1 - \%_CACHE)$$

Round **up** the number of IWSVA servers to the nearest whole number.

Option 2: Number of 2 x 6-core server (2 x 1.9 GHz Intel E5-2420, 16 GB RAM)

Using the variables obtained or assumed from Table 8 and LDAP/HTTPs' impact ratio, calculate the number of servers required:

$$\text{Number_Servers} = \frac{(\text{CONNECTIO NS_USER} \times \% _ \text{CONCURRE NT} \times \text{USER_POPUL ATION})}{5,400} \times (1 - \% _ \text{CACHE})$$

Option 3: Number of 2 x 8-core server (2 x 2.2 GHz Intel E5-2620, 24 GB RAM)

Using the variables obtained or assumed from Table 8 and LDAP/HTTPs' impact ratio, calculate the number of servers required:

$$\text{Number_Servers} = \frac{(\text{CONNECTIO NS_USER} \times \% _ \text{CONCURRE NT} \times \text{USER_POPUL ATION})}{7,500} \times (1 - \% _ \text{CACHE})$$

Round **up** the number of IWSVA servers to the nearest integer number.

IWSVA 6.0 Sizing Example

For a deployment with:

- USER_POPULATION = 10,000
- %_CONCURRENT = 0.20
- CONNECTIONS_USER = 3.5
- %_CACHE = 0.20
- THROUGHPUT_TOTAL=180 Megabits per second

Who desires sizing using the following existing server configuration:

- Two CPU (two, dual-core, 3.0 GHz Intel Xeon 5160 CPU's)
- Memory per server is 4 GB

The number of required servers is as follows:

$$\text{Number_Servers} = \frac{(3.5 \times 0.20 \times 10,000)}{1,700} \times (1 - 0.20) = 3. \text{ With rounding up, this equals 3.0}$$

For this network, three (3) servers are required to meet their needs to ensure scanning capacity meets environmental conditions.

Note: If throughput information is also available, the throughput capabilities of the solution should be compared to the environment needs prior to making a recommendation.

Simply compare the THROUGHPUT_TOTAL variable (if available) to the calculated throughput for the recommended solution. The calculated throughput should be greater than the THROUGHPUT_TOTAL variable. If it is not, the number of recommended servers should be adjusted accordingly.

Conclusion: For customers using IWSVA version < 6.0, if no need to enable new features and just keep legacy features enable, customers only need to adopt memory = num of CPU threads x 0.5 + 3.

If they want to enable new features in IWSVA 6.0, suggest use calculation above to adopt hardware resource.

For new customers using IWSVA 6.0, just use calculation above to adopt hardware is enough.

Step 3: Determine the Storage

The administrators can use the information described below to calculate the required disk space for log storage.

For storage type, we recommend a fast disk subsystem. (For example, 15K RPM SAS disk in RAID 1+0 configurations)

Note: As the number of users and events increases, the DISK_IO will be more and more busy, it becomes more important to use a fast disk subsystem to increase Log performance

Dimensions 1: How much disk space is needed if I have XXX log entries?

From above sizing about Log entries V.S. Disk size, the average log size is 140MB per million log entries, so we can calculate needed disk size (GB):

$$Disk_Size(GB) = \frac{MAX_LOG_TOTAL}{1,000 \times 1,000 \times 1,024} \times 140$$

IWSVA 6.0 Storage Sizing Example

For a customer with:

- MAX_LOG_TOTAL = 2,200,000,000

The number of required disk size is as follows:

$$Disk_Size = \frac{2,200,000,000}{1,000 \times 1,000 \times 1,024} \times 140 = 300.78125 \text{ with rounding up, this equals 301 GB}$$

Dimensions 2: How much disk space is needed if I have XXX users?

$$Disk_Size(GB) = \frac{NUM_OF_USER \times NUM_OF_ACCESS \times DAYS_OF_LOG}{1,000 \times 1,000 \times 1,024} \times 140$$

IWSVA 6.0 Storage Sizing Example

For a customer with:

- *USER_POPULATION* = 10,000

Following is IWSVA default value:

- *NUM_OF_ACCESS* = 6,500 (**default**)
- *DAYS_OF_LOG* = 30 (**default**)

The number of required disk size is as follows:

$$\text{Disk_Size} = \frac{10,000 \times 6,500 \times 30}{1,000 \times 1,000 \times 1,024} \times 140 = 266.6015625 \text{ with rounding up, this equals 267 GB}$$

Dimensions 3: How many days of log can be kept on IWSVA?

The default value for log days on IWSVA is 30 days. Sometimes user wants to modify the log days on IWSVA to fits the hardware. If the user has known the disk space available for logs and the number of users, then he can get the number of days to keep log on IWSVA by one of the following formulas:

$$\text{DAYS_OF_LOG} = \frac{\text{Disk_Size} \times 1,000 \times 1,000 \times 1,024}{\text{NUM_OF_ACCESS} \times \text{NUM_OF_USER} \times 140}$$

IWSVA 6.0 Storage Sizing Example

For a customer with:

- *USER_POPULATION* = 5,000
- *Disk_Size* = 128 (GB)

Following is IWSVA default value:

- *NUM_OF_ACCESS* = 6,500 (**default**)
- *LOG_LEVEL* = 5 (**default**)

The days of log can be kept is as follows:

$$\text{DAYS_OF_LOG} = \frac{120 \times 1,000 \times 1,000 \times 1,024}{6,500 \times 5000 \times 140} = 27.00659340659341 \text{ with rounding up, this equals 27 days}$$

Step 4: Determine the Standalone log server

Standalone log server can receive log from multiple IWSVA instances. Using standalone log server can also avoid IWSVA content scanning performance impact induced by log query and report generation operation.

The performance of log server is mainly determined by the handling capacity of log entries generated per day in the environment. The number of log sources registered has negligible impact on performance.

Recommended Configuration:

- ***If the log entries are greater than 10 million per day:***

CPU: 8 Threads

MEM: 16GB

Disk Size: Please refer to the section “Determine the storage”

*If log entries are greater than 1 billion per month, it is recommended to keep only **2-month** logs for optimal performance.*

- ***If the log entries are less than 3 million per day:***

CPU: 4 Threads

MEM: 8GB

Disk Size: Please refer to the section “Determine the storage”

Note:

1. The fast CPU and Disk IO can help on the log handling capacity; the more CPU threads can help query logs and generate reports in parallel.
2. IWSVA as log server can also be installed on a VMware ESXi virtual environment, and can leverage the VMware’s redundancy and fault tolerant features to create a robust and scalable deployment.

Appendix A

How Tests Were Conducted

Product performance was determined based on a workload where each active user accesses 12 Web sites sequentially. This workload was deemed representative of that of an actual enterprise. The range of object sizes (.jpg, .png, .css, .gif, and .js) ranged between 9 and 174 KBytes, with .htm pages ranging between 3 and 143 KBytes. Think time was maintained at 5 seconds, making this a test of moderately aggressive Internet surfing behavior. For the application control’s impact for IWSVA 6.0, we use one scenario, which is to use pure HTTP traffic to check application control’s impact for features.

The ICAP mode deployment with external caching servers in general greatly improves the capacity of the environment by reducing the amount of network communication to the IWSVA server.

Note: IWSVA also includes a **Web Reputation** feature. This feature relies on DNS queries to Trend Micro data centers for each new URL request. Reputations are cached for a period of 35 minutes by default and new reputation requests for that URL are provided without the need for additional queries.

This Sizing Guide assumes that the customer environment has sufficient DNS infrastructure to handle the query load that results from deploying one or more IWSVA units.

What Configuration Changes do to Sizing

Configuration changes to IWSVA affect sizing in a number of ways. These impacts are summarized below:

Reporting

- Real-time reports can take a significant amount of time to complete in high-workload environments if there are inadequate amounts of free CPU to process each request. For this reason, it is recommend to create schedule report for non-peak workload periods.
- Environments requiring efficient real-time reporting during high-workload periods are advised to size their servers for less than 100% CPU utilization to keep the end-user and administrative experiences positive. Alternatively, the standalone log server can be deployed on a separate host/vm to handle all log processing and reporting actions. This frees up the IWSVA instance to dedicate its resources to processing traffic.
- Using high-performance RAID arrays with fast hard disks will improve performance significantly.

Caching

Using caching in an environment improves system performance. With a properly sized ICAP 1.0 solution, the capacity of the environment increases proportionally to the cache percentage:

- A 25-percent cache allows each server to increase capacity by a factor of 1.3.
- A 50-percent cache allows each server to increase capacity by a factor of 2.
- A 75-percent cache allows each server to increase capacity by a factor of 4.
- A 90-percent cache allows each server to increase capacity by a factor of 10.

These performance factors are based on an off box external ICAP server.

Performance Criteria for Tests

Trend Micro conducted the tests with the requirement that all test results and sizing recommendations meet the following conditions:

- Hosts and servers have zero TCP Connection failures
- Hosts and servers have zero HTTP Transaction failures
- Hosts must experience an average page load time of no more than 2000ms (2 seconds)

Although the IWSVA servers can provide more connections and transactions than listed in the sizing tables, the page load latency will be above 2 seconds and will not reflect real-world expectations where users expect faster Internet response times.

Scalability and Accuracy

The performance ability of IWSVA depends on the quantity and type of CPU being used and also the feature sets enabled. Higher MHZ and more CPU can bring better performance results.. However, when the Application Control feature is enabled, performance increases flatten out above 12 CPU cores. Since 12 cores can support up to 12,000 users in a single IWSVA instance, it is unlikely that a customer would have more than that many users' traffic flowing through a single instance. Segmenting the network to have Internet traffic from separate sub-nets flowing to multiple concurrent IWSVA instances is one way to address this potential scalability issue.

The testing procedure and methodology used in this report is accurate, reproducible and well documented. The results are precise by ± 5 percent.

Hardware Tested

Tables 16-22 provide details of the hardware used in this Sizing Guide.

Table 9 4 thread Server

Component Type	Value
Chassis	Dell R210II
CPU	Intel® Xeon® Processor E3-1240
CPU Speed	3.30 GHz
Cores per CPU	4
Threads per cores	2
Number of physical CPU	1
Total Threads	8
Memory	8 GB
Storage	SATA SSD, SAS
Network	Broadcom BCM 5716

Table 10 24 thread Server

Component Type	Value
Chassis	Dell R420
CPU	Intel® Xeon® Processor E5-2420
CPU Speed	1.9 GHz
Cores per CPU	6
Threads per cores	2
Number of physical CPU	2
Total Threads	24
Memory	16 GB
Storage	SAS SSD, SATA SSD, SAS
Network	Broadcom® NetXtreme 5709c

Table 11 32 thread Server

Component Type	Value
Chassis	Dell R720
CPU	Intel® Xeon® Processor E5-2660
CPU Speed (Total Allocated)	2.2 GHz
Cores per CPU	8
Threads per cores	2
Number of physical CPU	2
Total Threads	32
Memory	24 GB
Storage	SAS SSD, SATA SSD, SAS
Network	Broadcom® BMC5709C

Glossary

Active Users – The number of actual users requesting Web content through an HTTP Web browser (such as Microsoft Internet Explorer) at any time.

Connection Latency – The amount of time between the user's first click in a Web browser until the time data begins appearing on the screen.

Default Configuration – The default configuration of IWSVA is with antivirus, Web-reputation, URL filtering, and Anti-Botnet (enabled in IWSVA 6.0 by default) active.

HTTP 1.1 Connection – A method that enables the use of one connection to send or receive multiple HTTP requests or responses. HTTP 1.1 allows multiple requests to be made through a single connection.

Requests per second – The rate at which HTTP objects (for example .jpg, .gif or .html files) are requested and processed.

Think Time – The time between browser clicks for an active user.

Throughput – The amount of digital data per time unit that is delivered over a physical or logical link, or that is passing through a gateway scanning device. This is expressed as either Bytes per second or bits per second (8 bits = 1 Byte).

Total Page Download Latency – The average total time to download a workload-specific Web site after initial connection.

User Population – The total number of users with Internet access to be supported by the IWSVA deployment.

About Trend Micro Incorporated

Trend Micro Incorporated, a global leader in Internet content security, focuses on securing the exchange of digital information for businesses and consumers. A pioneer and industry vanguard, Trend Micro is advancing integrated threat management technology to protect operational continuity, personal information, and property from malware, spam, data leaks and the newest Web threats. Trend Micro's flexible solutions, available in multiple form factors, are supported 24/7 by threat intelligence experts around the globe. A transnational company, with headquarters in Tokyo, Trend Micro's trusted security solutions are sold through its business partners worldwide.

For more information, please visit www.trendmicro.com.