

Antivirus Protection for the Sun StorageTek™ 5000 NAS Appliance

Tim Thomas
Sun Microsystems
August 2007

2007 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, CA 95054 USA

All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD appliances, licensed from the University of California.

Sun, Sun Microsystems, Solaris, Sun StorageTek, Sun StorEdge, and the Sun logo are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries.

UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company, Ltd.

Windows is a registered trademark of Microsoft Corporation in the United States and other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun's Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK Guise and otherwise comply with Sun's written license agreements.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-1987), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a). DOCUMENTATION IS PROVIDED AS IS AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS HELD TO BE LEGALLY INVALID.

Table of Contents

- Introduction.....1**
- What Are Computer Viruses?.....2**
- How Does Antivirus Software Work?.....3**
- Overview of the Antivirus Scanning Process.....4**
- Managing Sun StorageTek 5000 NAS Appliance Antivirus Features.....6**
 - Configuring the NAS Antivirus Agent6
 - Exclusion Lists.....7
 - Maximum File Size to Be Scanned.....7
 - Scanning Application Files and Directories.....7
- Managing Scan Engine Policies.....8**
 - Scan Policy.....8
 - Additional Scan Engine Settings.....8
 - Updating Virus Definitions.....8
 - Certified Antivirus Software for the Scan Engines.....9
- Advanced Topics.....10**
 - Logging and Notification of Virus Detection.....10
 - Managing Quarantined Files.....10
 - Scanning Compliance Volumes.....11
 - What Happens If No Scan Engines Are Available?.....11
- Performance and Availability Considerations.....12**
- External Antivirus Scanning.....13**
 - Antivirus Software on Client Systems.....13
 - On-Demand Full System Scans.....13
- Frequently Asked Questions.....14**
- Summary.....16**
- For More Information.....17**

Introduction

The Sun StorageTek™ 5000 NAS Appliance family has features that protect Common Internet File System (CIFS) clients from viruses in files stored on the appliance. These features are a standard part of Sun StorageTek NAS OS, the Sun StorageTek 5000 NAS Appliance's Operating System, and are known as the NAS Antivirus Agent (NAS AVA).

This paper discusses the features of the NAS AVA as implemented in Sun StorageTek NAS OS 4.21.

What Are Computer Viruses?

A virus is a small piece of software that piggybacks on real programs. For example, a virus might attach itself to a spreadsheet program. Each time the spreadsheet program runs, the virus runs too, and it has the chance to reproduce by attaching to other programs or possibly wreak havoc by, for example, destroying data or sending hundreds of emails.

Computer viruses mostly target Microsoft operating systems however computers running other operating systems can be infected directly or affected indirectly by viruses.

Here are some examples of computer viruses and how they are transmitted:

- **Email Viruses** - An email virus moves around in email messages as an attachment and is activated when a naive user opens it. The virus usually replicates by automatically mailing itself to dozens of people in the victim's email address book.
- **Worms** - A worm is a small piece of software that uses computer networks and security holes to replicate itself. A copy of the worm scans the network for another machine that has a specific security hole. It copies itself to the new machine using the security hole, and then starts replicating from there, as well.
- **Trojan horses** - A Trojan horse is simply a computer program. The program claims to do one thing (it may claim to be a game) but instead does damage when you run it (it may erase all the files from a hard disk for example). Trojan horses have no way to replicate automatically, though they may travel around as email attachments.

How Does Antivirus Software Work?

Antivirus software searches files for infections using sophisticated search techniques and a catalog of known virus definitions. The virus definitions are updated frequently, typically automatically, by the virus scanner using a subscription service.

There are two common approaches to virus scanning:

On-Demand Scan - The On-Demand scan method searches all or part of a file system when requested, checking files of selected types and/or modification dates. This is typically performed at scheduled intervals.

Real-Time Scan - The Real-Time scan method scans files as they are accessed. Files can be scanned when they are opened and/or after they are closed.

On-Demand Scanning is simple, but has the disadvantages that files recently updated might not be scanned before use and full scans can take a long time to complete on large systems.

The real-time scan method has the benefit of ensuring that files are scanned with the latest virus definitions before being used. This approach is more effective at detecting viruses before they are able to compromise data and has the additional benefit of not generating the very heavy I/O loads of On-Demand scans. This is the method supported by the NAS AVA.

If required, On-Demand scans of the Sun StorageTek 5000 NAS Appliance can be performed by a CIFS client scanning shares off the appliance.

Overview of the Antivirus Scanning Process

The NAS AVA supports real-time virus scanning of files stored on a Sun StorageTek 5000 NAS Appliance when accessed by CIFS clients. A file accessed over CIFS is a candidate to be scanned:

- When it is opened;
- When it is closed, IF it was modified.

If the NAS AVA determines that a file needs to be scanned, it transfers the file to an external server running antivirus software; this server is known as a Scan Engine. The NAS AVA communicates with the Scan Engine using the Internet Content Adaptation Protocol (ICAP). ICAP is becoming an industry standard for communicating with Scan Engines. You can read about ICAP at <http://www.i-cap.org>.

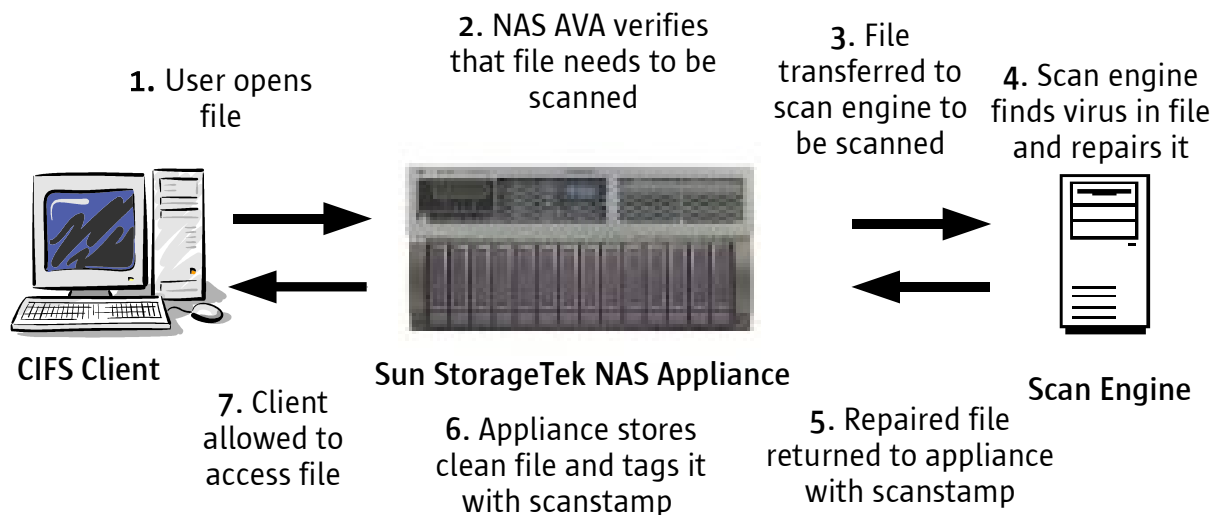


Figure 1: Scan-On-Open workflow for an infected file with Scan Engine set to repair the file. Sun StorageTek 5320 shown.

When a Scan Engine receives a file it scans it and then returns a scan status to the NAS AVA which it can act upon.

- If the file was not infected, the NAS AVA will grant access to the file.
- If the file was infected, the NAS AVA will either quarantine the file (which denies access to CIFS clients) or replace it with a repaired file sent to it by the Scan Engine. Repairing a file is the process of removing a virus, it is also known as curing and cleaning.

When the scan process has finished, the file is tagged by the NAS AVA with its scan status; this is known as a scanstamp. The purpose of scanstamps is to prevent files known to be clean from being scanned again. The scanstamp remains current so long as the virus definitions on the Scan Engines are not updated.

The next time an operation is requested on the file that makes it a candidate to be virus scanned the NAS AVA checks to see if the file's scanstamp is current and that the file has not been modified or renamed since it was last scanned: If these conditions are met, the operation is allowed to proceed without scanning the file; if not, then the file must be scanned before allowing the operation.

Managing Sun StorageTek 5000 NAS Appliance Antivirus Features

Configuring the NAS Antivirus Agent

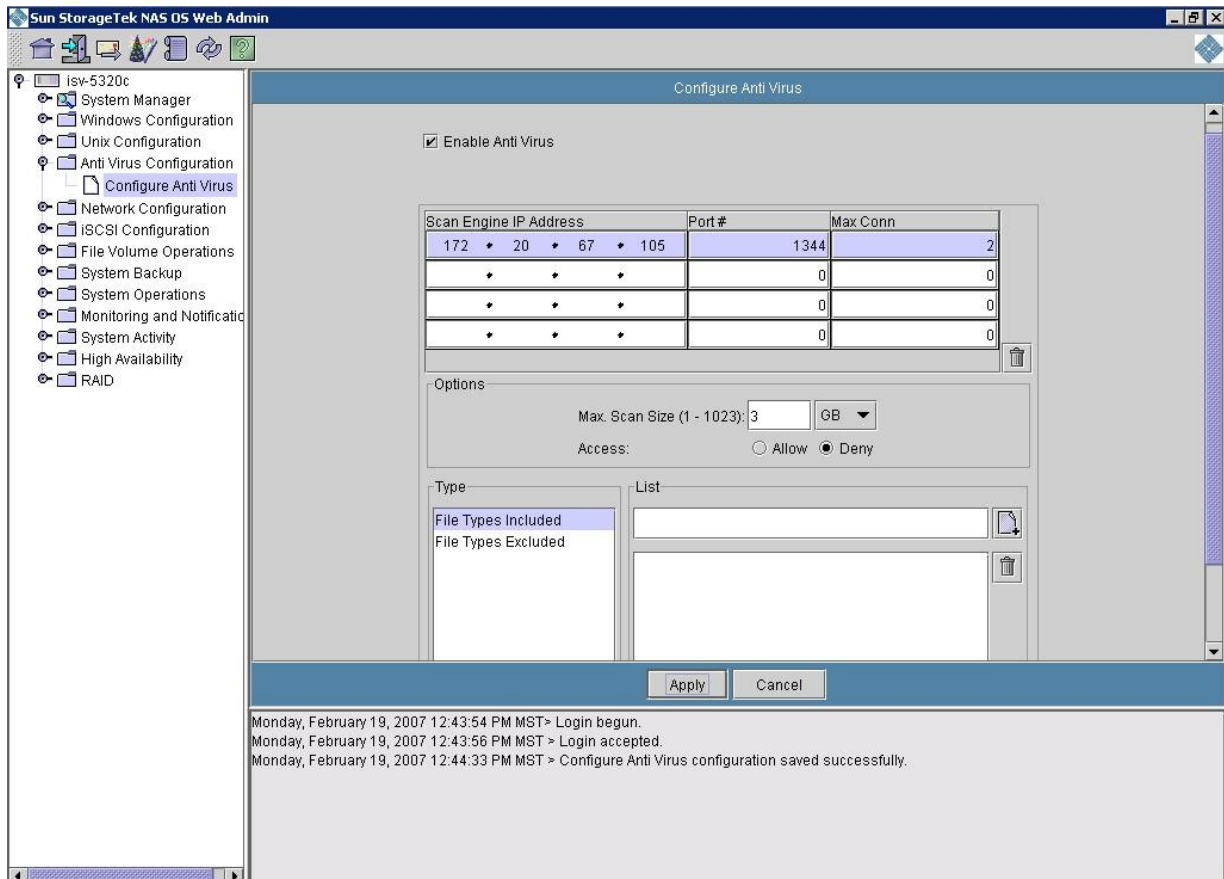


Figure 2: Antivirus protection set-up screen for the NAS AVA

The NAS AVA is a standard part of Sun StorageTek NAS OS; it does not have to be separately licensed. It is enabled by ticking the "Enable Antivirus" box, as shown in Figure 2.

The details of at least one Scan Engine must be entered before the configuration details can be saved. Up to four scan engines are supported per Sun StorageTek 5000 NAS Appliance and scan requests are round-robin load-balanced between them. *Max Conn* sets the number of concurrent scan requests allowed against the Scan Engine. *Max Conn* defaults to two, but can be much larger; the maximum setting is a function of the Scan Engine's antivirus software.

Exclusion Lists

By default all files on a Sun StorageTek 5000 NAS Appliance accessed via the CIFS protocol are candidates for virus scanning. Configuring Exclude Lists reduces the number of files that need to be scanned by focusing virus scans on the files, shares and volumes viewed to be at risk. Exclusions can be made at various levels of granularity:

- **By File Type:** Files can be excluded (or included, if that better suits requirements) by file type.
- **At the Share Level:** Whole Shares can be excluded. This takes precedence over File Type Exclusion.
- **At the Volume Level:** Whole Volumes, and so all the shares taken off them, can be excluded. This takes precedence over Share and File Type Exclusion.

Organizations using the antivirus scanning features of the Sun StorageTek 5000 NAS Appliance must decide what file types, shares and volumes should be excluded from scans based on their own security policies.

Maximum File Size to Be Scanned

The administrator must set the maximum size of file that should be sent to a Scan Engine, and whether to Deny or Allow access to a file that violates the policy, see Figure 2. This value should be set to less than or equal to the maximum sized file supported by the Scan Engine. Most Scan Engines support a maximum file size of 2 GB.

Scanning Application Files and Directories

Before configuring the NAS AVA to scan files that are part of an application it is important to consult the application vendor to find out what their recommendations are; they may be that some or all of the application's files are excluded from real-time scanning or that no virus scanning is required at all. Applications like email systems usually have their own integration points with antivirus software.

If an application's directory structure needs to be excluded from scanning then it must be placed on a share that is excluded from virus scans. If only parts of the application need to be excluded from scanning then the application can be spread over multiple shares or volumes and some of those excluded as required.

Managing Scan Engine Policies

Scan Policy

The NAS AVA's policies determine if a file should be sent to a Scan Engine to be scanned. The policies set on a Scan Engine determine how a file is processed when it is received. Typically, a number of Scan policies are available:

- **Scan Only:** No attempt is made to repair infected files. All infected files will be quarantined.
- **Scan and Delete:** Delete all infected files. The NAS AVA cannot delete a file: If this policy is selected, all infected files will be quarantined.
- **Scan and Repair:** Attempt to repair a virus infected file. Quarantine the file if it cannot be repaired.
- **Scan and Repair or Delete:** As above, but delete the file if it cannot be repaired. There are two cases to consider:
 1. **The file is a container:** The Scan Engine software will, if necessary, delete individual files that cannot be repaired when repairing containers such as zip files. This is done on the Scan Engine, so is beyond the control of the NAS AVA.
 2. **The file is not a container:** The NAS AVA cannot delete a file: if this policy is selected the infected file will be quarantined.

Additional Scan Engine Settings

Additional settings that can be made on Scan Engines include:

- How long (in seconds) to work on a file before giving up.
- How to handle complex container/archive files.
- Maximum file size to scan.
- How to handle encrypted files (e.g. zip files with passwords).

It is possible to define File Type Exclude Lists on a Scan Engine, but it is more efficient to define them on the Sun StorageTek 5000 NAS Appliance as that reduces the number of scan requests and minimizes network traffic.

Updating Virus Definitions

Antivirus software running on a Scan Engine updates its virus definitions autonomously. Deciding on the frequency of checking for updates is part of configuring the software on the Scan Engine.

Certified Antivirus Software for the Scan Engines

Antivirus software from Symantec, McAfee, Trend Micro and Computer Associates is certified with the Sun StorageTek 5000 NAS Appliance. Support for McAfee was introduced with Sun StorageTek NAS OS 4.21 Maintenance Update 1.

Consult the Sun StorageTek 5000 NAS Appliance support matrix and/or Release Notes for exact details of supported antivirus software and appliances.

Advanced Topics

Logging and Notification of Virus Detection

When a virus is detected the event is logged in a number of places:

- An entry is added to the NAS OS system log. This records: the name of the infected file; the name of the virus; how it was processed, which will depend on the Scan Engine software policies.
- A similar entry is made in a file called *virus.log* in the *.quarantine* directory at the root of the volume where the file is stored.
- On the Scan Engine in the log maintained by the antivirus software.

If required, some antivirus software products can be configured to send email and/or SNMP traps when an infected file has been found. For large organizations higher level software tools are available from the antivirus software suppliers to aggregate these notifications.

Managing Quarantined Files

To quarantine a file, the NAS AVA sets an attribute on the file which results in open, rename, read and execute access being denied. The file is still visible in its original location and can be deleted, but any attempt to copy, move, modify, execute or rename the file is denied, regardless of the protocol used.

As well as a GUI, the Sun StorageTek 5000 NAS Appliance has a command line interface that is accessed via *telnet*. A quarantined file can be deleted or removed from quarantine using NAS OS commands. See the Sun StorageTek 5000 NAS Appliance documentation for details.

Files Accessed via NFS

Scanning of a file is not initiated on NFS access. This is normal for NAS Appliances as CIFS clients running Microsoft Windows are the usual vector for viruses.

If a file is modified over NFS it will be scanned if it is later opened by a CIFS client. Quarantined files are not accessible over NFS.

Scanning Compliance Volumes

WORM files are read-only; if a WORM file is found to be infected with a virus no attempt is made to repair the file. The scan status sent back to the Sun StorageTek 5000 NAS Appliance will indicate that the file is infected and access to the file will be denied. The file cannot be quarantined or repaired.

For compliance records, it is recommended that a process be put in place whereby IT informs the organization's Records Manager of any retrieved file that contains a virus. Depending on the Enterprise Content Management application being used at the customer site, it may then be possible to repair the file (on a 'safe' PC) and store it back in the application as a new, duplicate record, with the same metadata and retention as the original. Audit notes should be added to the new and old records by the records manager to explain the duplication and enable audit. The old (original) record would be given new security to prevent its search/retrieval by regular users, and would be destroyed according to its original retention schedule, along with the duplicate. In this scenario, the original infected file is still available as evidence if absolutely required, but regular users are shielded from its damaging effects since it is hidden from them.

What Happens If No Scan Engines Are Available?

If no Scan Engines are available, due to a network or other hardware failure for example, then access to some files will be denied: If a file has never been scanned or has been modified/renamed since last scanned, access is denied; if a file has been scanned before and virus definitions have not been updated since it was scanned, access is allowed.

Performance and Availability Considerations

The Real-Time scanning of files when they are opened results in CIFS clients seeing latency of anything from a few milliseconds to a number of seconds.

Scanning a modified file when it is closed greatly reduces the probability that a file will need to be scanned when it is opened. This has been implemented as an asynchronous operation on the Sun StorageTek 5000 NAS Appliance, so it does not add latency to the close operation. If the file is opened again before the scan has completed the open operation must wait.

Overall performance overhead caused by virus scanning will depend very much on the workload placed on the Sun StorageTek 5000 NAS Appliance and how the NAS OS, Scan Engines and network are configured. Here are a number of factors to consider:

- **Scan Engine Specification:** Check the antivirus software vendor's documentation for the required Scan Engine hardware specification. The minimum specification is usually 1 GHz CPU, 1 GB memory and Gigabit Ethernet.
- **Have at least two Scan Engines:** This has performance benefits, as scan requests are load-balanced across all available Scan Engines, but also affects availability: if only one Scan Engine is present, and it fails, then access to some files will be denied as discussed previously. Scan Engines can be added, up to a maximum of four, and removed without disruption to I/O so long as at least one is kept live at all times.
- **Tune Number of Concurrent Requests allowed Per Scan Engine:** Tune *Max Conn* to make sure that the scan engines are being well utilized. The default of two is too low.
- **Exclusion Lists:** Make use of Exclusion Lists to minimize the number of files that need to be scanned.
- **Scanning Large Files:** Restrict the scanning of large files. At the very least, ensure that the NAS AVA does not send files to a Scan Engine larger than the maximum size it can handle.
- **Scanning Policy:** Consider choosing a policy on the Scan Engine to scan, but not repair files. This will speed up the scanning process Vs more complex scan policies.
- **Network Speed:** Ensure that the network connections between the Scan Engines and the Sun StorageTek 5000 NAS Appliance are running at Gigabit speed. This speeds the transfer of files between the Appliance and the Scan Engines.
- **Network Latency:** Keep network latency between the Appliance and the Scan Engines to a minimum.

Note that most of the NAS AVA and Scan Engine antivirus software settings are run-time configurable.

External Antivirus Scanning

Antivirus Software on Client Systems

Files written to a Sun StorageTek 5000 NAS Appliance by a CIFS client will be virus scanned by the NAS AVA. However, if the local file systems on a client are unprotected and a virus is activated on a client, it can damage local files.

To prevent this from happening, antivirus software should be deployed as normal on clients to protect local drives from infection when using the Antivirus features of a Sun StorageTek 5000 NAS Appliance.

On-Demand Full System Scans

Rarely accessed files may go for long periods without being scanned. These files may be infected and this infection can still be spread via backup tapes and appliance-to-appliance replication. To make sure that all of the files on a Sun StorageTek 5000 NAS Appliance are periodically checked for infection some organizations may wish to perform full system On-Demand scans on a regular basis.

The antivirus features of Sun StorageTek 5000 NAS Appliance do not support On-Demand scans. On-Demand scans have to be performed by a CIFS client of the Sun StorageTek 5000 NAS Appliance with a local installation of antivirus software that is configured to scan files on the appliance via CIFS shares.

The same caveats apply to On-Demand full system scans of application directory structures as for Real-Time scans.

On-Demand scans do not update scanstamps, so even if a file is found to be clean during an on-demand scan, it will still be real-time scanned by the NAS AVA when accessed again.

Frequently Asked Questions

Q1: How many Scan Engines are supported per Sun StorageTek 5000 NAS Appliance?

A1: Up to four are supported. A minimum of two are recommended.

Q2: What happens if all of the Scan Engines become unavailable?

A2: If a file has never been scanned, or has been modified or renamed since last scanned, access is denied; if the file has been scanned before, and virus definitions have not been updated since that scan, access is allowed.

Q3: What protocol does the Sun StorageTek 5000 NAS Appliance use to communicate with Scan Engines?

A3: The Internet Content Adaptation Protocol (ICAP) is used. This is becoming an industry standard for communicating with Scan Engines. You can read about ICAP at www.i-cap.org.

Q4: Do Scan Engines ship with the Sun StorageTek 5000 NAS Appliance?

A4: Scan Engines, and the antivirus software to run on them, must be purchased separately from the Sun StorageTek 5000 NAS Appliance.

Q5: Where can I find the specifications for the Scan Engines?

A5: In the documentation for the antivirus software you select. A typical configuration is a minimum of a 1 GHz CPU, 1 GB memory and a Gigabit ethernet adapter. Depending on the product, the Scan Engines can be servers running Microsoft Windows, the Solaris™ Operating System, or Linux.

Q6: Are whole files transferred between the Sun StorageTek 5000 NAS Appliance and the Scan Engines?

A6: It depends: some scan engines have a preview feature that means that only part of a file needs to be sent initially, the Scan Engine then decides if the rest of the file is required.

Q7: How will I know if an infected file has been found?

A7: When a virus is detected an entry is added to the system log that records the name of the infected file, the name of the virus, and what disposition was selected for the file i.e. was it repaired or quarantined. A similar entry is made in a file called *virus.log* in the *.quarantine* directory at the root of the volume where the file is stored. Some Scan Engine software products can be configured to send email and/or SNMP traps when an infected file has been found.

Q8: What happens when a file is quarantined?

A8: An attribute is set on the file that denies open, rename, read and execute access. The file can still be seen, but it cannot be opened, renamed, and so on, as described above. The file can be deleted.

Q9: Is there any tuning I can do of the NAS AVA?

A9: See the Performance and Availability section of this document.

Q10: Do I still need antivirus software on client systems?

A10: Yes. Local drives still need to be scanned for viruses.

Q11: How do I scan NFS files on the Sun StorageTek 5000 NAS Appliance for viruses?

A11: Use full system On-Demand scans.

Q12: UNIX and Linux don't get viruses, right?

A12: Viruses are currently rare on UNIX and Linux but they do exist. Today, virus writers focus most of their efforts on Microsoft Windows and antivirus software vendors target the removal of those viruses.

Q13: If a file not infected with a virus is sent to be scanned by a Scan Engine is it sent back to the Sun StorageTek 5000 NAS Appliance along with the scan status?

A13: No, only the scan status is sent back.

Q14: How do I update the virus definitions that the antivirus software uses?

A14: Antivirus software running on a Scan Engine updates its virus definitions autonomously in a similar manner to the way that a personal copy of antivirus software running on a Windows PC updates itself. The Sun StorageTek 5000 NAS Appliance does not control this.

Q15: Can Scan Engines be shared between multiple appliances?

A15: Yes. Scan Engines can offer remote scanning services to multiple NAS appliances and compatible applications but care must be taken not to overload them. Note that in a Sun StorageTek 5000 NAS Appliance Cluster both heads share the same set of Scan Engines.

Summary

Computer viruses are a menace to the security of personal and corporate data. The antivirus features of the Sun StorageTek 5000 NAS Appliance provide a scalable solution that efficiently scans files for viruses in real-time. This is generally more secure and less time consuming and I/O intensive than traditional scheduled On-Demand scans.

For more information about Sun storage products go to <http://www.sun.com/storage>.

For More Information

Storage Administration hub on the BigAdmin System Administration Portal:

<http://www.sun.com/bigadmin/hubs/storage/>

Training:

Storage and Information Life Cycle Management Courses

<http://www.sun.com/training>

Sun Product Documentation:

<http://docs.sun.com>

Sun Wikis: <http://wikis.sun.com>

- BigAdmin wiki
- Storage Administration wiki

Support:

- Sun Services & Solutions:
<http://sun.com/solutions>
- Support Information From SunSolve:
<http://sunsolve.sun.com>