SANscreen Operational Guidelines:

Guidance and Best Practices

Michael Steinberg, NetApp
February 2010

TABLE OF CONTENTS

1 DOCUMENT PURPOSE

This document provides guidance and best practices to customers running NetApp's SANscreen software. This document is based on the combined experiences of NetApp's Professional Services and Global Support Center personnel who work with SANscreen on a daily basis.

1.1 Background Information

Like all enterprise software packages, SANscreen requires maintenance after the initial deployment is completed. Given the dynamic nature of today's SAN environments, there are usually a number of ongoing changes taking place within a customer's data center. Examples of such changes are:

- Firmware upgrades for switches/fabrics
- Firmware upgrades for arrays
- Management software upgrades
- Addition or subtraction of hardware from the SAN environment
    - Arrays
    - Switches
    - Hosts
- Addition or removal of storage from hosts

SANscreen is integration software. Although NetApp makes every effort to ensure compatibility between SANscreen and the various third party products SANscreen integrates with, there are instances when the above-mentioned changes could cause SANscreen Data Source issues to arise. In addition, ongoing administrative tasks should be tracked within SANscreen to ensure that the tasks are completed properly and that the change is correctly reflected in SANscreen.

1.2 Maintenance Tasks Discussed within This Document

This document provides information about the following maintenance tasks:

- Data Source maintenance
- Host identification
- Tape identification
- Annotation maintenance
- Service Assurance violation and policy maintenance
- Automation of SANscreen and Data Warehouse backups

2 DATA SOURCE MAINTENANCE

Data Sources are the most critical component when trying to maintain a SANscreen environment. Because Data Sources are the primary source of information for SANscreen, it is imperative to make sure Data Sources are maintained in a running state. In addition, it is important to know when to add new Data Sources to a SANscreen environment and at what point you should consider expanding your SANscreen environment as it grows. Finally, the addition and/or removal of SAN devices from your environment may require Data Source modification as well.

2.1 Data Source Monitoring

The consistent acquisition of data from SAN devices is imperative in order to properly maintain a healthy SANscreen environment. SANscreen integrates with several third party SAN vendor products. As a result, it is possible that changes to your SAN devices may cause Data Source acquisition to fail. For example, firmware revisions and configuration changes (for example, using new administrative functions) may cause Data Sources that were working as expected to suddenly fail to acquire data from their target devices.

Data Source status can be monitored by using SNMP alerts or SMTP notifications. These options are discussed below. Data Source status can also be monitored in the Data Source view.

SNMP Alerts

SANscreen supports the use of SNMP alerting for Data Source failures. SANscreen comes with a MIB file that should be given to the IT personnel responsible for maintaining the company's SNMP management infrastructure.

Table 1 below illustrates the MIB structure for Data Source monitoring.

Table 1 – SNMP MIB Structure for Data Sources

SNMP Counter Description	OID
Acquisition Unit Name	.1.3.6.1.4.1.17187.1.3.1.2
Acquisition Unit Status (OK/Not Connected)	.1.3.6.1.4.1.17187.1.3.1.3
Data Source Active (True/False)	.1.3.6.1.4.1.17187.1.3.1.4
Data Source Name	.1.3.6.1.4.1.17187.1.3.1.5
Time of Last Successful Acquisition	.1.3.6.1.4.1.17187.1.3.1.6
Vendor of SAN device being monitored	.1.3.6.1.4.1.17187.1.3.1.7
Model of SAN device being monitored	.1.3.6.1.4.1.17187.1.3.1.8

If one of the Data Source acquisition attempts fails, an SNMP trap will be sent to the console as configured within the SANscreen portal. The flexibility of SNMP allows for conditional routing and handling to be configured based on any of the values specified above. For example, a Data Source failure on a Brocade switch Data Source can be sent to the Fabric team, while a Data Source failure on an EMC CLARiiON Data Source can be sent to the EMC storage team. Additionally, an SNMP console administrator using third party trap software could take the SANscreen SNMP traps and route them based on identification values in the SNMP trap itself.

SMTP Alerts

If SNMP is not configured in your environment, SMTP may be used to alert when Data Source failures occur. In order to configure SMTP alerting for Data Sources failures, you need to configure an SMTP gateway within the SANscreen portal.

Once SMTP has been configured, subscribe to the Data Source report within the SANscreen portal's Reports section. The report you wish to subscribe to is called "Data-Source Report". This report is illustrated below:

Figure 1 - Data Source Report

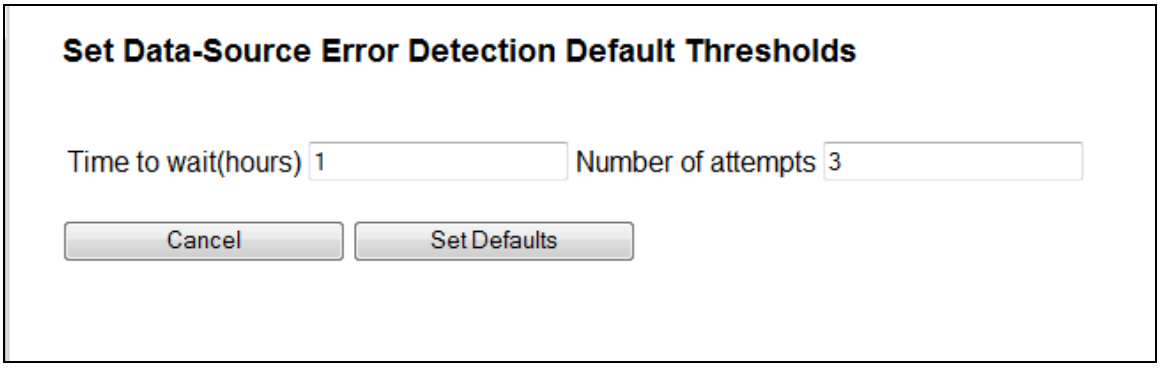

You configure SANscreen to generate the report when a Data Source fails a certain number of times in a given time interval. The default is to trigger the report when a Data Source fails three times within one hour. Note that this is a global setting; it applies to all Data Sources. There is no way to configure different settings for individual Data Sources in SANscreen 5.x.

SANscreen 6.0 includes an option where you can set the number of minutes after the Data Source begins before an alert is generated indicating the Data Source has not yet finished.

2.2 Creating New Data Sources

The creation of new Data Sources might be required when adding new hardware to the SAN fabric. The following devices might require new Data Sources to be configured:

- Fibre Channel Switches
- Storage Arrays
- VMware VirtualCenter/vSphere instances
- SRM instances

Table 2 below describes when it is necessary to add a new Data Source to the SANscreen environment.

Table 2 – Data Source Creation Matrix

Added Object	Data Source Type	New Data Source Required?	Notes
Storage Array	One-to-one	Yes	For example, CLARiiON arrays
Storage Array	Management Console	No	Assumes managed by currently configured array management console
Storage Array	Management Console	Yes	If a new management console is deployed to manage array
Switch	Non-fabric aware	Yes	For example, McData SNMP
Switch	Fabric aware	No	For example, Cisco SNMP
Switch	Management Console	No	For example, EFCM and switch is added to existing EFCM
Switch	Management	Yes	New EFCM is deployed to manage switch

Added Object	Data Source Type	New Data Source Required?	Notes
	Console		
VMware	VMware	Yes	If adding new VirtualCenter/vSphere instances, *not for ESX servers added to existing instances*
VM	VM	No	
SRM	ECC	Yes	If adding new ECC instances

If you configured an include filter in SANscreen to include only a few devices, then SANscreen will not display the addition.

2.3 Retiring Switches and Arrays

SANscreen does not delete SAN devices from Service Insight automatically. This ensures that SAN devices that are offline for maintenance do not get removed from the SANscreen database. Therefore, it is necessary to have a formal change plan for SANscreen in order to eliminate orphaned entries in the SANscreen database from remaining behind. The process for removing retired SAN devices varies based on the way the Data Source collects information about the device that is being retired.

Single Switches Managed by a Single Data Source

If a single switch is managed by a single Data Source, complete the following steps in order for the device to be removed from the SANscreen database:

1. Stop the Data Source for the target device.
2. Remove the Data Source from the target device.
3. Find the inactive device in Service Insight→Switches (the device will be in plain black text, not the normal blue text).
4. Remove the inactive device by selecting it, right clicking and selecting **Remove Inactive Device**.

For example, this process would apply to the removal of a single McData switch that is collected via the McData SNMP Data Source.

Single Switches Managed by a Fabric Aware Data Source

If a single switch is managed by a fabric-aware Data Source, complete the following steps in order for the device to be removed from the SANscreen database. (SANscreen 6.x includes an enhancement to this process.)

For example, this process would apply to the removal of a single McData switch that is collected via the EFCM/Connectrix Data Source.

1. Copy down the configuration of the current Data Source acquiring the device.
2. Create a new Data Source that has the exact configuration of the original Data Source (make sure the name is unique).
3. Force an acquisition on the newly created Data Source.
4. Stop the original Data Source.
5. Remove the original Data Source.

6. Find the inactive device in Service Insight→Switches (the device will be in plain black text, not the normal blue text).
7. Remove the inactive device by selecting it, right clicking and selecting **Remove Inactive Device**.
8. Stop the newly created Data Source and edit its properties so that it has the same name as the original Data Source (*optional*).
9. If Step #8 has been performed, restart the newly created Data Source.

Single Arrays Managed by a Single Data Source

If a single array is managed by a single Data Source, complete the following steps in order for the device to be removed from the SANscreen database.

1. For example, this process would apply to the removal of a single EMC CLARiiON array that is collected via the EMC CLARiiON CLI Data Source. Stop the Data Source for the target device.
2. Remove the Data Source for the target device.
3. Find the inactive device in Service Insight→Storage Arrays (the device will be in plain black text, not the normal blue text).
4. Remove the inactive device by selecting it, right clicking and selecting **Remove Inactive Device**.

Single Arrays Managed by a Management Console Data Source

If a single array is managed by a Management Console Data Source, complete the following steps in order for the device to be removed from the SANscreen database.

For example, this process would apply to the removal of a single Hitachi array that is collected via the HiCommand Data Source.

1. Copy down the configuration of the current Data Source acquiring the device.
2. Create a new Data Source that has the exact configuration of the original Data Source (make sure the name is unique).
3. Force an acquisition on the newly created Data Source.
4. Stop the original Data Source.
5. Remove the original Data Source.
6. Find the inactive device in Service Insight→Storage Arrays (the device will be in plain black text, not the normal blue text).
7. Remove the inactive device by selecting it, right clicking, and selecting **Remove Inactive Device**.
8. Stop the newly created Data Source and edit its properties so that it has the same name as the original Data Source (*optional*).
9. If Step #8 has been performed, restart the newly created Data Source.

2.4 Retiring Hosts and Tapes

SANscreen does not delete Host or Tape devices from Service Insight automatically. This is to ensure that Host and Tape devices that are offline for maintenance do not get removed from the SANscreen database. Therefore, it is necessary to have a formal change plan for SANscreen in order to eliminate orphaned entries in the SANscreen database from remaining behind. The process for removing retired Host and Tape devices is described below:

1. Force acquisition on the Data Source for the switch/fabric where the original host or tape device was connected.

2. Find the host or tape device that was removed in Service Insight (either under Service Insight→Hosts or Service Insight→Tapes). The device will be in black text.
3. Remove the inactive device by selecting it, right clicking, and selecting **Remove Inactive Device**.
4. If necessary, to prevent a path outage violation, remove any Host or Path specific Service Assurance policies that were applied to this host device.
5. As a best practice, remove any relevant zoning and masking entries for the removed host or tape device.

2.5 Eliminating Duplicate Data Source Reporting

It is possible for a single SAN device to be reported on by one or more Data Sources. Duplicate device reporting has an adverse affect on SANscreen accuracy as it relates to the following:

- Switch port counts
- Array capacities
- Violation and vulnerability reporting

Therefore, it is critical to avoid duplicate device reporting. There are currently no automated ways for a customer to detect duplicate device reporting from multiple data sources.

The best way to search for duplicate device reporting is to look for entries under the "Additional Data Sources" column within the Data Sources→Devices micro-view. Select each Data Source and enable the Devices micro-view. If you see entries for SAN devices (such as individual switches or arrays) under the "Additional Data Sources" column, you have duplicate reporting for this SAN device. *Note that it is acceptable to have entries in this column for SAN fabrics and VSANS*.

For example, you might have two HiCommand instances that manage the same USP array. In this case, the USP array will be reported on by both Data Sources. This results in duplicate reporting for this USP array. You need to determine from which HiCommand instance you wish to acquire the USP array.

If duplicate reporting for SAN devices is reported, you must identify the appropriate Data Source for acquisition of the affected SAN device. If you are unsure as to the appropriate Data Source, you should contact NetApp's SANscreen support. Once you determine the appropriate Data Source, you might need to configure an exclusion on the incorrect Data Source so that it does not reacquire the device.

Continuing the USP example from above, you would need to configure an exclusion for the specific USP array on the Data Source from which you *do not* wish to acquire. Another option would be to remove the USP array from one of the HiCommand instances, but this would require a change to the customer environment and thus might not be optimal.

2.6 Managing the Number of Data Sources per Server

- **When to split Data Sources between Acquisition Units on the same server**: Consider splitting Data Sources when the acquisition process consumes too much memory or reaches close to the 32-bit maximum amount of memory, which is about 1.2 GB, assuming that the machine has enough memory to perform one more 2 GB process.
- **When to split data sources on the same Acquisition Unit**: This results in one Data Source sampling devices A and B, while another Data Source samples devices C and D, rather than one Data Source sampling all A, B, C, and D. When the Data Sources work in parallel, you gain a faster response.

  Splitting Data Sources might result in lower RAM requirements, depending upon the specific Data Source. For example, if a Data Source brings in all the data for A, B, C, and D devices, the one Data Source puts in memory all four devices, yet the two Data Sources each holds only two

devices in memory. Then, if they are not sampling at the same time, the amount of RAM will be lower. However, if a Data Source provides information for only one device at a time, there might not be any advantage to using only one Data Source.

- **When to split data sources to a Remote Acquisition Unit (RAU) for scaling**: Consider splitting data sources when the machine does not have enough memory to run another process.

3 HOST IDENTIFICATION

Maintaining accurate host identification within SANscreen is critically important. Proper host resolution is required in order to obtain optimal SANscreen configuration for:

- Violation reporting
- Vulnerability reporting
- Capacity reporting in the Data Warehouse
- Performance metrics from Application Insight

This section provides best practices and guidance to assist a customer in maintaining proper host resolution within SANscreen.

3.1 How Does SANscreen Discover Hosts?

In the absence of SRM tools like ECC, SANscreen discovers hosts as they login to Fibre Channel switches. To SANscreen, a host device is represented by an HBA or group of HBAs performing a login to an FCP switch name server. SANscreen then represents the host device with the WWPN and WWNN of the HBA that has performed the name server login.

SANscreen discovers new hosts added to the fabric when data source polling occurs on the target FCP switch. For example, a new host with two HBA's is connected to "Fabric A" and "Fabric B". The new unidentified host will not be imported into SANscreen until the Data Sources responsible for both "Fabric A" and "Fabric B" are polled by SANscreen.

3.2 Host Resolution Techniques

Because SANscreen is an agent-less application, SANscreen does not have a way to pull host name information directly from the host. SANscreen has to tie soft attributes (human configurable) back to hard attributes (hard-coded WWN information) in order to identify host devices. Examples of soft attributes include storage alias entries, switch alias entries, and zoning entries.

Hosts that have not been identified are represented in SANscreen as Generic Devices. Unidentified devices can be seen by browsing to Service Insight → Generic Devices. Further analysis of unidentified devices can be performed by browsing to the Admin → FC Identify screen and filtering by the Type of "Unknown."

Host resolution can be accomplished using one of the following techniques:

- Auto resolution
- Identification via CSV files
- Manual identification

Auto Resolution

The preferred method of maintaining accurate host identification is to enable auto host resolution within SANscreen. Auto resolution uses soft attributes as described above in order to extrapolate host names. Figure 2 below illustrates the different methods that can be used to perform auto host resolution within SANscreen.

Figure 2 - Auto Resolution Techniques

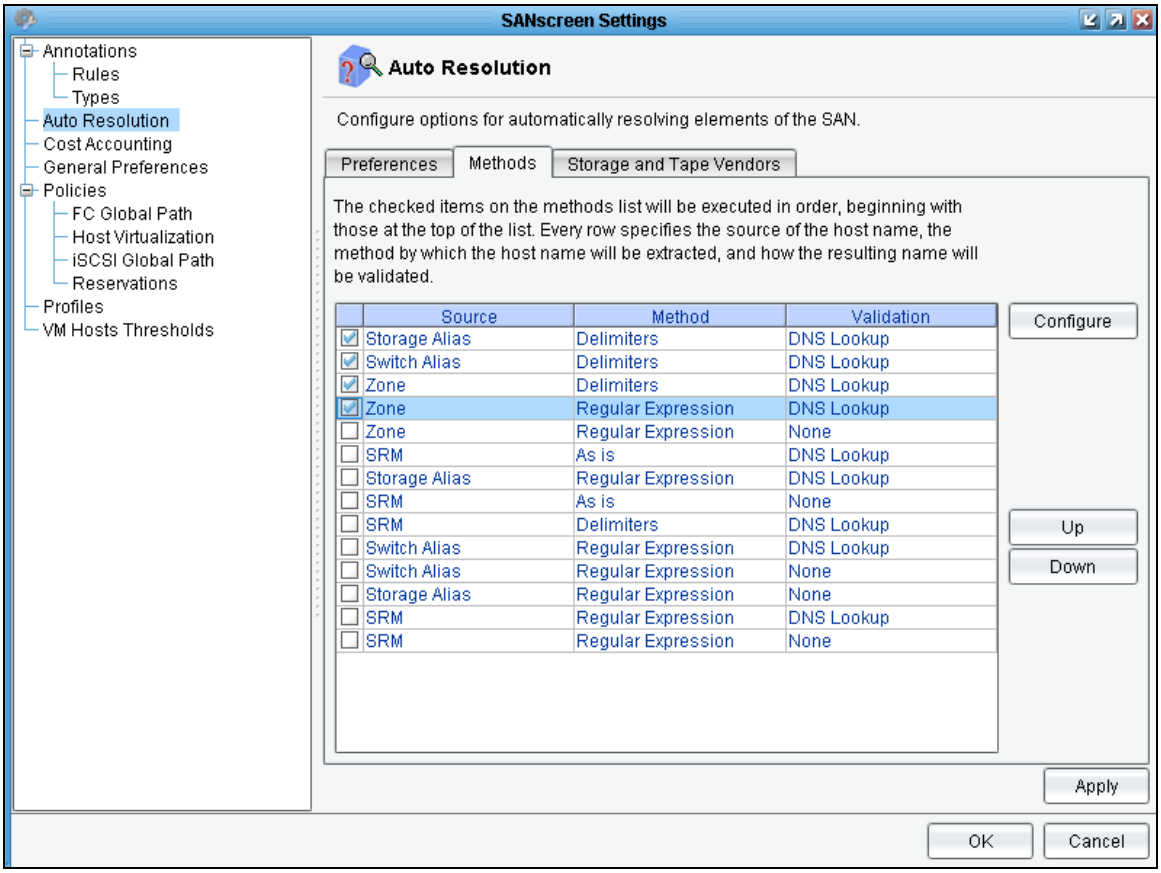

The auto host resolution methods illustrated above are the ones most commonly configured by NetApp Professional Services in the field. The use of DNS Lookup validation is critical to ensure that identified hosts match what is currently registered on the customer's network.

If DNS is not reliable or accessible in your environment, you might want to use Zones and Regular Expressions without DNS validation. However, keep in mind that selecting this option increases the probability that identified host names do not match the actual host names on the network.

Identification Using CSV Files

In some customer environments, soft attribute mappings are not accurate enough to obtain a viable initial load of host to WWN identification. In these instances, it is valuable to obtain a mapping of HBA WWN's from an authoritative source such as an enterprise CMDB.

SANscreen provides examples of CSV files that may be used for performing host identification. By default, these files are installed in the *c:\<sanscreen_installation_directlry>\imports* directory. An example of a CSV file used for host identification is illustrated in Figure 3 below.

Figure 3 - CSV File for Host Identification

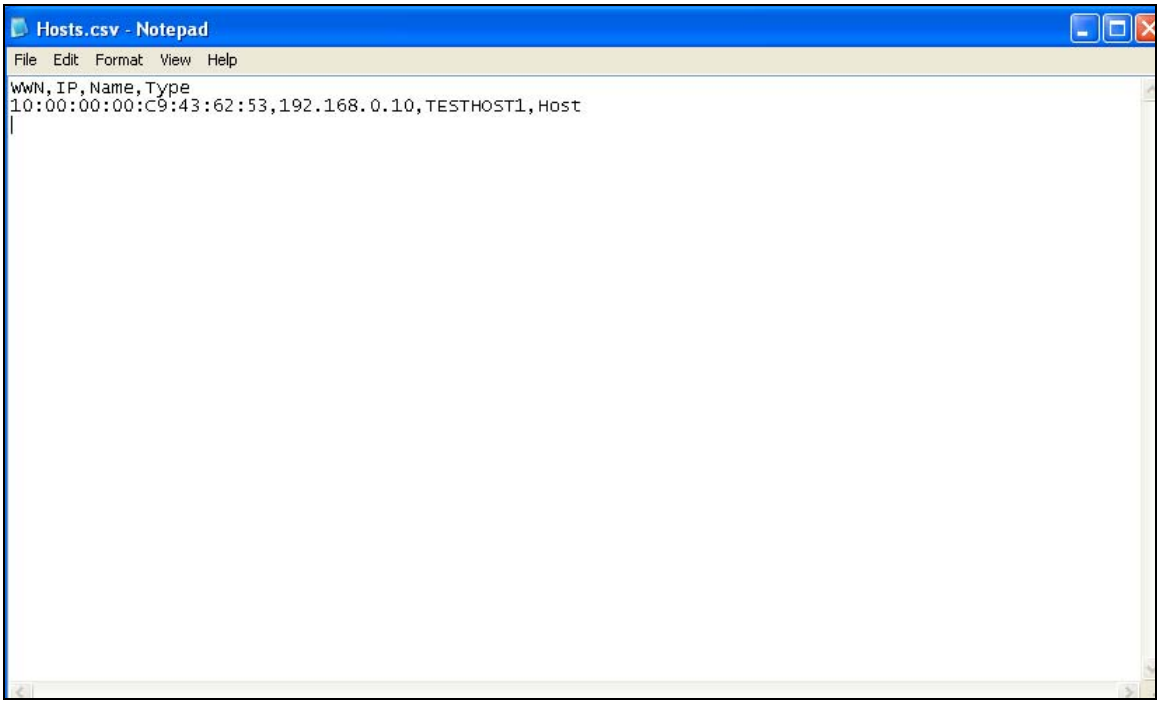

The fields at the top of the file are required and are defined as follows:

- WWN – World Wide Port Name of the HBA device
- IP – IP address of the host, can be in dotted decimal format or text format if IP address is unknown/not assigned
- Name – simple name of the host, not the FQDN of the host
- Type – set to Host in this example because we are identifying Host Devices

Once the CSV file is properly formatted, it can be imported into SANscreen:

1. From the Admin → FC Identify screen, right-click anywhere on the screen.
2. Select **Identify from File**. Browse for the CSV file you create and open that file.
3. There is no visual confirmation that the import occurred but at the bottom of the screen, you will see the “Update Changes” button enabled.
4. Click **Update Changes** to apply the changes from the imported CSV file.

Note that the WWPN you enter into the CSV file must exist in SANscreen. You cannot load WWPN’s into SANscreen via a CSV file. SANscreen must have already discovered the WWPN during normal Data Source acquisition.

The order of the identification is important and is retained based on the following order:

- VM identification is always the highest precedence and cannot be overridden.
- Manual identification
- Auto identification according to the order displayed in the Auto Resolution screen.

CSV file creation is also described in NOW article KB39690.

Manual Identification

A final option for host identification is to select individual entries in the FC Identify screen and identify them manually. The manual identification option is usually performed when the specific entry does not identify via auto host resolution and the list of entries is small enough where it is not worth the effort to format a CSV file.

Manual identification is performed as follows:

1. Select the specific entry (or entries) in the Admin → FC Identify screen.
2. Select **Identify Selected**.
3. Enter the proper IP address, name and device type in the dialog box (as illustrated below in Figure 4).
4. Click **OK**.
5. Click **Update Changes** in the FC Identify screen.

Figure 4 - Manual Device Identification

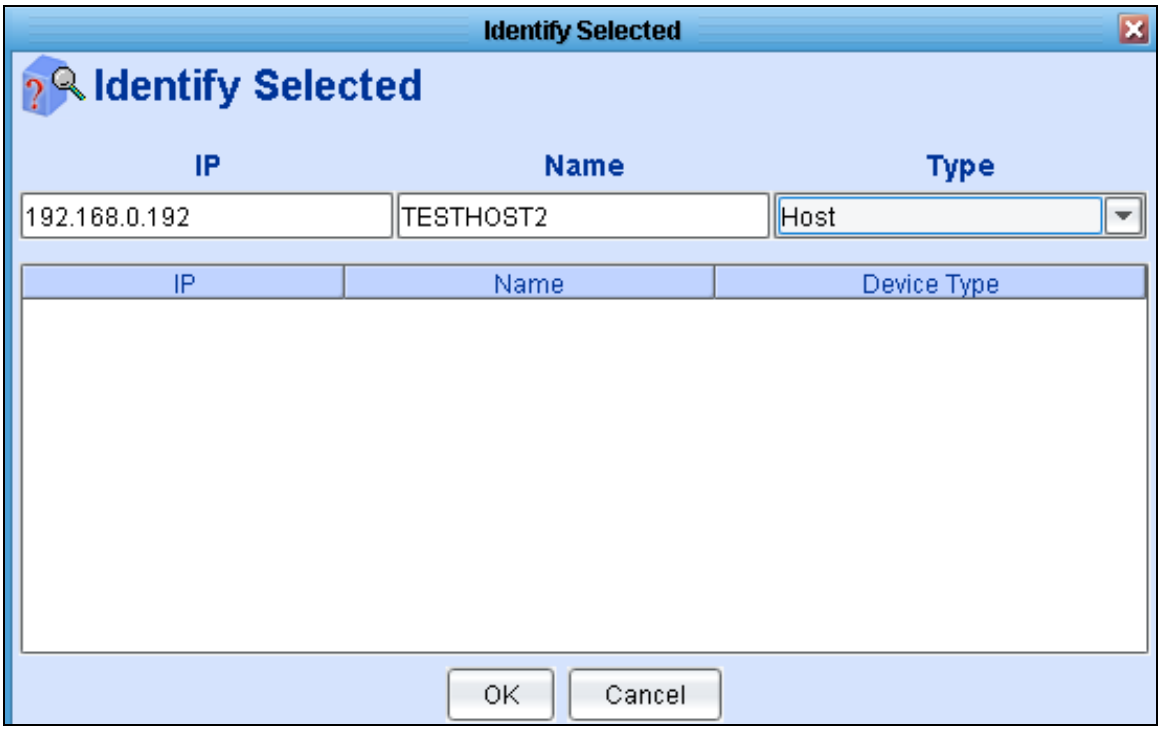

3.3 Handling Host Renaming

The renaming of host machines is a common occurrence within any enterprise. Optimally, SANscreen detects when a host is renamed and reflects these changes within Service Insight. However, there are some caveats that must be taken into account in order for SANscreen to properly handle host renaming via auto resolution.

When a host is renamed on the network, SANscreen only detects the name change if the soft attribute used to initially automatically identify the device is updated to reflect the change. Some examples will help to illustrate.

Table 3 – Host Name Change on Network, Auto-Resolution Fails to Detect Change

Event on Network/SAN	SANscreen Activity	Notes
Name change on host	None	Host was initially identified using Storage Alias with Delimiters
Storage Alias is not changed on array	Data Source acquisition occurs but no change in storage alias	Because Storage Alias is not updated on array, auto resolution does not reflect name change
Zoning entry is changed on switch	Data Source acquisition occurs and new zoning entry is acquired	Because auto resolution used Storage Alias for this host, new zone entry is not used for new host name

In this example, the name change on the network will not be reflected within SANscreen. This is because we used Storage Alias initially to auto identify the host and we did not update the Storage Alias to reflect the new name change. At this point, we have two options:

1. Update the Storage Alias to reflect the new name (preferred)
2. Remove the identification from the device in FC Identify and allow auto resolution to detect the new name via the zone entry. This assumes that zone identification was performed before the storage alias change.

Table 4 – Host Name Change on Network, Auto-Resolution Succeeds

Event on Network/SAN	SANscreen Activity	Notes
Name change on host	None	Host was initially identified using Storage Alias with Delimiters
Storage Alias is changed on the array	Data Source acquisition occurs and change in storage alias is captured	Host name will be reflected in SANscreen because Storage Alias technique detected change
Zoning entry is changed on switch	Data Source acquisition occurs and new zoning entry is acquired	No affect on SANscreen with regards to reflecting the name change via auto host resolution

In this example, the name change on the network will be reflected within SANscreen. This is because we used Storage Alias to auto identify the host and we did update the Storage Alias to reflect the new name change. No further actions are required.

4 TAPE IDENTIFICATION

Unlike host devices, tape devices cannot be acquired via the auto resolution process. In addition, there are no Data Sources for tape devices. This includes VTL devices that are shipped as part storage arrays but have VTL personalities. For example, NetApp VTL devices may successfully be acquired using the NetApp ONTAP Data Source, but this will lead to an incorrect device identification of Storage when in reality SANscreen should see the VTL device as a tape.

Maintaining accurate tape identification within SANscreen is critically important. Proper tape identification is required in order to obtain optimal SANscreen configuration for:

- Violation reporting
- Vulnerability reporting
- Performance metrics from Application Insight

This section provides best practices and guidance to assist a customer in maintaining proper tape identification within SANscreen.

4.1 How Does SANscreen Discover Tapes?

SANscreen will discover tapes as they login to Fibre Channel switches. To SANscreen, a tape device is represented by a tape port performing a login to an FCP switch name server. SANscreen then represents the tape device with the WWPN and WWNN of the tape port that has performed the name server login.

SANscreen discovers new tapes added to the fabric when data source polling occurs on the target FCP switch. For example, a new host with two tape ports is connected to “Fabric A” and “Fabric B”. The new tape will not be imported into SANscreen until the Data Sources responsible for both “Fabric A” and “Fabric B” are polled by SANscreen.

4.2 Tape Identification Techniques

Because there are no Data Sources for tapes, there is no way for SANscreen to acquire the proper identification of tape devices via acquisition. In addition, auto resolution does not apply to tape devices so there is no automated way to tie soft attributes to WWPN of the tape ports. Therefore, a manual approach must be used for tape devices. Tape device identification can occur using one of the following methods:

- Identification via CSV files
- Manual identification

Identification Using CSV Files

Similar to host entries discussed earlier, it is possible to obtain a mapping of tape device WWN’s from an authoritative source such as an enterprise CMDB.

SANscreen provides examples of CSV files that may be used for performing tape identification. By default, these files are installed in the *c:*\<*sanscreen_installation_directlry*>*imports* directory. An example of a CSV file used for host identification is illustrated in Figure 5 below.

Figure 5 - CSV File for Tape Identification

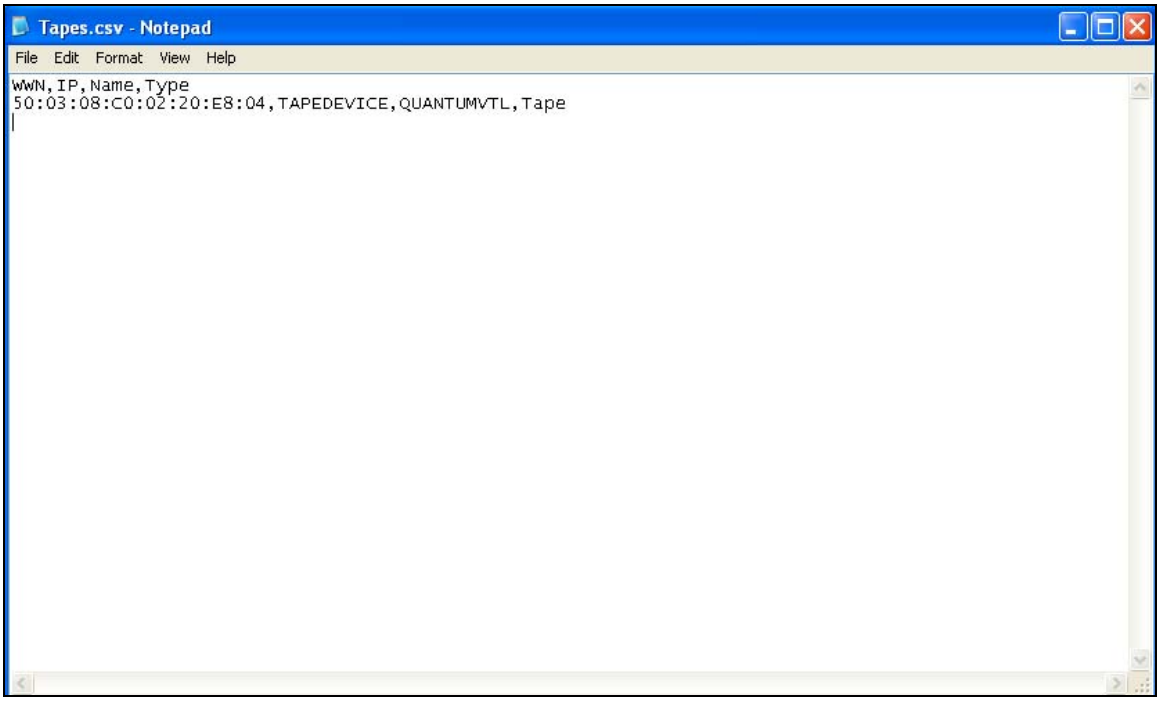

The fields at the top of the file are required and are defined as follows:

- WWN – World Wide Port Name of the tape port
- IP – IP address of the tape, in this example we use a text entry because the tape device does not have an IP
- Name – simple name of the tape device
- Type – set to Tape in this example because we are identifying Tape Devices

Once the CSV file is properly formatted, it can be imported into SANscreen.

1. From the Admin → FC Identify screen, right click anywhere on the screen.
2. Select **Identify from File**.
3. Browse for the CSV file you create and open that file.

   There is no visual confirmation that the import occurred but at the bottom of the screen, you will see the "Update Changes" button enabled.
4. Click **Update Changes** to apply the changes from the imported CSV file.

Note that the WWPN you enter into the CSV file must exist in SANscreen. You cannot load WWPN's into SANscreen via a CSV file. SANscreen must have already discovered the WWPN during normal Data Source acquisition.

CSV file creation is also described in NOW article KB39690.

Manual Identification

Another option for tape identification is to select individual entries in the FC Identify screen and identify them manually. The manual identification option is usually performed when the list of entries is small enough where it is not worth the effort to format a CSV file.

Manual identification is performed as follows:

1. Select the specific entry (or entries) in the Admin → FC Identify screen.
2. Select **Identify Selected**.
3. Enter the proper IP address, name and device type in the dialog box (as illustrated below in Figure 6).
4. Click **OK**.
5. Click **Update Changes** in the FC Identify screen.

Figure 6 - Manual Tape Device Identification

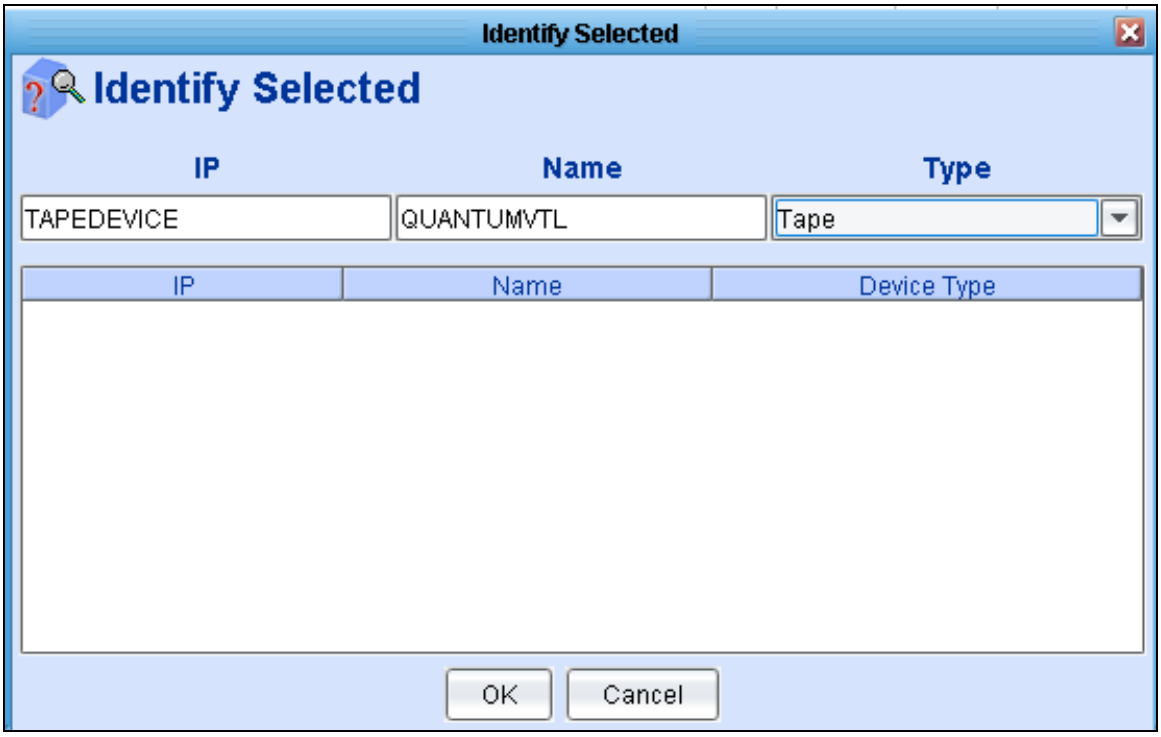

4.3 Storage Vendors Tab and Tape Identification

SANscreen comes configured with the OUI numbers for several common storage vendors. A storage vendor should be identified on the Storage Vendors tab in the Auto Resolution settings in order to make sure a storage device is not improperly configured as a host device during auto resolution. When auto resolution runs, any device whose OUI matches one of the applied storage vendors will be ignored by auto resolution so that it does not get erroneously identified as a host device.

An example should help to clarify this. Looking at Figure 7, you will see that Data Domain is not configured as a storage vendor within the customer's Storage Vendor tab.

Figure 7 - Storage Vendors Tab

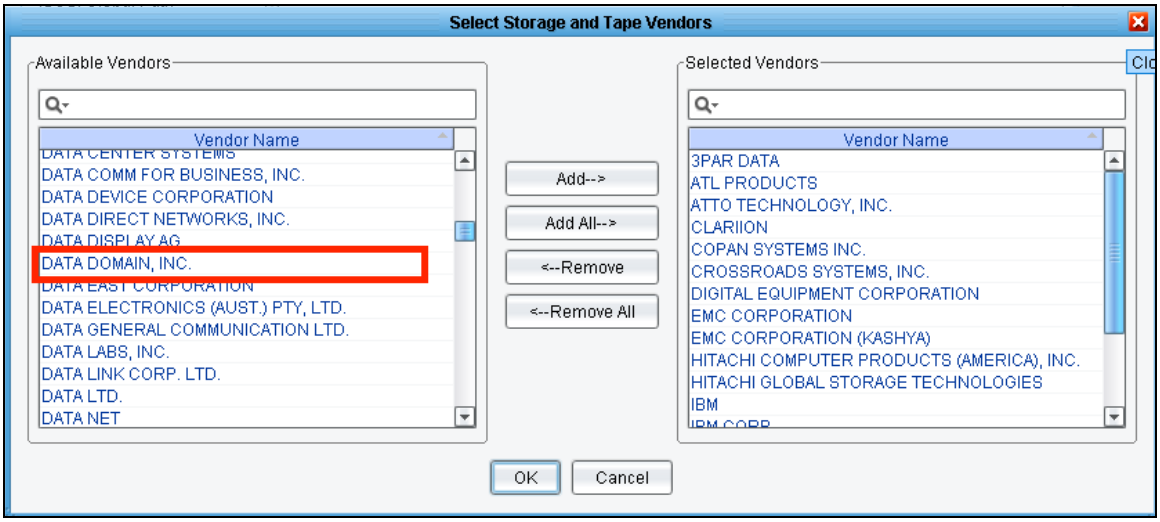

Data Domain appears under the Available Vendors side of the dialog, but not under the Selected Vendors side of the dialog. As a result, SANscreen could mistakenly identify any SAN devices that match the Data Domain OUI as host devices.

By selecting Data Domain and adding it to the Selected Vendors side, SANscreen will know to skip auto resolution for any devices matching the Data Domain OUI, thus allowing the device to be properly identified as a tape device.

5 ANNOTATION MAINTENANCE

Annotations are metadata that describe SAN devices within SANscreen. Annotations are used for the following purposes:

- Application sharing violations
- Data Warehouse reporting

SANscreen provides default annotations. Custom annotations can be created if required for customer reporting needs.

5.1 Commonly Used Annotations

SANscreen comes with several pre-created annotations, but there are certain annotations that are commonly configured. These annotations are described in Table 5 below.

Table 5 - Commonly Used Annotations

Annotation Name	Description	Use Cases
Application	Application name used for VM	Only used for VM's within VMware
Business Unit	Business Units within the organization	Commonly used for reporting capacity and chargeback
Application Group	Combination of Business Unit, Application, and Application Priority	Used for reporting capacity by application, also used for Application Sharing violations
Data Center	Physical location of SAN devices	Used for reporting capacity by Data Center
Tier	Classification of storage, as defined by the customer	Used for reporting capacity by tier, chargeback by tier, storage tier analysis
Cost Accounting	Cost per GB of storage or per port for switches	Used for chargeback reporting
Storage Tier Rules	Ties specific storage vendors or arrays to tiers	Used for chargeback reporting, capacity by tier and storage tier analysis
Volume Tier Rules	Ties specific volume definitions using criteria such as RAID level, disk type, disk speed to tier levels	Used for chargeback reporting, capacity by tier and storage tier analysis

At a minimum, the above mentioned annotations should be defined for proper metadata population to enable application sharing violation tracking and proper reporting out of the Data Warehouse.

To define annotations, complete the following steps.

1. From the SANscreen Client Tools menu, select **SANscreen Settings** > **Annotation** > **Types**.
2. Select **Add** and add the annotation.
3. Select the resources, for example, switches, to which you want to be able to assign the annotation.

To set annotation on a device, complete the following steps.

1. In the SANscreen client, display a device view. For example, select the Hosts view.
2. Right-click on a device and select **Set Annotation**. You can assign the annotation to multiple devices at once, if you select multiple devices.

3. Select the annotation, for example, data center, that you want to assign to the device.
4. Some annotations require additional values, such as the building or floor. Enter that information.
5. Click **OK**. The annotation is now attached to the device. You can later filter or group devices by the annotation.

5.2 Application Sharing Violations

SANscreen Service Assurance can enforce proper volume sharing for hosts. For example, hosts in high availability clusters often need to be masked to the same volumes to allow for failover. However, hosts in unrelated applications usually have no need to access the same physical volumes. In addition, regulatory policies may require customers to explicitly disallow unrelated applications from accessing the same physical volumes.

When Application Sharing is enabled, hosts in the same Application Group will be allowed to access the same set of physical volumes.

Let's look at an example, as illustrated in Table 6 below. For this example, assume both hosts are masked to a volume "0036".

Table 6 - Volume Sharing

Host	App Group	Host	App Group	Violation?
Apple	Finance	Grape	Finance	No
Apple	Finance	Grape	R&D	Yes
Apple	Not Defined	Grape	Not Defined	Yes

If you wish to enforce volume sharing via Service Assurance, you must properly define your Application Groups so that clustered hosts are placed into the same Application Groups. You need to also define the policy and enable the option indicating that the application is not ignored for sharing.

5.3 Data Warehouse Reporting

A full discussion of Data Warehousing concepts is beyond the scope of this document. A Data Warehouse consists of Fact Tables (things you measure) and Dimension Tables (descriptive information on objects that are related to facts).

Data Warehouses are used for what is commonly called "slice and dice" operations. By applying more granular dimensions to measured facts, a Data Warehouse allows for an end user to see a finer grain of what is being measured. The corollary is when you remove dimensions from measured facts; you will see a less granular view of the fact that is being measured.

In SANscreen, some dimension tables include the annotations that are populated within SANscreen. The annotations become the metadata used for populating the Data Warehouse dimension tables. When annotation data is not properly populated in SANscreen, you might see one of the following issues arise when running Data Warehouse reports:

- Pre-packaged dashboards might not run
- Report data might consist of several "N/A" entries (Not Available)

If you experience any of these issues, you should go back and populate more annotations within SANscreen and then perform an ETL (Extract, Transform and Load) from SANscreen to the Data Warehouse. For details, see the *SANscreen Data Warehouse User Guide*.

6 SERVICE ASSURANCE MAINTENANCE

Service Assurance reports violations as a result of hosts not being in compliance with policies. Opportunities exist for violations to be raised either when Data Source acquisition occurs or when configuration changes occur within SANscreen. Therefore, it is critical that SANscreen be monitored so that violations introduced into the environment can either be:

- Fixed in the SAN environment itself
- Configured as exceptions using discrete policies

6.1 Monitoring for Violations

In addition to monitoring violations on the Violations view, there are two other primary methods:

- SNMP alerts
- SMTP alerts

SNMP Alerts

SANscreen supports the use of SNMP alerting for Service Assurance violations. SANscreen comes with a MIB file that should be given to the IT personnel responsible for maintaining the company's SNMP management infrastructure.

SANscreen reports on violations only when they occur, for example when a host cannot access its storage. SANscreen does not report a violation when it is resolved, for example, when the host finally accesses the storage.

Table 7 below illustrates the MIB structure for Service Assurance violations.

Table 7 - SNMP Monitoring for Violations

SNMP Counter Description	OID
Time the violation occurred	.1.3.6.1.4.1.17187.1.2.1.2
Initiator of the violation, either a host or generic device	.1.3.6.1.4.1.17187.1.2.1.3
Target of the violation, either an array or tape device	.1.3.6.1.4.1.17187.1.2.1.4
Target volume of the violation	.1.3.6.1.4.1.17187.1.2.1.5
Violation type (defined below)	.1.3.6.1.4.1.17187.1.2.1.6
IP address of the initiating device	.1.3.6.1.4.1.17187.1.2.1.7

Violation Types are defined in Table 8 below.

Table 8 - Violation Types

Violation Type	Description
1001	Volume exists but is not approved
1002	Volume approved but does not exist
1003	Path has a Single Point of Failure
1004	Path does not have redundancy
1005	Number of switch hops exceeded
1006	Number of host ports is less than required

Violation Type	Description
1007	Shared volume is being accessed
1008	Number of storage ports is less than required

SMTP Alerts

If SNMP is not configured in your environment, SMTP can be used to alert when Service Assurance violations occur. In order to configure SMTP alerting for violations, you will need to configure an SMTP gateway within the SANscreen portal. Once SMTP has been configured, the next step is to subscribe to the Violations report within the SANscreen portal's Reports section. The report you wish to subscribe to is called "Violation Report." This report is illustrated below:

Figure 8 - Violation Report Subscription

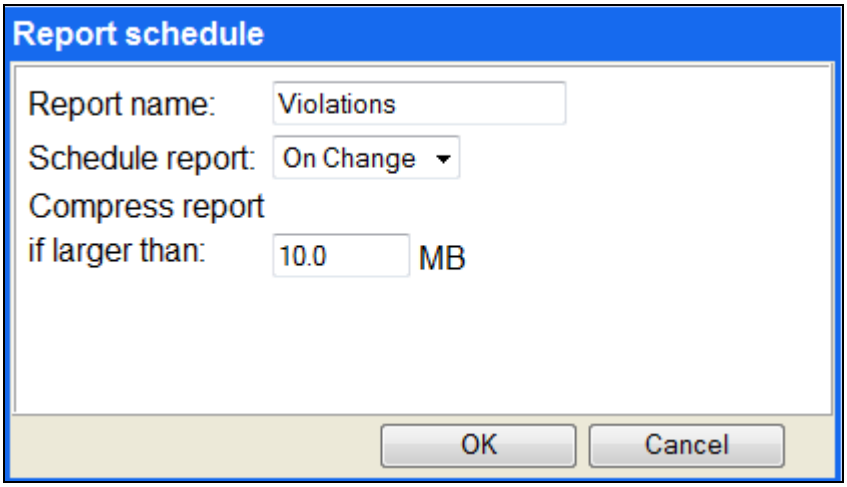

The report is scheduled to be generated "on change." This means that each time a violation occurs the report will be updated and subsequently emailed to the personnel who subscribe to it.

It is important to note that one email will be triggered for each violation that occurs. For example, if a host is masked to ten volumes and one of the host HBAs has an issue, ten different violations will be raised. In this case, ten separate emails will be generated. As a result, SMTP alerting for violations might generate quite a bit of mail for users who subscribe to this report.

Another important factor with SMTP alerting for violations is that there is no way to filter out violation reporting for scheduled maintenance. If a host is scheduled for volume migrations, several path outage violations will be raised for the source volumes, resulting in a significant amount of mail generated.

Given the factors mentioned above, SNMP is the preferable method for Service Assurance violation monitoring.

6.2 Using Global Policies

SANscreen allows for global policy settings as defined below:

Global Policy Type	Description
Storage	Default policy for volume access
Volume Type Exceptions	Policies enforced by volume type as discovered by SANscreen (i.e. RAID5)
Volume Capacity Exceptions	Policies enforced when volume capacity is below a specified threshold (i.e. Symmetrix Gatekeeper volumes)

Global Policy Type	Description
Tape	Default policy for initiator access to tape devices

The global Storage policy should be set to match the customer's standard configuration for host access to storage. Figure 9 below illustrates the most common setting for the global Storage policy.

Figure 9 - Typical Storage Policy Settings

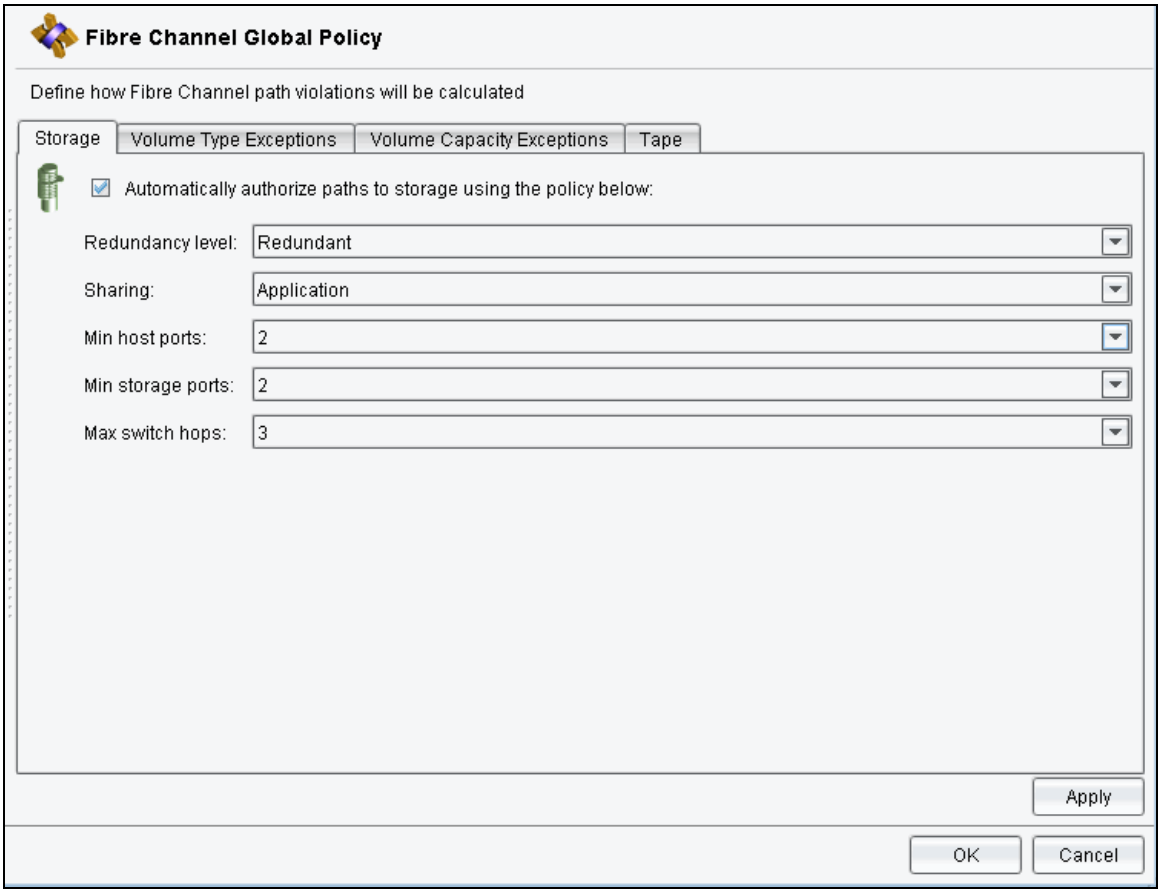

6.3 Using Discrete Policies

There are times when global policy settings will not accurately reflect the configuration of a host that is being reported on by SANscreen. For example, a common scenario is a test/staging server that is located on a production fabric. The server might only have a single HBA and there is no desire to add another HBA for redundancy.

In this example, a standard global Storage policy that enforces redundant connectivity would report a violation on this host. To clear this violation, a discrete policy should be configured so that the exception for this host does not result in a violation. There are two types of discrete policies that can be set within SANscreen:

- Path Policies
- Host Policies

Path Policies

The most granular policy that can be set in SANscreen is a Path policy. A Path policy is set from a specific initiator to a specific volume. For example, if redundancy is not required for host Apple to volume 0036, then a Path policy could be configured for this path.

Keep in mind that this path policy is only for one specific volume; if multiple volumes accessed by the same initiator have the same redundancy requirement, you will need to create multiple Path policies. However, in this case a better option would be to configure a Host policy.

Host Policies

Host policies are configured when an entire host has specific policy requirements for all volumes that are accessed. For example, let's assume host Apple has access to twenty different volumes on an array. Apple does not require redundancy to any volumes that it may access. In this example, you could conceivably create twenty different Path policies. However, a better solution would be to create a Host policy and configure that Host policy to not require redundancy.

By design, SANscreen enables you to set policies on a more global scale and then override them on specific objects. For example, you can set a global policy, but override it for a specific host. Additionally, you can create a host policy and then override it on a specific path.

7 BACKUP PROCEDURES

7.1 SANscreen Server Backup

SANscreen can be backed up either manually from the web portal or in an automated way using the Windows Scheduler and a batch file. Both methods are discussed here.

Manual Backup of SANscreen Server

SANscreen can be manually backed up by browsing to the web portal on the SANscreen server and selecting Backup/Restore from the menu on the left. You will be presented with the following backup options as illustrated below in Figure 10:

Figure 10 - Manual Backup of SANscreen Server

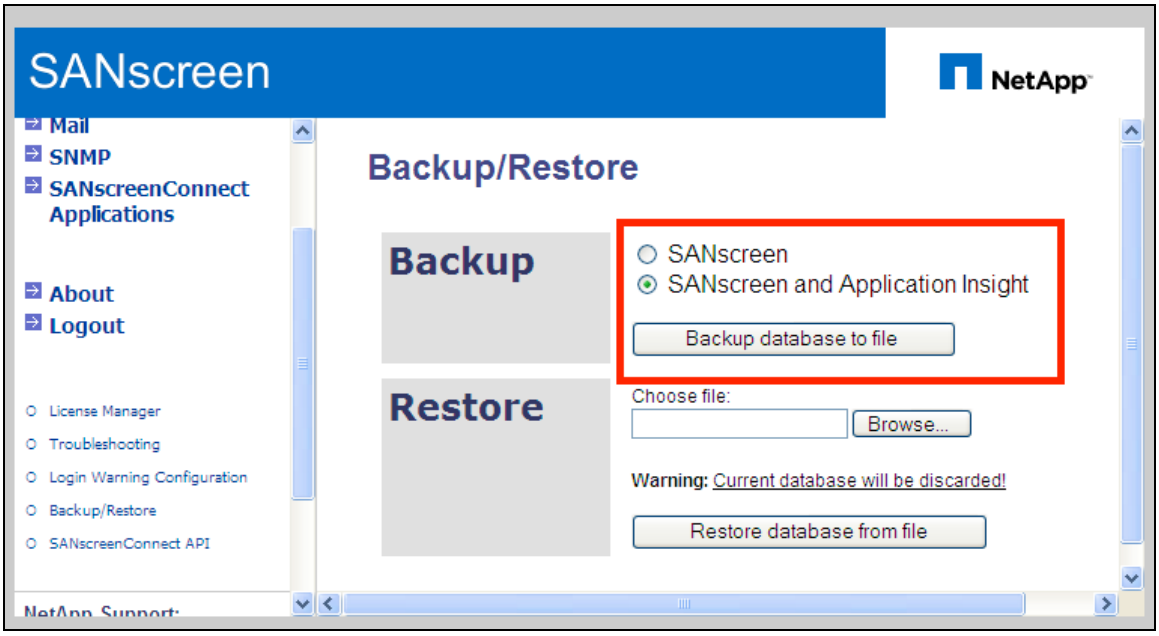

If you select “SANscreen”, you will only backup everything except for Application Insight for Switches and Application Insight for Array Performance. If you select “SANscreen and Application Insight” you will backup the entire product.

The type of backup you select depends on the purpose of the manual backup. If NetApp Global Support requests the backup, they will specify which backup you should perform. If you are performing a manual backup for your own purposes, you are free to determine which backup to perform.

When performing a SANscreen upgrade, the SANscreen upgrade process includes a backup; however, you might want to perform a manual backup in order to properly restore the database after the upgrade has been performed.

Automated Backup of SANscreen Server

SANscreen can be scheduled to be automatically backed up using the Windows Scheduler and a batch file. By default, SANscreen installs the necessary batch file in the “*c:\<SANscreen-install-directory>\backup*” directory. The batch file is called “backup.cmd”.

By default, SANscreen will place all of the resulting database and registry entries that are backed up in the “*c:\<SANscreen-install-directory>\backup\data*” directory. The customer should coordinate with the network backup team to ensure that this directory is backed up nightly so that if a DR recovery is required, the required files can be restored.

7.2 Data Warehouse Database Backup

The SANscreen Data Warehouse can be backed up either manually from the web portal or in an automated way using the Windows Scheduler and a batch file. Both methods are discussed here.

Manual Backup of SANscreen Data Warehouse

The SANscreen DwH can be manually backed up by browsing to the web portal on the SANscreen DwH server and selecting Backup/Restore from the menu on the left. You will be presented with the following backup options as illustrated below in Figure 11:

Figure 11 - Manual Backup of SANscreen DwH

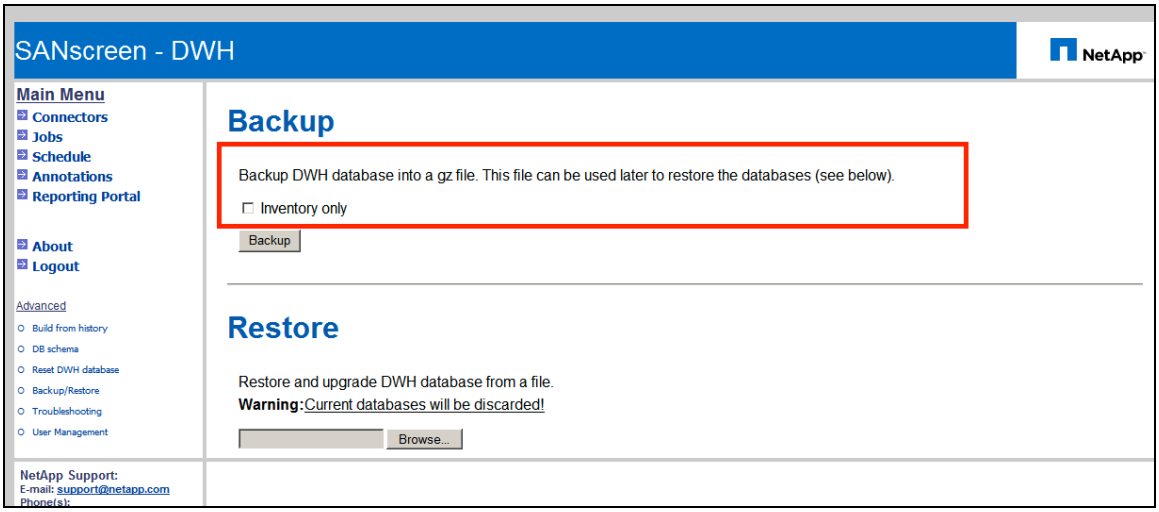

If you select the “Inventory only” checkbox, you will only backup the Inventory Data Mart. Unless you have a specific need (or are directed by the Netapp GSC), you should leave this box unchecked so that all of the DwH Data Marts are backed up.

Automated Backup of SANscreen Data Warehouse

The SANscreen Data Warehouse can be scheduled to be automatically backed up using the Windows Scheduler and a batch file. By default, SANscreen installs the necessary batch file in the “*c:*\<*SANscreenDwH-install-directory*>*backup*” directory. The batch file is called “backup.cmd”.

By default, SANscreen will place all of the resulting database and registry entries that are backed up in the “*c:*\<*SANscreenDwH-install-directory*>*backup**data*” directory. The customer should coordinate with the network backup team to ensure that this directory is backed up nightly so that if a DR recovery is required, the required files can be restored.

7.3 Data Warehouse Content Store Backup

When performing the DwH backup discussed in Section 7.2 above, any custom reports and Framework model changes are not backed up as part of this process. A separate backup procedure must be performed in order for custom reports and Framework customizations to be backed up.

In order to schedule nightly backups of customized content, complete the following steps:

1. Go to the Reporting Portal and log on using the “oadmin” account.
2. Select **Administer SANscreen Reporting Content**.
3. Select the **Configuration** tab.
4. Select the **Content Administration** link on the left panel.

5. Click the **New Export** button on the right side of the page. A backup of SANscreen Reporting Content is called an export. By creating a new Export, we are creating a new job to backup custom content.
6. Give the export a descriptive name, such as “Full SANscreen DwH Backup.”
7. The Description and Screen Tip fields are optional but can be populated.
8. Click **Next** to proceed through the New Export wizard.
9. Select the dialog box labeled **Select the entire Content Store**. **Do not** select the box labeled “Include user account information.”
10. Click **Next** to proceed through the New Export wizard.
11. The “New Archive” radio button should be selected with the name of the Export populated from a previous step above.
12. Click **Next** to proceed through the New Export wizard.
13. Enter a password for the content store. When you export the entire content store, a password must be entered so that the contents are secured.
14. Click **OK** to complete the Export configuration.
15. Review the settings to ensure proper configuration.
16. Click **Next** to proceed.
17. Under the Action dialog, select **Save and schedule** and click **Finish**.
18. Select the proper backup schedule as per your requirements. It is recommended to select a daily backup to be performed during the early AM hours.
19. Click **Finish** to complete the scheduling configuration for the Export job.
    The Export package will now be displayed in the Content Administration screen.

At this point, the Export job is configured and scheduled. When the Export is performed, the Export will be backed up to the *“<SANscreen-DwH-Install-Directory>\cognos\c8\deployment*” directory. The name of the file will be “*<Name-of-Export-Job>.zip*”. The customer should coordinate with the network backup team to ensure that this directory is backed up nightly so that if a DR recovery is required, the required files can be restored.