Technical Report

Clustered Data ONTAP CIFS Auditing Quick Start Guide

Sharyathi Nagesh, NetApp
June 2013 | TR-4189

Summary

This technical report is about native auditing implementation in clustered Data ONTAP® with specific focus on Common Internet File System (CIFS). This document serves as a reference for customers and partners who want to use this feature. Native auditing helps to monitor file activities in NAS environments for diagnostic or reporting purposes. This report covers information on audit configuration, event support, and log format.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Introduction

Native auditing implementation for clustered Data ONTAP is supported from version 8.2 onward. This report describes how to configure auditing in clustered Data ONTAP, access log files, and interpret log information. Native auditing provides a file auditing framework that supports both CIFS and NFS protocols. Auditing in CIFS is based on New Technology File System (NTFS), system access control lists (SACLs), or NFS 4.x access control lists (ACLs).

Native auditing helps to generate and manage audit logs securely and on time. This feature specifically helps to meet industry requirements such as compliance, reliable auditing, intrusion detection, and close to real-time alerting.

1.1 Introduction to Clustered Data ONTAP

Clustered Data ONTAP supports scale-out architecture used to add multiple NetApp® nodes that provide scalability for storage capacity and performance.

Figure 1) Data ONTAP: a scale-out architecture.

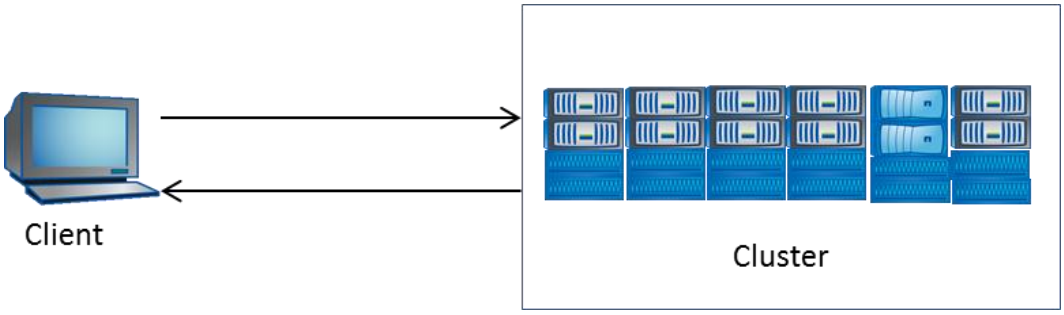

1.2 Introduction to Data ONTAP Global Namespace

The namespace offered by a virtual storage serve, or Vserver, is called a NetApp global namespace. It acts as a container for all storage object servers by the Vserver and identifies each such object with a unique identity. NetApp global namespace supports combining volumes across the cluster to provide a single namespace. Junction points provide means to join volumes together, creating a single namespace. This provides additional flexibility in laying out namespaces when compared to 7-Mode.

Figure 2) Global namespace in clustered Data ONTAP.

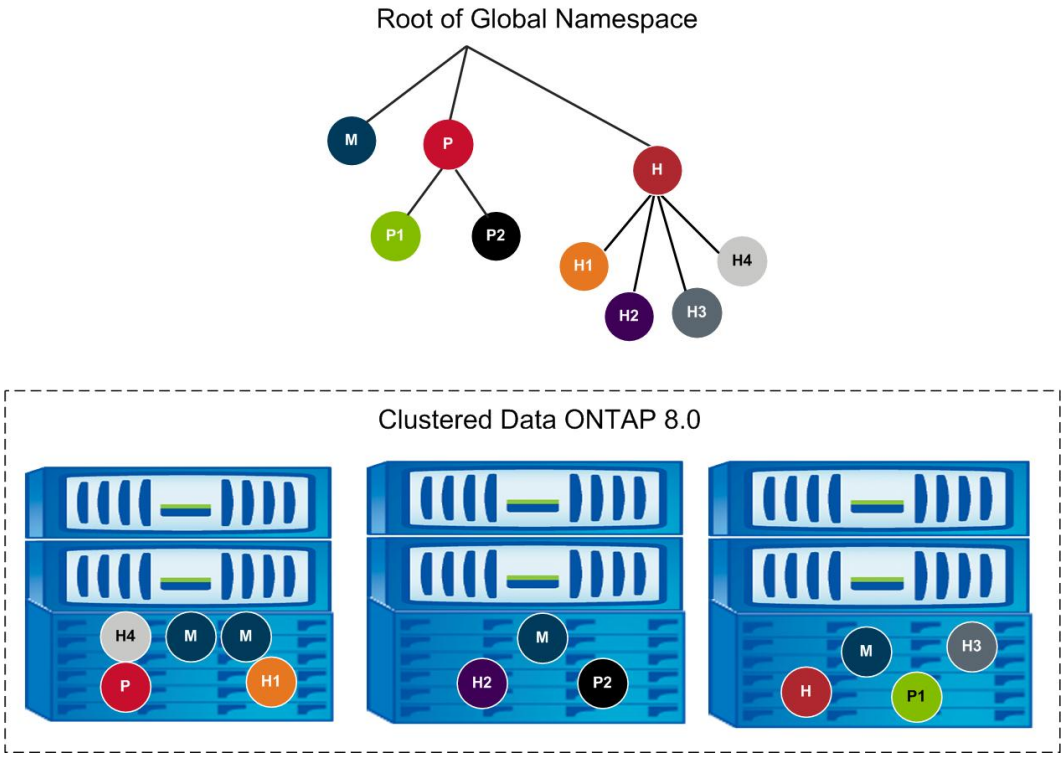

The global namespace created using junction points has the following characteristics:

- Stitching volumes together is transparent to clients.
- CIFS shares can be created on the volume or on qtrees within the volume.

1.3 Introduction to Data ONTAP Native Auditing Implementation

Native auditing has been added to the Data ONTAP 8.2 version to support both CIFS and NFS protocols. Either CIFS or NFS license is required to configure native auditing. To support reliable auditing, audit information is stored on the disk instead of in memory so that in the event of a node or cluster crash, the latest audit information is committed to the disk.

To enhance performance and user experience, this audit information is stored in a specific location in each aggregate. This location is referred as the staging volume. The log records in the staging volume are consolidated into a single log file on a periodic basis. The location and consolidated log file are specified during audit configuration. This is explained in section 2.1.

Creation of staging volumes is transparent to end users. On creating an audit policy on any one Vserver of the cluster, a staging volume will be created on all the aggregates in the cluster. Thereon, all other Vservers will use the existing staging volumes. Each staging volume will consume 2GB of free space, and this needs to be provisioned while configuring auditing. Without the required free space in the aggregate, audit configuration will fail. The staging volume is deleted only when all the Vservers on the cluster delete audit policy.

For example, in Figure 3, the Vserver is spread across two nodes and three aggregates; enabling auditing will create a staging volume in each of the aggregates and by default will take up 2GB of space.

Figure 3) Staging volume creation in clustered Data ONTAP.

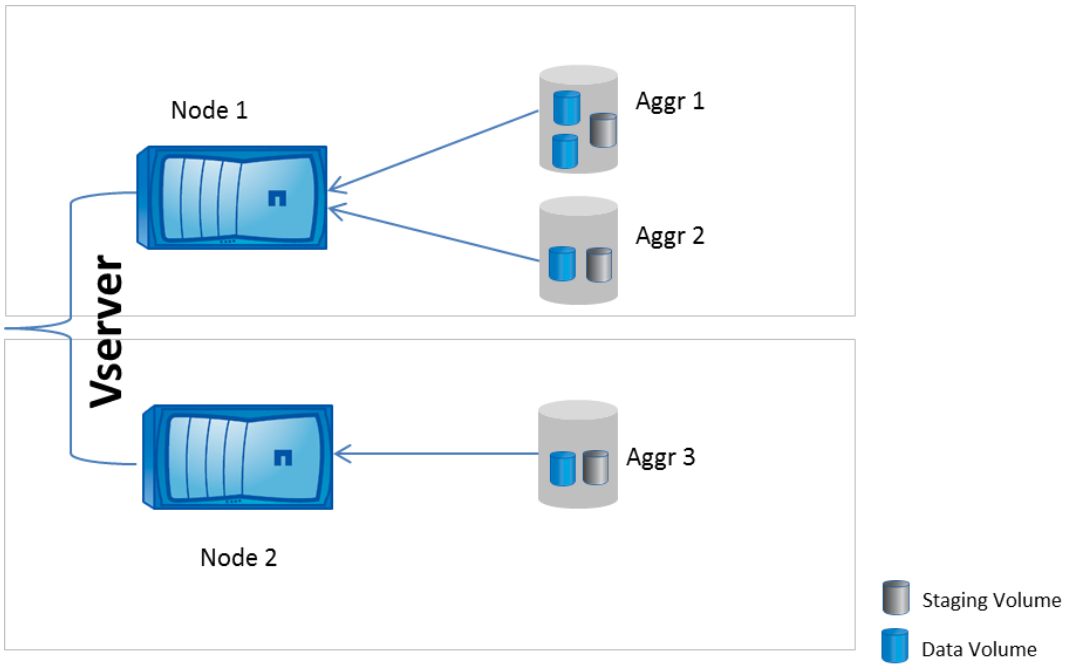

Log consolidation is scheduled every 10 seconds, and scheduling depends on the available CPU bandwidth in the user space. Log consolidation is not configurable. However, log rotation can be configured based on either size or time. This is explained in section 2.1.

Note: On-demand view of the log files is not supported in this version.

2 Configuration of Native Auditing

This section introduces the configuration required to enable auditing on clustered Data ONTAP for Vserver context and configuration of SACLs on files and folders.

2.1 Configuration of Native Auditing on Data ONTAP CLI

Configuration of auditing over Data ONTAP is enabled through the `Vserver audit` command. With this command you can enable or disable auditing, define log location files, manage log rotation, and so on.

Configuration of auditing can be done either under cluster admin or Vserver vsadmin credential. With cluster admin credential, this can be applied on any Vserver in the cluster, while Vserver vsadmin credentials restrict you to only the specific Vserver context. Examples given in the report are for cluster context. In cluster context, you can access/modify/create the audit config for all the Vservers in the cluster. In the Vserver context, you have access only to the Vserver audit config. You require SeSecurity privileges to configure NTFS SACLs. By default, only the local user BUILTIN\administrator has this privilege.

To assign SeSecurity privileges to a user, run the following command:

```
vserver cifs users-and-groups privilege add-privilege -vserver <vserver name> -user-or-group-name
<user> -privileges SeSecurityPrivilege
```

Workflow of Configuring Audit Policy on Vserver

The following flow chart captures the workflow for enabling native auditing on clustered Data ONTAP 8.2 and later. This report primarily explains audit configuration through the CLI; equivalent operations are possible through ONTAPI® as well. ONTAPI information is captured in the appendix.

Figure 4) Workflow to configure audit policy.

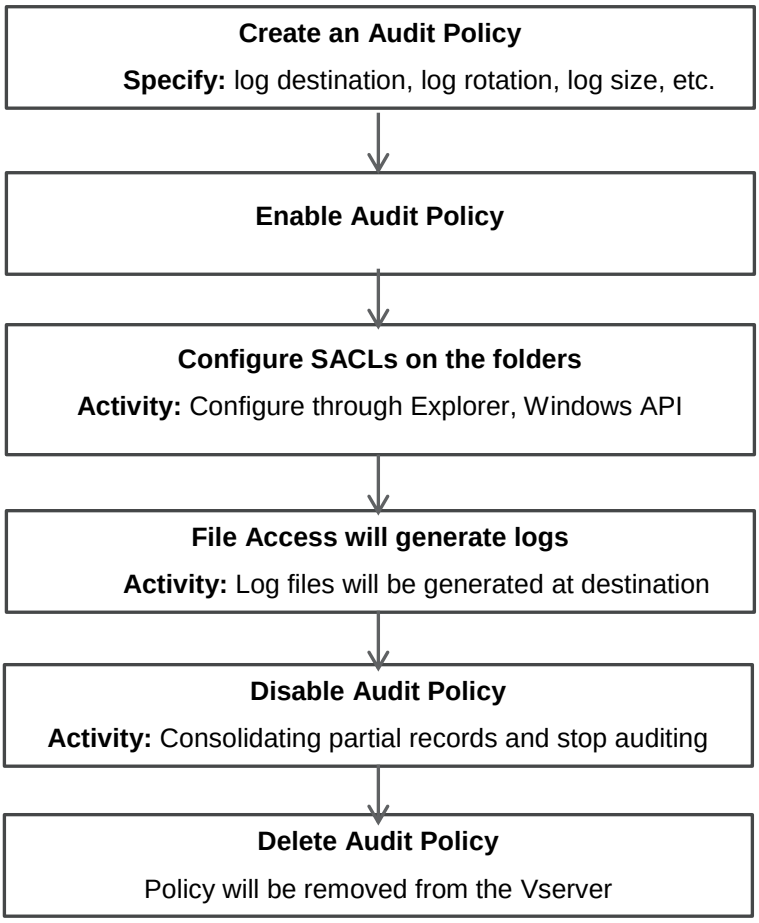

Create an Audit Policy on Vserver

The first step for enabling auditing on a Vserver is to create an audit policy. Vserver name, destination path for saving logs, and log rotation parameters are required as inputs. It is necessary to create a destination path for the audit. This command will either:

- Create new staging volumes if staging volume does not already exist in the data aggregate.
- Share existing staging volume in the data aggregate without compromising on multi-tenancy. Sometimes staging volume may be shared by multiple Vservers.

By default, the staging volume will consume 2GB of space. Without sufficient free space in each aggregate where the data volume resides, auditing will fail. Staging volumes will be created under a cluster Vserver context, and not under the data Vserver context. Staging volume is accessible only by Cserver admin. Cserver admin can resize staging volume in the diag-mode using vol resize option.

Only one active policy can be created per Vserver. The following example explains this.

Creating Policy Based on Log Size

In the following example, we are creating an audit policy for the specified Vserver with log location specified in the destination field. The destination path is a path to the folder location and should have been precreated. Size of log file is specified through the rotate-size field. Rotate-limit will specify the maximum number of log files that will be kept in the specified destination. Log files beyond this value will be overwritten. A value of zero indicates unlimited log files; in this case, number of log files will be limited by the available free space in the destination.

```
vserver audit create -vserver <vserver_name> -destination <unix path> -rotate-size <size MB> -
rotate-limit <Number of log files>
```

When log files reach the specified rotate-size, it triggers log rotation. Minimum value for rotate-size is 100MB.

Creating Policy Based on Time

This example illustrates creating an audit policy for a specified Vserver with the log location specified in the destination field with the following command:

```
vserver audit create -vserver <vserver_name> -destination <unix path> -rotate-schedule-minute
<minute of the hour> -rotate-limit <Number of log files>
```

Note: Rotate-limit will specify maximum number of log files that will be kept in the specified destination.

Rotate-schedule field will define how often the audit log file will be rotated. For more details about log rotation, refer to section 3.2.

Enable an Audit Policy on Vserver

After the audit policy is created, it needs to be enabled for audit action to begin. Enabling audit policy is a simple operation and can be executed as follows:

```
vserver audit enable -vserver <vserver_name>
```

Disable an Audit Policy on Vserver

This command will consolidate any partial audit records present in the staging volumes into the consolidated audit log file and stop any further logging of audit records:

```
vserver audit disable -vserver <vserver_name>
```

Delete an Audit Policy on Vserver

Deleting the audit policy can free space for data by deleting the staging volume. If the staging volume is used by another Vserver, the deletion will not be possible. Staging volumes will be deleted only when all the Vserver references are deleted.

An audit policy can be deleted as follows:

```
vserver audit delete -vserver <vserver_name>
```

Modify an Audit Policy on Vserver

The modify command modifies the parameters previously created for an audit policy. This command can be used to modify log destination location and rotation policies such as number of concurrent log files, log rotation triggers, and so on.

```
vserver audit modify -vserver <vserver_name>  -destination <unix path> -rotate-size 100MB -
rotate-limit 0
```

2.2 Configuration of SACLs on the Storage Object

After enabling audit policy at the Vserver level, configure SACLs on files, folders, or shares.

SACLs can be configured on files and folders as follows:

- By using client applications such as Windows Explorer
- From script/application using appropriate Windows APIs
- From fsecurity (file-directory) command through the CLI

SACLs can be configured on shares as follows:

- By setting SACLs on the root of the share from the Windows client

Note: Windows RPCs are currently not supported. Configuration through MMC or dependent application will not be possible.

Set of Object Access Operations Supported

Logging of file and folder access operations is supported by the auditing framework. Equivalent Windows object access operation ID is captured in Table 1 for better understanding. Both success and failure auditing is supported for each of these operations.

The supported events can be broadly classified as follows. The mapping between these events with Windows events is on a best effort basis. Some of the information present in Windows event might not be possible to provide in the Data ONTAP environment; for example, Windows audit records capture process ID and process name, which is not possible in Data ONTAP audit records.

Table 1) Supported access operations.

Windows Event ID	Event Name	Description
4656	Open object	A handle to an object is requested. This corresponds to event ID 560 in W2k3 and before.
	Create object	
4663	Read object	An attempt was made to access an object. This corresponds to event ID 567 in W2k3 and before. This event documents operations performed against data objects. This event logs operation that takes place between the open and close events for the object.
	Write object	
	Get object attributes	
	Set object attributes	
4664	Hard link	An attempt was made to create a hard link. Hard link is a pointer to another file in the same file system.
9999	Rename object	Added by NetApp. This captures the object rename operation. This is currently not supported by Windows as single event.
9998	Unlink object	Added by NetApp. This captures the object unlink operation. This is currently not supported by Windows as single event.

Note: Currently, NetApp does not support the close object event, event ID 4658, because it was creating unwanted notifications.

Note: During the delete operation, only event ID 4656 is generated, because it has all the information required for identifying the event.

Note: Even though 4656 is open event, the desired access field will specify the intent of opening the file from which the event can be identified.

Read and write events are optimized to log only the first read and write to avoid unwanted notifications.

For more information about description of security events, refer to MS KB Article ID: 947226: Description of security events in Windows Vista and in Windows Server 2008.

3 Managing the Audit Logs

3.1 Audit Log Record Format

Path of File in Notifications

Path information provided in logs will include only the relative path from the root of the containing volume. Users need to construct the absolute path information from the volume ID, also called msID, and information available in the file handler field of the log record.

Let us understand this with an example:

- If there are two volumes—vol0 and vol1—with vol0 joined on / and vol1 on /home/userA, the path /home/userA/division/team/prod has /home/userA in vol0 and /division/team/prod in vol1.
- When the file in /home/userA/division/team/prod is accessed, only the path /division/team/prod is available in the notification. The mount point of the volume vol1, which is /home/userA, is called the junction point of the volume vol1.
- To construct the absolute path name, information available outside the log records must be used. Clustered Data ONTAP can be queried over with `volume-get-iter` ONTAPI call with unique msID to retrieve its junction point. User developing this support can cache the msID to junction path mapping to avoid calling it every time. Since the namespace will not change frequently, one-time operation to build the namespace should be sufficient.

Note: When a new volume is added, Vserver has to be queried again to find the junction point. Rarely, if the volumes are remounted on a new junction path, the global namespace will undergo change. In such instances, periodic querying with `volume-get-iter` to update the volume–junction path mapping is required.

Schema of Log Records

Information in the log files is stored in NetApp proprietary format. Align the format of log records closely with that of Windows .evtx format.

The naming convention of the consolidated log file follows this format, which is not configurable:

```
audit_<vservername>_D<yyyy>-<MM>-<DD>-T<HH>-<MM>-<SS>_milliseconds.Xml
```

3.2 Audit Log Rotation

When log size and log rotation parameters are not specified, the default values will be taken.

The default value is log rotation based on log size of 100MB. New logs will be created as long as the destination volume has free space. the number of concurrent files kept for log management can be changed with rotate-limit parameter.

The log-rotation can be configured for time or size.

Log Rotation Based on Time

Log rotation is based on calendar date and time. The parameters supported are:

- Month: Month of the year

- Day: Day of the week or month
- Time: Specific hour and minute of the day. Specifying minutes is mandatory. For example, if you specify minute as 45, every 45th minute of the hour, a new log file will be generated.

This will create new log files on specific days of the week:

```
vserver audit modify -vserver <vserver_name> -destination <unix path>  -rotate-schedule-month
February,March -rotate-schedule-dayofweek Sunday -rotate-schedule-hour 22 -rotate-schedule-minute
45 -rotate-limit <Number of log files>
```

Log rotation can be based on calendar date and time. The parameters supported are:

```
vserver audit modify -vserver <vserver_name> -destination <unix path>  -rotate-schedule-month
February,March -rotate-schedule-day 22 -rotate-schedule-hour 22  -rotate-schedule-minute 45 -
rotate-limit <Number of log files>
```

Log Rotation Based on Log Size

Log rotation can be based on log size. This can be configured as follows:

```
vserver audit modify -vserver <vserver_name> -destination <unix path> -rotate-size <size MB> -
rotate-limit <Number of log files>
```

3.3 Accessing Audit Logs

Audit logs will be saved in the destination location specified during audit configuration. The logs can be accessed over the data access path. The destination path and the file can be accessed through CIFS shares. Access can be restricted with share-level ACLs or through folder- or file-level ACLs. Similar access is possible through the NFS export path as well.

Note: Access to audit logs is through a pull mechanism and retrieved over NFS, CIFS, or other file access protocol methods. Unlike syslog framework, logs cannot be accessed through push mechanism.

3.4 Partial Logs

During cluster failovers, audit engine will not be able to consolidate the complete Vserverized logs. In that case, audit log file name will indicate that it is a partial file. As soon as the node boots up, the audit engine will consolidate the records and order the node chronologically.

Appendix

Audit Guarantee Feature

When auditing is highly critical, either because of organizational policies or because of regulatory requirements, the auditing guarantee feature should be used. This will make sure the log records are written to disk before file operation completes. If log records cannot be committed to the disk either because of insufficient space or because of some issues, client I/O will be blocked. This feature is enabled by default.

This feature is available in the diag-mode and can be configured as follows:

```
vserver audit modify -vserver <vserver_name> -destination <unix path> -rotate-size 100MB -rotate-
limit 0 -audit-guarantee true|false
```

Relevant ONTAPI Interfaces for Configuring Auditing

The auditing features can be configured either through the command line interface (CLI) or through APIs. Data ONTAP APIs (ONTAPI or Manage ONTAP®) supported with auditing allow configuring auditing remotely. Information about ONTAPI interfaces can be found from NM-SDK documentation available at NetApp Developer Community. The developer forum is a useful reference for developers with technical queries.

Note: Cluster ONTAPI interfaces are supported from NM-SDK 4.2 and later.

Table 2) List of audit ONTAPI interfaces added to clustered Data ONTAP.

ONTAPI Interfaces	Description
fileservice-audit-config-get	This API will provide audit config details of a particular Vserver.
fileservice-audit-config-get-total-records	This API gets the total number of audit config entries/records in the table.
fileservice-audit-config-get-iter	This API provides audit config details for all the Vservers.
fileservice-audit-config-create	This API creates audit config for a particular Vserver.
fisleservice-audit-config-destroy	This API deletes audit config for a particular Vserver.
fileservice-audit-config-modify	This API modifies the audit config for a particular Vserver.
fileservice-audit-enable	This API enables auditing for a particular Vserver.
fileservice-audit-disable	This API disables auditing for a particular Vserver.

Executing the ONTAPI Interfaces

Some of the ONTAPI interfaces run only in clustered context and some only in Vserver context, but there are a few ONTAPI interfaces that run in both. Keep this in mind before calling ONTAPI.

- **Cluster APIs.** These APIs are executed against the `cluster-mgmt` IP using cluster administration credentials.
- **VServer APIs.** These APIs are executed either by:
  - Calling ONTAPI against a VServer LIF using Vserver admin credentials

    or
  - Calling ONTAPI against the cluster-mgmt IP with cluster admin credentials, but using tunneling

Using Fsecurity to Configure SACLs on Files and Folders

The following options can be used to configure SACLs on files and folders:

- From Windows Explorer or using any tool that configures SACLs based on Windows APIs.
- By using the fsecurity command in clustered Data ONTAP. The fsecurity command is renamed to file-directory and is available under Vserver security command structure.
- Configuring SACLs through any application that uses Windows RPC, for example, MMC, is not currently supported.

Best Practices and Recommendations

- Make sure that enough free space is available to create the staging volume. Each aggregate, including aggr0, requires 2GB free space to enable auditing.

- Make sure that the destination path and volume have enough space for audit logs. If filled up, it will start affecting the client operation.
- The space guarantee feature is enabled by default. In case the destination volume gets filled up, it will affect the client operation. In this case, disable the guarantee feature.
- Enabling SACLs will affect the performance. NetApp recommends the user to judiciously choose the SACLs. Configure only for create and delete operations if that is sufficient to meet the security guidelines.

Note: Be aware that creating SACLs will enable write auditing as well. Writes are optimized to log only the first write operation within the open-close context to enhance performance.

References

The following references were used in this technical report:

- File Access Protocol Management Guide for Data ONTAP 8.2
- For description of security events in Windows Vista® and in Windows Server® 2008: http://support.microsoft.com/kb/947226

Version History

Version	Date	Document Version History
Version 1.0	June 2013	Initial release