Technical Report

NetApp Best Practices for Clearing a Data Spill in Data ONTAP 7-Mode

Mike Scanlin, NetApp
January 2013 | TR-4121

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Solution Architecture

1.1 Clearing a Data Spill in Data ONTAP 7-Mode

Overview

A data spill is a security incident that results in the transfer of classified or sensitive information (for example, personally identifiable information [PII] or contract-sensitive information) to unaccredited and unauthorized information systems (ISs), applications, or media. NetApp provides procedures for postincident eradication of a data spill on a NetApp® FAS (Data ONTAP® operating in 7-Mode) system.

NIST Special Publication 800-88, Guidelines for Media Sanitization, defines clearing as a media sanitization process that protects the confidentiality of information against a robust keyboard attack. Clearing must prevent information retrieval by data, disk, or file recovery utilities and resist keystroke recovery attempts executed from standard input devices and data-scavenging tools. In accordance with NIST SP 800-88, overwriting storage space on media with nonsensitive data is an approved clearing method, but simple deletion is not.

NetApp provides procedures for clearing a data spill on a NetApp FAS (7-Mode) system and for preventing recovery of spilled data by means of keyboard attack (recovery of information from a hard drive by using simple keyboard commands).

Note: The procedures outlined in this document are intended for ISs that contain data categorized as Sensitive, Classified, and up to and including Secret. Consult with the organization's information security office or Designated Approving Authority (DAA) before performing these procedures on an IS with a classification above Secret. When IS confidentiality requirements require protection against a laboratory attack (recovery of information from a hard drive by means of disassembly, special forensic tools, and specially trained personnel), NetApp recommends using the NetApp encryption solutions.

NetApp clearing methods comply with the NIST SP 800-88 guidelines for clearing data from magnetic disk media. The NetApp Disk Sanitization procedure applies overwrite guidance according to the clearing and sanitization matrix in the Department of Defense (DoD) Defense Security Service (DSS) DSS Industrial Security Letter (ISL) 2007-1. The procedures outlined in this document promote sanitization decisions consistent with an organization's risk response strategy and requirements for IS confidentiality and availability.

Preconditions for using this information are that:

- A data spill has occurred and has been identified.
- The data spill is contained.
- A preliminary assessment of the scope, nature (inadvertent or malicious), and classification of the data spill is complete.

NetApp provides operational procedures for the following use cases:

- Prepare a contaminated NetApp FAS (7-Mode) system for single-pass overwrite or disk sanitization.
- Perform a single-pass overwrite of zeros to each block of every disk to clear (overwrite) data on disks in a NetApp FAS (7-Mode) system.
- Apply DoD-specified overwrite patterns to clear (overwrite) data on disks in a NetApp FAS (7-Mode) system.
- Selectively remove and retain data on disks in support of statutory compliance or forensic analysis requirements.
- Back up contaminated Snapshot™ copies to tape media using the `Dump` command.

  Note: For more details about these use cases, refer to the "Description of Use Cases" section.

Manage Risk Rather Than Definitions

The current debate about what qualifies as sanitization involves several competing definitions. For example:

- According to NSA/CSS Storage Device Declassification Manual, Dec 2007, sanitization is the removal of information from the storage device so that data recovery using any known technique or analysis is prevented. Sanitization includes the removal of data from the storage device, as well as the removal of all labels, markings, and activity logs. The method of sanitization varies, depending upon the storage device in question, and may include degaussing, incineration, shredding, grinding, embossing, chemical immersion, and so on.
- According to the Committee on National Security Systems (CNSS) Instruction No. 4009, National IA Glossary, 26 April 2010, sanitization is a general term referring to the actions taken to render data written on media unrecoverable by both ordinary and, for some forms of sanitization, extraordinary means.
- According to the DoD 5220.22-M, National Industrial Security Program Operating Manual (NISPOM), February 28, 2006, sanitization is the process of removing the data from media before reusing the media in an environment that does not provide an acceptable level of protection for the data that was in the media before sanitizing. IS resources shall be sanitized before they are released from classified information controls or released for use at a lower classification level.
- According to the NIST Special Publication 800-88, September 2006, sanitization refers to the general process of removing data from storage media, such that there is reasonable assurance that the data may not be easily retrieved and reconstructed.

Rather than debate definitions, NetApp focuses on managing risk and providing data spill remediation alternatives consistent with NIST SP 800-37, Guide for Applying the Risk Management Framework to Federal Information Systems. The concepts and operational procedures described in this document enable organizations to "…more effectively manage information system–related security risks in highly diverse environments of complex and sophisticated cyber threats, ever-increasing system vulnerabilities, and rapidly changing missions." Figure 1 illustrates the risk management framework (RMF).

Figure 1) Risk management framework (from NIST Special Publication 800-37).

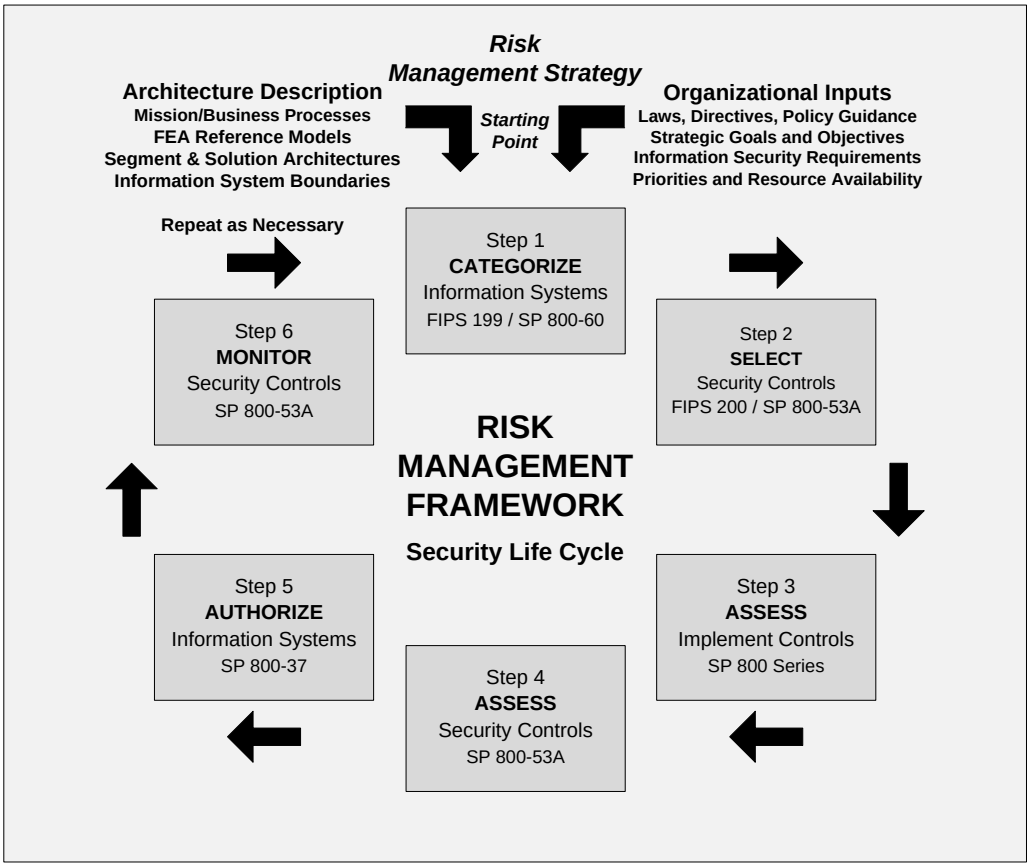

NetApp Provides Efficiency, Performance, Availability, and Security

NetApp storage technologies balance efficiency, performance, availability, and security consistent with an organization's business needs and risk management objectives. NetApp recognizes that in certain environments, "the systems' data content or mission purpose is of such value that aggressive tradeoffs in favor of security outweigh the potential negative consequences to other useful system attributes such as legacy applications or interoperability with other systems." (Source: NIST, SP 800-70, National Checklist Program for IT Products—Guidelines for Checklist Users and Developers) In all instances, NetApp people, process, and technology enable customers to make the right decision for their business and operational environment.

Data ONTAP and WAFL Provide Enhanced Confidentiality for Simple Deletes

NetApp core technologies that provide efficiency, performance, and availability also offer an added measure of information assurance to the eradication phase of a data spill. In "Security and Usability—Designing Secure Systems That People Can Use," author Simson Garfinkel outlines a sanitization taxonomy, as shown in Table 1, that informatively differentiates file deletion from overwriting and highlights the protection that Write Anywhere File Layout (the WAFL® system) provides if sanitization must be deferred for any period of time.

Table 1) A sanitization taxonomy (from "Security and Usability—Designing Secure Systems That People Can Use" (used by permission of the author).

Level	Type of Data	Description
0	Regular files	Information contained within the file system, including file names, file attributes, and file content. By definition, there has been no attempt to sanitize the information that is contained within Level 0 files. Level 0 also includes information that is written to the disk as part of any sanitization attempt. For example, if a copy of Windows® 95 is installed on a hard drive in an attempt to sanitize the drive, then the files contained within the `C:\WINDOWS` directory would be considered Level 0 files. No special tools are required to retrieve Level 0 data.
1	Temporary files	Temporary files, including print spooler files, browser cache files, files for helper applications, and files in recycle bins. Most users either expect that these files will be deleted automatically in time or are not even aware that these files exist. **Note:** Level 1 files are a subset of Level 0 files. Experience has shown that it is useful to distinguish this subset, because many naive users will overlook Level 1 files when they are browsing a computer's hard drive to see if it contains sensitive information. Although no special tools are required to retrieve Level 1 data, special training is required so that the operator knows where to look.
2	Deleted files	When a file is deleted from a file system, most operating systems do not overwrite the blocks on the hard disk on which the file is written. Instead, they simply remove the reference to the file from the containing directory. The file's blocks are then placed on the free list. These files can be recovered using traditional undelete tools such as Norton Utilities.
3	Retained data blocks	Data that can be recovered from a disk but that does not obviously belong to a named file. Level 3 data includes information in slack space, swap space for virtual memory, and Level 2 data that has been partially overwritten so that an entire file cannot be recovered. One common source of Level 3 data is disks that have been formatted with the Windows `FORMAT` command or the UNIX® `newfs` command. Even though these commands give the impression that they overwrite the entire hard drive, they do not, and the vast majority of the information on a formatted disk can be recovered with Level 3 tools. Level 3 data can be recovered using advanced data recovery tools that can unformat a disk drive and using special-purpose forensics tools.
4	Vendor-hidden data	Data blocks on the drive that can be accessed only by using vendor-specific commands. Level 4 includes the drive's controlling program, blocks used for bad-block management, and the host-protected area of modern hard drives.
5	Overwritten data	Many individuals maintain that information can be recovered from a hard drive even after it is overwritten. Level 5 is reserved for such information.

Note: Author Simson Garfinkel contends that "[s]imply overwriting user data with one or two passes of random data is probably sufficient to render the overwritten information irrecoverable" due to the high density of modern hard disk drives. (Garfinkel & Shelat, 2003, p. 21) According to "Overwriting Hard Drive Data: The Great Wiping Controversy," the authors "…categorically state that there is a minimal (less than a 0.01%) chance of recovering any data on a NEW and unused drive that has a single raw wipe pass (not even a low-level format). In the cases where a drive has been used (even being formatted for use) it is not possible to recover the information—there

is a small chance of bit recovery, but the odds of obtaining a whole word are small." (Wright, Kleiman, & Sundhar, 2008, p. 251)

When operations prohibit immediate sanitization of a contaminated storage system, WAFL's data layout algorithm allows a simple file delete to provide a significantly higher level of confidentiality than that of most other operating systems. Enabling a disk sanitization license and restricting console access to a trusted administrator prevent both inadvertent and malicious access to contaminated data that has been deleted from a NetApp storage array. In keeping with other NIST SP 800-53 management, operational, and technical controls, WAFL enables a simple file delete to mitigate risk sufficiently until conditions permit a comprehensive clearing of the contaminated storage system. Figure 2 shows how WAFL provides enhanced confidentiality for simple deletes.

Figure 2) Enhanced confidentiality for simple deletes with WAFL.

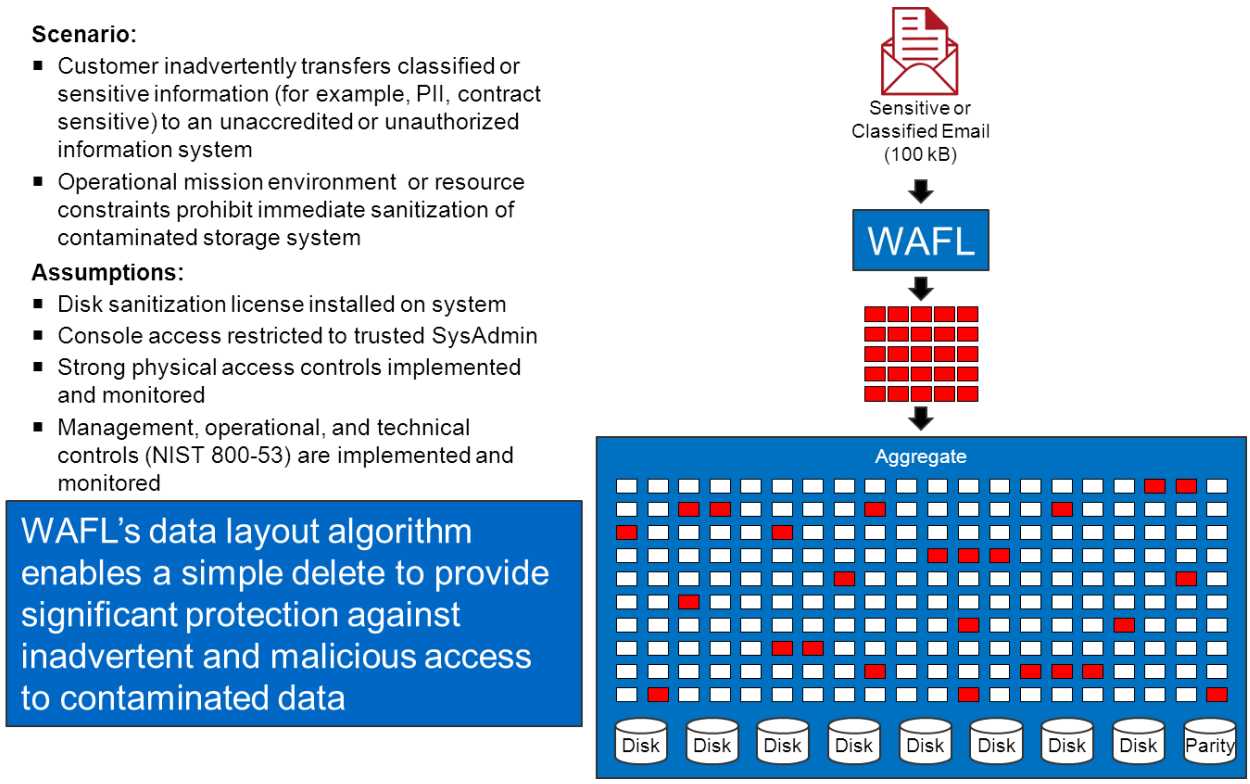

Description of Use Cases

NetApp provides operational procedures for the following use cases.

Prepare a contaminated NetApp FAS (7-Mode) system for single-pass overwrite or disk sanitization.

This use case prepares a contaminated aggregate on a NetApp FAS (7-Mode) system for sanitization in accordance with NIST SP 800-88 guidelines for clearing data from magnetic disk media. At end state, users have been migrated to a target (clean) aggregate, operations have resumed on the target aggregate, and the source (contaminated) aggregate is ready to be cleared (overwritten) by means of single-pass overwrite or NetApp Disk Sanitization.

Perform a single-pass overwrite of zeros to each block of every disk to clear (overwrite) data on disks in a NetApp FAS (7-Mode) system.

This use case clears (overwrites) all disks in a contaminated aggregate in order to prevent data recovery by means of a robust keyboard attack. At end state, the contaminated aggregate and all related volumes are destroyed. All disks that were RAID group members of the contaminated aggregate are overwritten once with zeros and designated as hot spares.

Note: Remnant magnetic patterns might exist on storage media after clearing, but no decipherable data or information is recoverable by means of a robust keyboard attack.

Disks are available for reuse consistent with organizational policy and risk response strategy. Serial numbers of all cleared disks are manually recorded and posted to the `/etc/messages` file.

Note: Single-pass overwrite clears a contaminated aggregate in one third of the time required to complete a default three-pass overwrite using NetApp Disk Sanitization.

Apply DoD-specified overwrite patterns to clear (overwrite) data on disks in a NetApp FAS (7-Mode) system.

This use case clears (overwrites) an aggregate in order to prevent data recovery by means of a robust keyboard attack. At end state, all disks that were RAID group members of the contaminated aggregate are overwritten with a DoD-specified overwrite pattern and designated as hot spares.

Note: Remnant magnetic patterns might exist on storage media after clearing, but no decipherable data or information is recoverable by means of a robust keyboard attack.

Disks are available for reuse consistent with organizational policy and risk response strategy. Serial numbers of sanitized disks are written to the `/etc/sanitized_disks` file.

Selectively remove and retain data on disks in support of statutory compliance or forensic analysis requirements.

This use case retains and preserves select data from a contaminated aggregate when required for compliance with statutory requirements or for postincident forensic analysis of a data spill. At end state, select contaminated data has been migrated (to a new volume or aggregate) for retention and preservation. The source (contaminated) aggregate is ready for disk sanitization.

Back up contaminated Snapshot copies to tape media using the `Dump` command.

This use case retains and preserves select data from a contaminated aggregate when required for compliance with statutory requirements or for postincident forensic analysis of a data spill. At end state, select contaminated data has been migrated to tape media and properly secured for the duration of the retention period. The contaminated source aggregate is ready for disk sanitization.

Single-Pass Overwrite, Disk Sanitization, and Selective Disk Sanitization

Note: The following section and procedures do not apply to drives with NetApp Storage Encryption built into them. They are covered in a different section of the documentation.

NetApp recommends following best practices for creating and backing up aggregates containing data to be sanitized.

Best Practice

- Make sure aggregates containing sensitive data are not larger than required for functionality and performance. Oversized aggregates consume more time, disk space, and bandwidth during sanitization.
- Avoid backing up aggregates containing sensitive data to aggregates containing large amounts of nonsensitive data. This practice reduces the resources required to move nonsensitive data before sanitizing sensitive data.
- Employ disciplined Snapshot copy management:
  - Apply the principle of least privilege to Snapshot copies. By default, every volume contains a directory named `.snapshot` (NFS) or `~snapshot` (CIFS) through which users can access previous versions of files in that directory by means of an NFS or a CIFS client. Restrict access to Snapshot copies as consistent with user duties and functions.
  - Implement a Snapshot policy consistent with IS availability requirements that prevents unregulated proliferation of Snapshot copies.

Single-Pass Overwrite

Single-pass overwrite clears all data from SAS and SATA disks by performing a single-pass overwrite of zeros to each block of every disk assigned to the owning storage controller. The procedure leverages studies suggesting that a single-pass overwrite can prevent data recovery by means of a laboratory attack as well as a keyboard attack. For details about research findings, refer to Overwriting Hard Drive Data: The Great Wiping Controversy.

Single-pass overwrite clears a contaminated aggregate in one-third of the time required to complete a default three-pass overwrite using NetApp Disk Sanitization.

Disk Sanitization

Disk sanitization clears data from SAS and SATA disks by overwriting disks with DoD-specified overwrite patterns in order to prevent recovery of the original data by means of a robust keyboard attack.

Note: A disk sanitization license must be installed in order to perform disk sanitization.

The disk sanitization process uses three successive default (or user-specified) overwrite patterns for up to seven cycles per operation. Depending on the disk capacity, the overwrite patterns, and the number of cycles, the process can take several hours. Sanitization runs in the background. It is possible to start, stop, and display the status of the sanitization process.

After the sanitization process is started, Data ONTAP begins sanitizing each of the specified disks. The process consists of a disk format operation, followed by the specified overwrite patterns repeated for the specified number of cycles.

Note: The formatting phase of the disk sanitization process is skipped on advanced technology attachment (ATA) disks.

If the sanitization process is interrupted by power failure, system panic, or the user, the sanitization process must be repeated from the beginning in order for sanitization to take place. When the sanitization process is complete, the specified disks are in a sanitized state but are not automatically returned to spare status.

Disk sanitization is subject to the following limitations:

- It is not supported in takeover mode for systems in a high-availability (HA) configuration.

  Note: If a storage system is disabled, it remains disabled during the disk sanitization process.
- It cannot be carried out on disks that were failed because of read or write problems.
- It does not perform its formatting phase on ATA drives.

- If a random pattern is used, it cannot be performed on more than 100 disks at one time.
- It is not supported on array logical unit numbers.
- It is not supported on solid-state drives.
- If both SCSI Enclosure Services (SES) disks are sanitized in the same Embedded Switched Hub (ESH) shelf at the same time, errors occur on the console about access to that shelf, and shelf warnings are not reported for the duration of the sanitization. However, data access to that shelf is not interrupted.

Selective Disk Sanitization

Selective disk sanitization clears data in specified files or volumes while preserving all other data located on the affected aggregate for continued user access. Selective disk sanitization is composed of three primary tasks:

1. Deletion of files, directories, or volumes containing the data to be sanitized from the aggregate that contains them.
2. Migration of data to be preserved to a new set of disks in a destination aggregate on the same storage system (using the `ndmpcopy` command or qtree SnapMirror® technology).
3. Destruction of the original aggregate and clearing (overwriting) of all disks that are RAID group members in that aggregate by using NetApp Disk Sanitization.

1.2 NetApp FlexVol

FlexVol Limits and Maximums

Refer to the NetApp product documentation associated with the specific version of Data ONTAP for a detailed explanation of product limits and scaling considerations consistent with various storage efficiency and other optional settings.

FlexVol Default Snap Reserve

Be sure to take into account the space that snap reserve can consume (20% of the volume by default when a FlexVol® volume is created).

FlexVol Language

By default, any new FlexVol volumes use the language of the root volume.

Thin Provisioning

Thin provisioning is covered in a dedicated section.

2 Operational Procedures

Use Case	Procedure
Prepare a contaminated NetApp FAS (7-Mode) system for single-pass overwrite or disk sanitization.	Prepare Contaminated NetApp FAS (7-Mode) System for Single-Pass Overwrite or Disk Sanitization
Perform a single-pass overwrite of zeros to each block of every disk to clear (overwrite) data on disks in a NetApp FAS (7-Mode) system.	Clear Contaminated NetApp FAS (7-Mode) System by Using Single-Pass Overwrite

Use Case	Procedure
Apply DoD-specified overwrite patterns to clear (overwrite) data from disks in a NetApp FAS (7-Mode) system.	Remove Data from Disks in NetApp FAS (7-Mode) System by Using Disk Sanitization
Selectively remove and retain data on disks in support of statutory compliance or forensic analysis requirements.	Selectively Remove and Retain Data on Disks by Using Selective Disk Sanitization
Back up contaminated Snapshot copies to tape media using the `dump` command.	Back Up Contaminated Snapshot Copies to Tape Media by Using Dump Command

2.1 Prepare Contaminated NetApp FAS (7-Mode) System for Single-Pass Overwrite or Disk Sanitization

This procedure prepares a contaminated aggregate on a NetApp FAS (7-Mode) system for sanitization in accordance with NIST SP 800-88 guidelines for clearing data from magnetic and optical disk media. The goal of this procedure is to rapidly migrate and resume operations on a new aggregate while preparing the contaminated aggregate for clearing (overwriting) by performing single-pass overwrite or disk sanitization. To prepare a contaminated aggregate for sanitization, complete the following steps.

1. Apply and enforce user access restrictions to the storage system volume or logical unit number (LUN) consistent with IS confidentiality requirements and the organization's Cyber Incident Response Plan.
2. Stop all applications that write to the aggregate that is being sanitized.
3. Back up volume-specific storage controller configuration.

```
config dump -v
```

4. From a Windows or a UNIX client, delete from the active file system the directories or files containing data to be selectively sanitized. Use the appropriate Windows or UNIX command, as shown in this example.

```
rm /unixdir/unixfile.doc
```

5. Remove NFS and CIFS access to all volumes in the contaminated aggregate.
6. From the Data ONTAP CLI, delete all volume Snapshot copies of the FlexVol volumes that previously contained the deleted files and directories.

```
snap delete -V -a <<var_vol01>>
```

7. Delete all Snapshot copies referring to data blocks in the contaminated source file.

   Note: Include all hourly, nightly, and weekly versions.

```
snap delete <<var_vol01>> name
```

8. Create a clean target aggregate with sufficient space to store all data from the contaminated aggregate.

```
aggr create <<var_aggr01>>-d disk1 [ disk2 ... ] [ -d diskn [ diskn+1 ... ] ]
```

9. On the storage controller, enable the request daemon.

```
myhost> ndmpd on
```

10. Migrate clean files to the destination on the target aggregate.

```
myhost> ndmpcopy /vol/source_path
        /vol/destination_path
```

Note: This command migrates data from a source path (`source_path`) to a different destination path (`destination_path`) on the same storage controller (`myhost`).

11. Copy clean Snapshot copies (those that do not refer to data blocks in the contaminated source file) to the target volume.
12. Migrate users to the target aggregate and resume operations on the target aggregate.
13. Offline the contaminated volume(s).

```
vol offline {volname|plexname}[-t cifsdelaytime]
```

Note: If a volume contains CIFS shares, use the `-t` option to warn users before taking the volume offline. The `cifsdelaytime` argument specifies the number of minutes to delay before taking the volume offline, during which time CIFS users are warned of the pending loss of service. A time of 0 means that the volume should be taken offline immediately and without warning. CIFS users can lose data if they are not given a chance to terminate applications gracefully. The root volume cannot be taken offline.

Note: If you attempt to take a volume offline while any files contained by that volume are open, the `volume offline` command fails and displays the names (or inodes, if i2p is disabled) of the files that are open, along with the processes that opened them.

14. Destroy the contaminated volume(s).

```
vol destroy {volname|plexname} [-f]
```

15. Offline the contaminated aggregate.

```
aggr offline {aggrname|plexname} [-t cifsdelaytime]
```

Note: If the aggregate is embedded in a traditional volume that has CIFS shares, use the `-t` option to warn users before taking the aggregate (and hence the entire traditional volume) offline. The `cifsdelaytime` argument specifies the number of minutes to delay before taking the embedded aggregate offline, during which time CIFS users of the traditional volume are warned of the pending loss of service. A time of 0 means take the aggregate offline immediately with no warnings given. CIFS users can lose data if they are not given a chance to terminate applications gracefully.

16. Destroy the contaminated aggregate.

```
aggr destroy {aggrname|plexname} [-f]
```

Note: Before destroying the aggregate, traditional volume, or plex, the user is prompted to confirm the operation. The `-f` flag can be used to destroy an aggregate, traditional volume, or plex without prompting the user.

Note: Upon completion of the `aggr destroy` procedure, all disks are designated as spare disks.

Note: For contaminated storage environments that maintain SnapMirror or MetroCluster™ relationships with a secondary storage system, perform this procedure on the secondary storage system as well.

2.2 Clear Contaminated NetApp FAS (7-Mode) System by Using Single-Pass Overwrite

This procedure clears all data from SAS and SATA disks by performing a single-pass overwrite of zeros to each block of every disk assigned to the storage controller. The goal of this procedure is to provide data confidentiality and prevent recovery of cleared data by a robust keyboard attack. To clear all data from SAS and SATA disks, complete the following steps.

Note: Complete the procedure titled "Prepare Contaminated NetApp FAS (7-Mode) System for Single-Pass Overwrite or Disk Sanitization" before performing this procedure.

Note: NetApp recommends PuTTY (or a similar terminal emulator) for this procedure to perform error-free copy and paste (as opposed to keystroke) transcription of disk serial numbers to system log files in step 3.

1. Zero all disks previously contained in the contaminated aggregate.

```
disk zero spares
```

Note: The `disk zero spares` command zeroes out all nonzeroed RAID spare disks. The command runs in the background and may take several hours to complete, depending on the number of disks to be zeroed and the capacity of each disk. Spare disks that are in the process of zeroing are still eligible for use as creation, extension, or reconstruction disks. After the command is invoked, the `aggr status -s` command can be used to verify the status of the spare disk zeroing.

2. Once all disks previously contained in the contaminated aggregate are zeroed, display the configuration information.

```
sysconfig -a
```

3. Copy disk serial numbers from step 2 and paste them after the `logger` command.

Note: The `logger` command manually records and posts important changes to system configuration in the `/etc/messages` file. In this example, the `logger` command is used to record and log the serial numbers of all zeroed disks and other textual information important to data confidentiality and postincident handling of the data spill.

```
logger The following disk serial numbers refer to disks that were RAID group members of a
previously contaminated aggregate. All disks were cleared using a single pass overwrite of zeros:
6000c290c4dd03d337020938f3a330159; 6000c298bbc405823fde97e1b5102801;
6000c29144004cbf52af89f55917fec2; 6000c296729a82a69d7d60935b963fa5;
6000c29703f141733f44071a387550c2
```

2.3 Remove Data from Disks in NetApp FAS (7-Mode) System by Using Disk Sanitization

Disk sanitization is the process of clearing data by overwriting SAS and SATA disks with DoD-specified overwrite patterns. The goal of this procedure is to provide data confidentiality and prevent recovery of cleared data by a robust keyboard attack. To clear data from disks using disk sanitization, complete the following steps.

1. Assign ownership to unassigned disks (this applies to storage systems using software-based disk ownership).

```
disk show –n
disk assign all
```

2. Verify that the disks to be sanitized do not belong to a RAID group in any existing aggregate.

```
sysconfig –r
```

Note: Copy and save the list of spare disks for reference. Disks to be sanitized should be listed with `spare` status. If the expected disks are not displayed, they have not been assigned ownership. A disk must have ownership assigned to it before it can be sanitized.

3. Sanitize the specified disk or disks of all existing data.

```
disk sanitize start [-p pattern1|-r [-p pattern2|-r [-p pattern3|-r]]][-c cycle_count] disk_list
```

The identifier `-p pattern1 -p pattern2 -p pattern3` specifies a cycle of one to three user-defined hex-byte overwrite patterns that can be applied in succession to the disks that are being sanitized. The default pattern is three passes, using `0x55` for the first pass, `0xaa` for the second pass, and `0x3c` for the third pass.

Where:

- `-r` replaces a patterned overwrite with a random overwrite for any or all of the passes.
- `-c cycle_count` specifies the number of times the specified overwrite patterns will be applied. The default value is one cycle. The maximum value is seven cycles.

– `disk_list` specifies a space-separated list of the IDs of the spare disks to be sanitized.

Note: Do not turn off the storage system, disrupt storage connectivity, or remove target disks while sanitizing. If the sanitizing process is interrupted while target disks are being formatted, the disks must be reformatted before sanitizing can finish. If the sanitizing process must be aborted, use the `disk sanitize abort` command. If the specified disks are undergoing the disk formatting phase of sanitization, the abort will not occur until the disk formatting is complete. When the sanitizing process is interrupted, Data ONTAP displays a message that sanitization was stopped.

4. Verify the status of the disk-sanitizing process.

```
disk sanitize status [disk_list]
```

5. When disk sanitization is complete, release sanitized disks for reuse as spare disks.

```
disk sanitize release disk_list
```

Note: Do not interrupt disk sanitization. Rebooting the storage system or removing and reinserting a disk that has been sanitized causes that disk to be designated as broken.

Note: The specified disks are sanitized and designated as hot spares. The serial numbers of the sanitized disks are written to `/etc/sanitized_disks`.

In this example, the following command applies the default three disk sanitization overwrite patterns for one cycle (for a total of three overwrites) to the specified disks:

```
disk sanitize start 0c.00.3 0c.00.4 0c.00.5 0c.00.6 0c.00.7 0c.00.8 0c.00.9 0c.00.10 0c.00.11
```

In this example, the following command creates three disk sanitization overwrite patterns for six cycles (for a total of 18 overwrites) to the specified disks:

```
disk sanitize start -c 6 0c.00.3 0c.00.4 0c.00.5 0c.00.6 0c.00.7 0c.00.8 0c.00.9 0c.00.10
0c.00.11
```

2.4 Selectively Remove and Retain Data from Disks by Using Selective Disk Sanitization

Selective disk sanitization is used when a contaminated SAS or SATA storage system contains data that must be preserved to comply with statutory requirements or for postincident forensic analysis of a data spill.

Note: NetApp recommends using PuTTY (or a similar terminal emulator) for this procedure to perform error-free copy and paste (as opposed to keystroke) transcription of disk serial numbers to system log files in step 14.

To selectively remove and retain data from disks, complete the following steps.

1. Apply and enforce user access restrictions to the storage system volume or LUN consistent with IS confidentiality requirements and the organization's Cyber Incident Response Plan.
2. Stop any applications that write to the aggregate that is being sanitized.
3. From a Windows or UNIX client, delete from the active file system the directories or files containing data to be selectively sanitized. Use the appropriate Windows or UNIX command. For example:

```
rm /unixdir/unixfile.doc
```

4. Remove NFS and CIFS access to all volumes in the aggregate.
5. From the Data ONTAP CLI, delete all volume Snapshot copies of the FlexVol volumes that previously contained the deleted files and directories.

```
snap delete -V -a <<var_vol01>>
```

6. Note the names of the volumes containing data that should be preserved.

7. Display the free disk space, noting the total size and space used.

```
df -g <<var_vol01>>
```

8. If there is not enough free space to create an aggregate to contain the migrated volumes at their current size and the volumes have free space, decrease the size of each volume individually.

```
vol size <<var_vol01>> new_size
```

Note: The new size must be larger than the used space in the volume.

9. Create an aggregate to which the preserved data will be migrated.

```
aggr create <<var_aggr01>> disks
```

10. For each FlexVol volume containing data to be preserved, create a corresponding FlexVol volume in the new aggregate.

```
vol create dest_vol dest_aggrsize
```

11. For each FlexVol volume containing data to be preserved, copy the data to the new aggregate.

```
ndmpcopy /vol/src_vol /vol/dest_vol
```

Where:

- `src_vol` is the FlexVol volume in the aggregate to be sanitized.
- `dest_vol` is the newly created FlexVol volume that corresponds to the `src_vol volume`.

Note: Before running the `ndmpcopy` command, verify that the files or directories to be sanitized have been deleted from the source volume.

12. Copy clean Snapshot copies (those that do not refer to data blocks in the contaminated source file) to the target volume.
13. Display configuration information for the source aggregate.

```
sysconfig -a
```

14. Copy disk serial numbers from step 13 and paste after the `logger` command.

Note: The `logger` command manually records and posts important changes to system configuration in the `/etc/messages` file. In this example, `logger` is used to record and log the serial numbers of all zeroed disks and other textual information important to data confidentiality and postincident handling of the data spill.

```
logger The following disk serial numbers refer to disks that were RAID group members of a
previously contaminated aggregate. All disks were cleared using a single pass overwrite of zeros:
6000c290c4dd03d337020938f3a330159; 6000c298bbc405823fde97e1b5102801;
6000c29144004cbf52af89f55917fec2; 6000c296729a82a69d7d60935b963fa5;
6000c29703f141733f44071a387550c2
```

15. Offline and destroy each FlexVol volume in the aggregate that is being sanitized.

```
vol offline src_vol
vol destroy src_vol
```

16. Offline and destroy the source aggregate.

```
aggr offline src_aggr
aggr destroy src_aggr
```

17. Rename the new aggregate, giving it the name of the aggregate that was destroyed.

```
aggr rename <<var_aggr01>> old_src_aggr_name
```

18. Rename each FlexVol volume in the new aggregate to the name of the original FlexVol volume.

```
vol rename dest_vol old_src_vol_name
```

19. Reestablish CIFS or NFS services.

20. Perform the procedure in Remove Data from Disks in NetApp FAS (7-Mode) System Using Disk Sanitization on the contaminated aggregate.

2.5 Back Up Contaminated Snapshot Copies to Tape Media by Using Dump Command

Note: The `dump` command is used when a contaminated storage system contains data that must be retained and preserved on tape for compliance with statutory requirements or for postincident forensic analysis of a data spill. Copying contaminated Snapshot copies to tape is not a security best practice. It should be considered only when the need to retain and preserve information on tape exceeds the potential risks to IS confidentiality. When performing this procedure, enforce physical security and other compensating controls that mitigate risk and protect the confidentiality of the data on tape media.

To back up contaminated Snapshot copies to tape, complete the following steps.

1. Apply and enforce user access restrictions to the storage system volume or LUN consistent with IS confidentiality requirements and the organization's Cyber Incident Response Plan.
2. Stop all applications that write to the aggregate that is being sanitized.
3. From a Windows or a UNIX client, delete from the active file system the directories or files containing data to be selectively sanitized. Use the appropriate Windows or UNIX command, as shown in this example.

```
rm /unixdir/unixfile.doc
```

4. Remove NFS and CIFS access to all volumes in the aggregate.
5. From the Data ONTAP CLI, delete all volume Snapshot copies of the FlexVol volumes that previously contained the deleted files and directories.

```
snap delete -V -a <<var_vol01>>
```

6. Note the names of any volumes containing data that should be preserved.
7. Enter the following command for each volume that should be preserved, noting the total size and space used.

```
df -g <<var_vol01>>
```

8. Make a Level 0 dump of the entire file system of volume `vol0` to a remote tape device.

   Note: Each tape file in the dump should be less than 2GB.

```
dump options [ arguments ... ] tree
```

9. Secure backup tapes for the duration of the retention period.
10. Reestablish CIFS or NFS services.
11. Perform the procedures in the section titled "Remove Data from Disks in NetApp FAS (7-Mode) System by Using Disk Sanitization" on the contaminated aggregate.

2.6 FlexVol

Table 2) Use cases for FlexVol volumes.

Use Case	Procedure
Use this procedure to disable a flexible volume. This may be a precursor to destroying a volume that is no longer needed. Leaving a flexible volume offline for a period of time may be advisable as a final confirmation step that users do not require access to the volume.	Take FlexVol Volume Offline

Use Case	Procedure
Use this procedure to destroy a flexible volume when the volume is no longer needed. This step is functionally irreversible. Do not execute this procedure without thoroughly verifying that the FlexVol volume is no longer being used.	Destroy FlexVol Volume

Take FlexVol Volume Offline

To take a FlexVol volume offline, complete the following steps.

1. Take the FlexVol volume offline.

```
vol offline <<var_vol01>>
```

Note: Only one FlexVol volume can be taken offline at a time.

2. Check the volume status.

```
vol status <<var_vol01>>
```

Destroy FlexVol Volume

To destroy a FlexVol volume, complete the following steps.

1. Take the FlexVol volume offline.

```
vol offline <<var_vol01>>
```

2. Remove the FlexVol volume.

```
vol destroy <<var_vol01>>
```

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®

www.netapp.com