Redpaper

Alex Osuna
Roland Tretau

Introduction to IBM System Storage N series Data ONTAP 8.0

This IBM® Redpaper™ document provides an introduction to Data ONTAP® 8.0 7-mode. It describes the enhanced functionality in the latest release and gives step-by-step guidance for new installation and to upgrade from Data ONTAP 7G.

IBM System Storage® N series systems employ Data ONTAP, a highly scalable and flexible operating system for network storage appliances used in heterogeneous host environments. Data ONTAP can reduce storage management complexity in your enterprise, delivering flexible management and supporting high availability and business continuance.

A brief history of Data ONTAP

Figure 1 on page 2 shows the genealogy of Data ONTAP 8.0, which was developed by Network Appliance™ Inc. From 2004 until 2006, NetApp® continued to sell SpinServer® systems, which included the SpinFS® file system running in a Linux® environment IBM and the N series product line did not. Data ONTAP GX was created as a combination of the WAFL® file system and SpinFS running in a FreeBSD environment. Data ONTAP 8.0 brings Data ONTAP 7G and Data ONTAP GX together into a single storage operating system, also running on FreeBSD.

© Copyright IBM Corp. 2010. All rights reserved.

Having two products provided a way to meet the needs of the N series customers who were happy with Data ONTAP 7G, while allowing customers with certain application requirements to use Data ONTAP GX to achieve even higher levels of performance, and with the flexibility and transparency afforded by its scale-out architecture. Although the goal was always to merge the two products into one, the migration path for Data ONTAP 7G customers to get to clustered storage would eventually require a big leap.

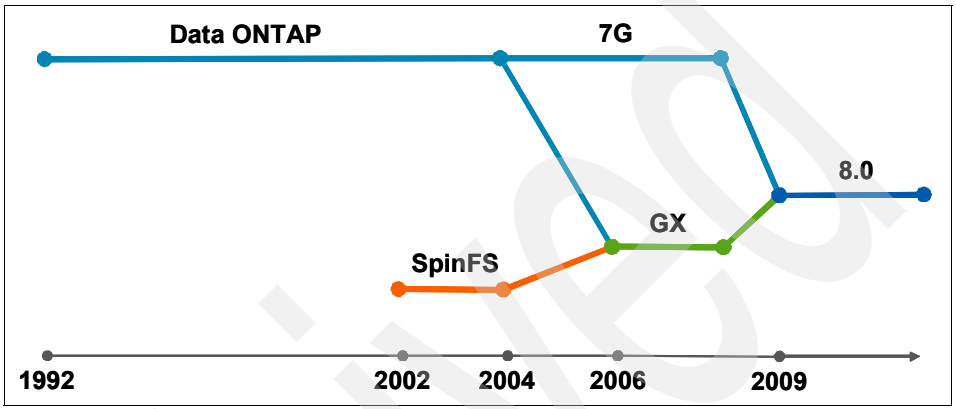

Figure 1 Data ONTAP history

Data ONTAP 8.0 allows Data ONTAP 7G customers to operate a Data ONTAP 8.0 “7-mode” system in the manner to which they are accustomed, while also providing a first step in the eventual move to a clustered environment. Data ONTAP 8.0 “cluster-mode” allows Data ONTAP GX customers to upgrade and continue to operate their clusters in the way they are already familiar with.

Components

Data ONTAP 8.0 is available in two image modes: 7-mode and cluster-mode. Customers must choose one mode when they purchase or install a storage system. Having the choice of two modes provides a simpler migration path for customers who want to use a legacy 7G system and those who do not have the business needs for clustering.

Customers who choose the 7-mode will have 7G-like functionality and a similar user interface. Customers opting for the cluster-mode will have GX-like functionality and a similar user interface. Only one mode can be in use at any given time.

Archived

If a customer decides to change from one mode to another, the change is a transition rather than an upgrade (or downgrade). Dual boot capabilities are not present, so the transition requires total reconfiguration of the storage system. This can include backup and restore of user data.

Typical approaches to install Data ONTAP 8.0 are:

- Upgrading from 7G to 7-mode
- Upgrading from Data ONTAP GX to cluster-mode
- New 8.0 installation - customer specifies the mode of operation

This IBM Redpaper focuses on Data ONTAP 8.0 7-mode.

7-mode stack

Although FreeBSD is familiar to Data ONTAP GX users, it is a departure from Data ONTAP 7G. The use of FreeBSD as the operating system for Data ONTAP 8.0 allows for some significant benefits. Besides the fact that Data ONTAP will continuously benefit from the third-party work taking place within the FreeBSD community, the clean separation of the operating system from the file system allows for focused innovation within the file system itself.

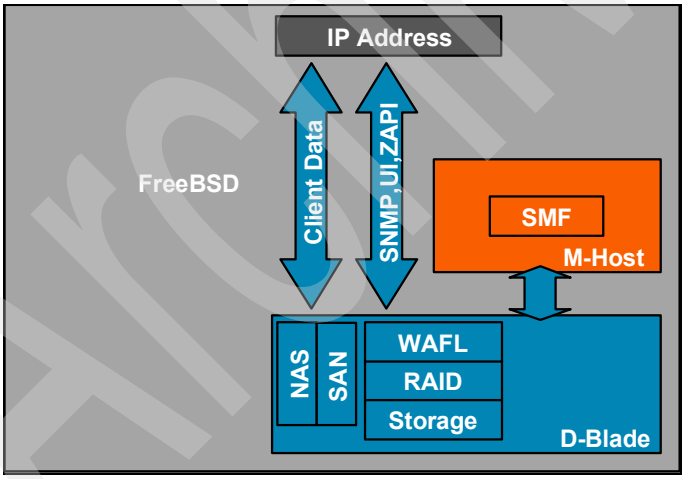

Figure 2 7-mode stack

Figure 2 shows the software stack that makes up Data ONTAP 8.0 7-mode. Although 7-mode runs in FreeBSD with a new data component called the D-blade and a new management component called the M-host, it acts very much like Data ONTAP 7G. The NAS and SAN protocols are handled by the D-blade.

Archived

In addition, there continues to be one logical interface for both client and administrative access to the node.

The D-blade is the data blade. It manages the storage attached to a node, and provides the WAFL (Write Anywhere File Layout) file system to map data containers and their associated metadata and attributes to disk blocks. In 7-mode, the D-blade services NAS and SAN protocol requests. It also provides the 7G compatible node user interface.

What if my system does not support Data ONTAP 8?

In case your system does not meet the hardware requirements for DOT8, you have to continue using Data ONTAP 7 software. The software will be maintained to provide the best possible interoperability with Data ONTAP 8 systems.

Changes and new features

This section discusses changes in terminology and functions as introduced by Data ONTAP 8.0 7-mode.

New terminology

This section discusses changes in terminology and new terms introduced with Data ONTAP 8.

Cluster and high availability terms

The following cluster and high availability terms have revised definitions:

- **Cluster:** In the Data ONTAP 7.1 release family and earlier releases, refers to an entirely different functionality: a pair of storage systems (sometimes called nodes) configured to serve data for each other if one of the two systems stops functioning.
- **HA (high availability):** In Data ONTAP 8.0, refers to the recovery capability provided by a pair of nodes (storage systems), called an HA pair, that are configured to serve data for each other if one of the two nodes stops functioning.

Archived

- **HA pair:** In Data ONTAP 8.0, refers to a pair of nodes (storage systems) configured to serve data for each other if one of the two nodes stops functioning. In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an *active/active* configuration.
- **CFO:** The term is now used for *controller failover* rather than cluster failover.

Data ONTAP terms

The following Data ONTAP terms have revised definitions:

- **Interface groups (ifgrps):** The naming conventions for 802.3ad link aggregation were not consistent. In Data ONTAP GX, 802.3ad link aggregation was called trunks, while in Data ONTAP 7G link aggregation was called vifs. There is now one name, *ifgrps* (for interface groups), for both 7-mode and cluster-mode. Interface groups are the grouping of several physical ports together to provide increased aggregate bandwidth and redundancy.
- **FreeBSD:** FreeBSD is now the Data ONTAP foundation.

Interface groups

Data ONTAP connects with networks through physical interfaces (or links). The most common interface is an Ethernet port, such as e0a, e0b, e0c, and e0d.

IEEE 802.3ad link aggregation is now supported by using interface groups. They can be single mode or multimode. In a single mode interface group, one interface is active while the other interface is on standby. In single mode, a failure signals the inactive interface to take over and maintain the connection with the switch.

In a multimode interface group all interfaces are active and share the same MAC address. There are two types of multimode operation:

- Static multimode interface group
- Dynamic multimode interface group

The **`ifgrp`** command refers to this setting as **`multi`**. Dynamic multimode interface groups can detect not only the loss of link status (as do static multimode interface groups), but also a loss of data flow. This feature makes dynamic multimode interface groups compatible with high-availability environments. The dynamic multimode interface group implementation in Data ONTAP is in compliance with IEEE 802.3ad (dynamic), also known as Link Aggregation Control Protocol (LACP). Dynamic multimode interface groups can detect the loss of link status, as well as a loss of data flow. However, a compatible switch must be used to

Archived

implement the dynamic multimode configuration. Example 1 shows options available to the **ifgrp** command.

Example 1 ifgrp command

```
TUCSON1> ifgrp
Usage:
        ifgrp create [single|multi|lacp] <ifgrp_name> -b [rr|mac|ip]
[<interface_list>]
        ifgrp add <ifgrp_name> <interface_list>
        ifgrp delete <ifgrp_name> <interface_name>
        ifgrp destroy <ifgrp_name>
        ifgrp {favor|nofavor} <interface>
        ifgrp status [<ifgrp_name>]
        ifgrp stat <ifgrp_name> [interval]
```

AutoSupport

AutoSupport has a new architecture in Data ONTAP 8 as shown in Figure 3. AutoSupport is now an M-Host (user space) process called *notifyd*. It collects information from the D-Blade, Management Gateway(mgwd), BSD commands, and files.

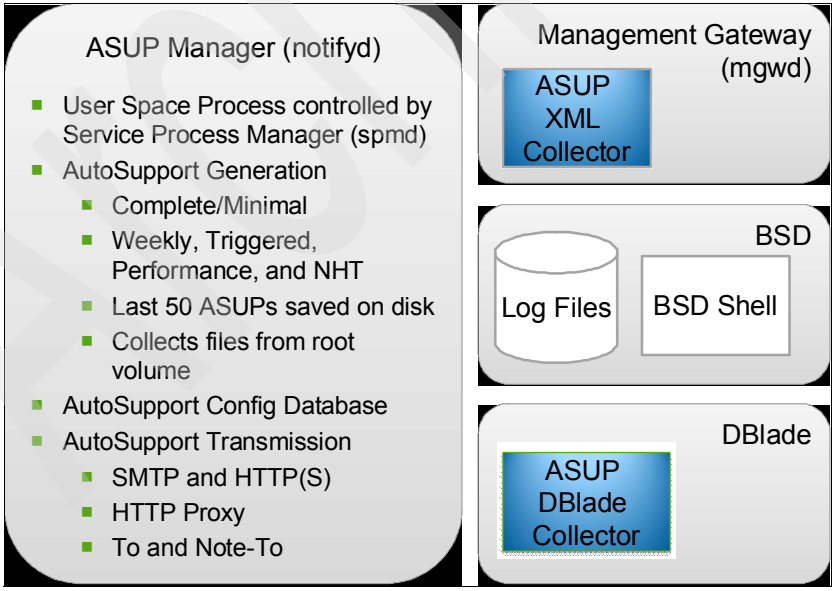

Figure 3 AutoSupport Architecture

Archived

Comparison

This section compares the old and new software releases to highlight major changes and differences.

7G and 7-mode

Because Data ONTAP 8.0 7-mode is very much like Data ONTAP 7G, the noticeable differences between them are few (see Table 1). The enhancements are the introduction of 64-bit aggregates, which allow for much larger aggregates, and new virtualization support for Hyper-V and Xen.

Table 1 Comparison of 7G and 7-Mode

Data ONTAP 7G	Data ONTAP 8.0 7-mode
32-bit aggregates	32-bit and 64-bit aggregates
Virtual interfaces (vifs)	Interface groups (ifgrps)
Active-active configurations	High-availability (HA) pairs
-	FreeBSD infrastructure
-	Hyper-V and Xen
IPSec	-
IPv6	-
SMB 2.0 (as introduced with 7.3.1)	-

Note that 7-mode has almost the exact same functionality as Data ONTAP 7G. Everything that is in Data ONTAP 7.3.0 and most of what is in 7.3.2 is available in Data ONTAP 8 7-mode.

Currently, Data ONTAP 8.0 7-mode does not provide support for:

- IPSec
- IPv6
- SMB 2.0
- SnapLock® software feature

New functionality of 7-mode

This section contains detailed information about the new functionality in Data ONTAP 8.0 7-mode.

Archive

Data ONTAP 8.0 7-mode is a storage solution that can help customers manage data in enterprise environments using a scalable and flexible storage operating system.

Data ONTAP 8.0 7-mode provides:

- Efficient storage
- High availability
- Business continuance
- Quality of service
- Reduced storage management complexity

SAN

The SAN functionality for block-based protocols FC and iSCSI contains the same functionality in Data ONTAP 8 7-mode as in Data ONTAP 7.3.

LUNs can be provisioned in aggregates larger than 16 TB by placing them in 64-bit aggregates. 64-bit aggregates are explained in detail in the next section.

Virtualization support has been improved by including native igroup and LUN support for Microsoft® Hyper-V and Citrix XenServer.

The `host ostype` parameter for the **`lun create`** command indicates the type of operating system running on the host that accesses the LUN. This parameter determines the following:

- Geometry used to access data on the LUN
- Minimum LUN sizes
- Layout of data for multiprotocol access

Note: Windows® Server 2008 is the only platform that runs Hyper-V. Windows Server 2008 hosts must have a LUN type of `windows_2008` or `hyper_v` to have aligned I/O.

The ostype indicates the type of host operating system used by all of the initiators in the igroup. All initiators in an igroup must be of the same ostype. The ostypes of initiators are solaris, windows, hpux, aix, netware, xen, hyper_v, vmware, and linux.

You must select an ostype for the igroup.

Archived

The **lun create** command has been enhanced by two new ostype options: **-hyper_v** and **-xen**. Example 2 displays the **lun** command options. The **-t** option can be used to specify the ostype.

Example 2 lun create command

```
lun create -s <size> -t <ostype> [ -o noreserve ] <lun_path>
lun create -f <file_path> -t <ostype> [ -o noreserve ] <lun_path>

Possible -t ostype options:
-aix
-hpux
-linux
-netware
-openvms
-solaris
-solaris_efi
-vld
-vmware
-windows
-windows_gpt
-windows_2008
-hyper_v
-xen
```

Like the **lun create** command, the **igroup create** command has been enhanced by two ostype options: **-hyper_v** and **-xen**. Example 3 displays the **igroup create** command options. The **-t** option can be used to specify the ostype of the initiators within the group. The type applies to all initiators within the group and governs the finer details of SCSI protocol interaction with these initiators. Valid arguments are listed in the example.

Example 3 igroup create command

```
igroup create { -f | -i } -t ostype [ -a portset ] initiator_group [
node ... ]

Possible -t ostype options:
-aix
-hpux
-linux
-netware
-openvms
-solaris
-vmware
-windows
```

Archived

```
-hyper_v
-xen
```

64-bit aggregates

Prior to Data ONTAP 8.0 7-mode, all aggregates were 32-bit, resulting in the 16 TB size limitation. As disk sizes increase, fewer spindles can be used; fewer spindles means lower performance (see Figure 4), hence system management would get more complex.

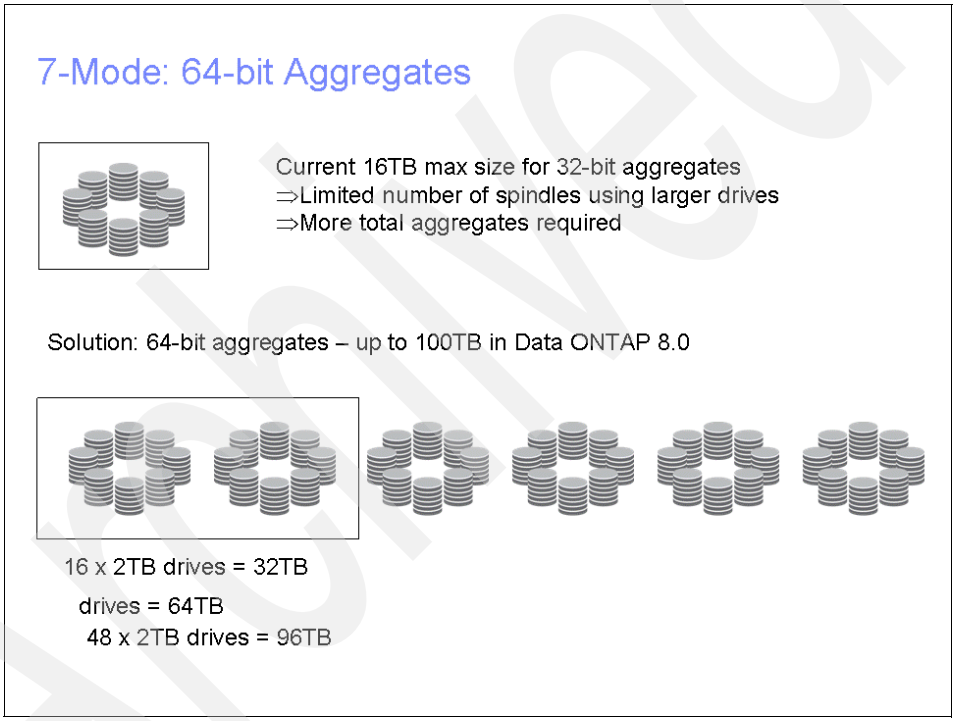

Figure 4 7 mode 64 bit agregates

The newly introduced 64-bit aggregate feature is able to overcome these limitations. No extra licenses are needed to enable and use the 64-bit aggregates feature. 64-bit aggregates are supported on normal N series and Gateway storage systems. They allow administrators to create aggregates up to 100 TB depending on the platform, as shown in Table 2.

The default aggregate type that is created in data Data ONTAP 8 7-mode is a 32-bit aggregate. Existing and newer 32-bit aggregates created in Data ONTAP 8

Archived

7-mode are still bound by the 16 TB size threshold, so you cannot add more disks to a 32-bit aggregate after the threshold is reached. A 64-bit aggregate provides the ability to scale future data growth.

Once created, 64-bit aggregates behave and can be used just like existing 32-bit aggregates. All the processes for creating and managing FlexVol® volumes inside the aggregate and accessing the volume are the same and use the same commands. All operations that can be performed on FlexVol volumes in 32-bit aggregates are also supported and work with FlexVol volumes in 64-bit aggregates. A 64-bit aggregate of any size can be created as long as it is less than the maximum size threshold stated in this document.

The newer 64-bit aggregates can coexist with new and existing 32-bit aggregates on the storage system. Therefore, you can create a 64-bit aggregate on a storage system that has existing 32-bit aggregates. This coexistence has no impact on the storage system in any way. Once the 64-bit aggregate is created, it is simply another aggregate on the storage system, with a higher size threshold, and can be used as a regular aggregate. 32-bit aggregates can also be created if necessary.

Table 2 64-bit aggregates size limits per system

Hardware platform	Maximum aggregate size
N7900	100 TB
N7700	70 TB
N6070	70 TB
N6060	50 TB
N6040	40 TB
N5600	50 TB
N5300	40 TB
N3040	30 TB

Important: System root volumes can only reside on 32-bit aggregates.

Traditionally, the address pointer is 32 bits in size; the maximum address that can be stored is (2^32 -1). Since each block that is referenced with this address is 4 KB in size, the maximum size of the aggregate that can be referenced is (2^32 -1) * 4 KB = 16 TB.

Archive

Data ONTAP 8.0 7-mode has increased the size of the address pointers to 64 bits, which makes possible a much bigger address space. This bigger address space in turn makes it possible to have aggregates that are much larger than the current 16 TB threshold.

64-bit aggregates have a bigger address format than 32-bit aggregates. Therefore, depending on the data set and the access pattern of the workload, the memory used for metadata when using 64-bit aggregates might be more than when using 32-bit aggregates. The exact amount of extra memory used in accessing data from 64-bit aggregates depends on the workload.

Interaction with other features

This section provides details on how 64-bit aggregates integrate with other Data ONTAP 8 7-mode functions.

Volume SnapMirror

The following information is valid for volume SnapMirror® in asynchronous mode, semi synchronous mode, and synchronous mode. Because volume SnapMirror works at the block level, the source and destination must both be FlexVol volumes in the same aggregate type.

Table 2 shows the volume SnapMirror combinations that are possible.

Table 3 Volume SnapMirror and 64-bit aggregates

	Destination volume in 32-bit aggregate	**Destination volume in 64-bit aggregate**
Source volume in 32-bit aggregate	Yes	No
Source volume in 64-bit aggregate	No	Yes

Qtree SnapMirror and SnapVault

Qtree SnapMirror and SnapVault® work at the logical file level and are not aware of the underlying physical storage structure or memory addressing scheme. Therefore, they can be used to transfer data between FlexVol volumes located in any kind of aggregates. You can establish a qtree SnapMirror relationship by using a qtree or a FlexVol volume in a 32-bit aggregate as the source and a qtree in a FlexVol volume in a 64-bit aggregate as the destination, and vice versa, as shown in Table 4.

Archived

Table 4 qtree SnapMirror and SnapVault and 64-bit aggregates

	Destination volume in 32-bit aggregate	Destination volume in 64-bit aggregate
Source volume in 32-bit aggregate	Yes	Yes
Source volume in 64-bit aggregate	Yes	Yes

NDMP and DUMP

Network Data Management Protocol (NDMP) and dump operation work at the logical file level and are not aware of the underlying physical storage structure or memory addressing scheme. Therefore, they work across different aggregate types.

You can use NDMP to copy data from a FlexVol volume in a 32-bit aggregate to a FlexVol volume in a 64-bit aggregate or a 32-bit aggregate.

You can create a dump of a volume in a 32-bit aggregate to create a backup copy and then restore from that backup to a volume in a 64-bit aggregate, and vice versa.

Aggr copy and vol copy

Aggr copy and vol copy operations work at the physical level. They work across similar types of aggregates and FlexVol volumes. Therefore, when using an aggr copy operation to copy data from one aggregate to another, both the source and destination aggregates must be of the same type. Similarly, for vol copy, both the source and destination must be FlexVol volumes contained in the same aggregate type.

Deduplication

Deduplication is supported on FlexVol volumes that reside in 64-bit aggregates. However, to enable deduplication on a FlexVol volume, the maximum FlexVol volume size must be less than or equal to the sizes shown in Table 5.

Table 5 Maximum FlexVol volume size for deduplication-enabled volumes

N series or Gateway platform	Deduplication maximum FlexVol volume size
N3400, N5300, N6040	4 TB
N5600, N6060, N6070, N7700, N7900	16 TB

Archived

MultiStore

MultiStore® is an optional software product that enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems in the network. Each storage system (sometimes referred to as a *filer*) created as a result of the partitioning is called a vFiler unit. A vFiler unit, using the resources assigned, delivers file services to its clients as a storage system does.

MultiStore works with 64-bit aggregates. However, if a vFiler unit has a FlexVol volume as its primary storage unit or as one of its storage units, you can perform disaster recovery and data migration operations only to a FlexVol volume that is located in the same aggregate type as its storage unit.

Other supported features

The following hardware and software features are supported:

- Controller failover (HA functionality is not affected by the aggregate type.)
- Performance Acceleration Module (PAM)
- FlexCache™
- SyncMirror®
- FlexShare™

Create 64-bit aggregates

You can use FilerView®, Systems Manager, or the command line interface (CLI) to create a 64-bit aggregate. Additional supported tools are SnapDrive® and Operations Manager. 64-bit aggregates of less than 16 TB can be created.

When using FilerView, the Aggregate Wizard shows the 64-bit option (Figure 5). Simply select the option to make the aggregate a 64-bit aggregate.

Archived

Figure 5 Aggregate wizard 64-bit option

Select **Manage Aggregate** and specify Filter by: Larger Aggregates to view details about the 64-bit aggregates (Figure 6 on page 15).

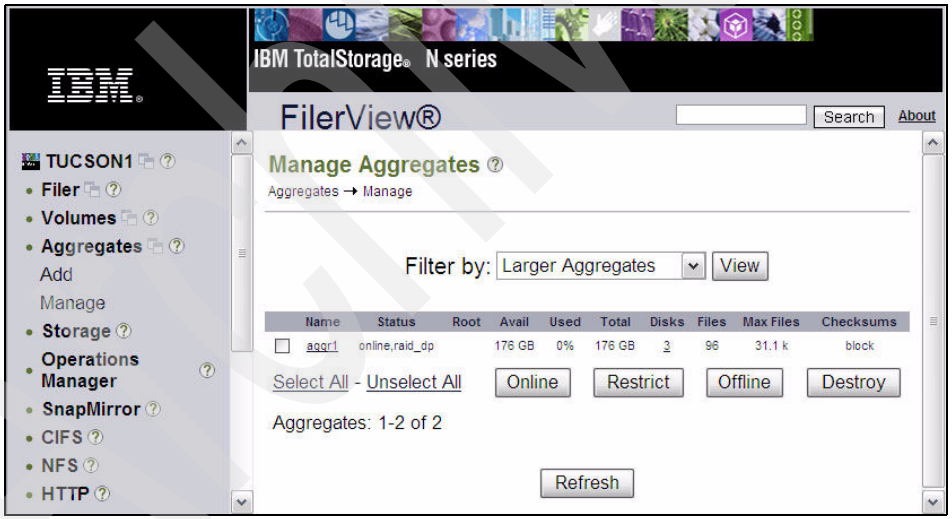

Figure 6 Manage aggregates

Select the aggregate of interest; in the resulting detail window the 64-bit check mark is set, indicating the 64-bit status (Figure 7).

Archived

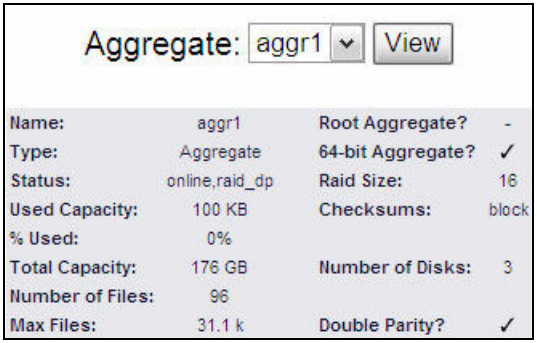

Figure 7 64-bit aggregate properties

The command to create a 64-bit aggregate is the same **aggr create** command that is used in Data ONTAP 7G for aggregate creation. The **aggr create** command has a new flag, **-B,** to specify the type of aggregate to create. It takes a value of either 32 or 64 and creates the respective type of aggregate. The **aggr create** command with **-B 64** creates a 64-bit aggregate; without the **-B** flag by default it creates a 32-bit aggregate.

Example 4 shows how to use the CLI to create an 64-bit aggregate.

Example 4 Create 64-bit aggregate

```
TUCSON1> aggr create aggr1 -B 64 3
TUCSON1> aggr status
          Aggr State          Status            Options
          aggr1 online        raid_dp, aggr
                              64-bit
          aggr0 online        raid_dp, aggr     root
                              32-bit
```

Migration to 64-bit aggregates

Currently, there is no direct migration path or conversion from 32-bit to 64-bit aggregates. The following options can be used to migrate the data:

- qtree SnapMirror (QSM)
- SnapVault
- ndmpcopy (snapshot copies cannot be migrated using this method).

If a FlexVol volume is in a 32-bit aggregate, but you need it to be larger than 16 TB to hold the data in the volume, the data has to be moved to a FlexVol volume located in a 64-bit aggregate. The following tools can be used to migrate

Archived

data from a FlexVol volume in a 32-bit aggregate to a FlexVol volume in a 64-bit aggregate:

- **Qtree SnapMirror or SnapVault:** Qtree SnapMirror can be used to copy data from a qtree or a FlexVol volume in a 32-bit aggregate to a qtree located in a FlexVol volume in a 64-bit aggregate. If the data on the source volume is laid out in qtrees, you need to have one qtree SnapMirror relationship per qtree from the source to the destination to move all the data from the source to the destination.
- **NDMP:** The **`ndmpcopy`** command can be used to copy volumes, qtrees, directories, or individual files from the source volume located in a 32-bit aggregate to the destination volume located in the 64-bit aggregate. If you have a single large directory or file that you want to copy from a volume in a 32-bit aggregate to a volume in a 64-bit aggregate, you can use this tool.

Both of these tools can be used to copy data that is present only in the active file system of the source volume. However, Snapshot™ copies that are present in the source volume cannot be copied to the destination using either of these tools.

The following options are *not* supported because they are block-based mirroring techniques:

- Volume SnapMirror (VSM)
- aggr copy
- vol copy

Access options

The default values of some options have changed for Data ONTAP 8. This section explains details about these options.

Secure console

Data ONTAP 8 uses SSH and SSL options as defaults. They allow administrators to configure a secure environment. The new default values shown in Example 5 differ from previous Data ONTAP 7.3.x versions.

Example 5 Default SSH and SSL values

```
options ssh.enable  on
options ssh2.enable  on
options ssh1.enable  off
options ssh.passwd_auth.enable  on
options ssh.pubkey_auth.enable  on
```

Archive

```
options httpd.admin.ssl.enable  on
```

Note: Upgraded systems will retain the security settings present prior to the upgrade.

Administration access

Starting with Data ONTAP 8, non secure protocols are disabled by default, as shown in Example 6. In addition, a root password is required for new installations.

Example 6 Default administration access options

```
options rsh.enable  off
options telnet.enable  off
options httpd.enable  off
options ftpd.enable  off
options httpd.admin.enable  off
```

Note: Upgraded systems will retain the security settings present prior to the upgrade.

The default requirements can be changed using the security password rules options shown in Example 7.

Example 7 Password security rules

```
TUCSON1> options security.passwd.rules
security.passwd.rules.enable on
security.passwd.rules.everyone off
security.passwd.rules.history 0
security.passwd.rules.maximum 256
security.passwd.rules.minimum 8
security.passwd.rules.minimum.alphabetic 2
security.passwd.rules.minimum.digit 1
security.passwd.rules.minimum.symbol 0
```

Named snapshot feature for SnapVault

The *named snapshot* feature for SnapVault allows you to back up data using SnapVault from any arbitrary snapshot at the disaster recovery site.

Archived

Previously, in Data ONTAP 7.3.1 and earlier releases, SnapVault could not back up data from a specified snapshot residing on the volume SnapMirror destination volume. SnapVault would only transfer data from the latest Volume SnapMirror (VSM) created snapshot.

In Data ONTAP 8 7-mode SnapVault can back up data from any arbitrary snapshot (either a user-specified or a scheduled snapshot) from the volume SnapMirror destination. SnapVault backup from the disaster recovery site continues to be the same as the SnapVault backup from primary storage system to secondary, with the following restrictions:

- For a SnapVault scheduled update from a DR site, administrators need to set up the SnapVault primary schedule at the volume SnapMirror source.
- In case of a SnapVault update from a named snapshot (that is, **snapvault update -s <*snapname*>**), the administrator needs to make sure that the named snapshot exists at the VSM source. To prevent any snapshot getting deleted by Data ONTAP applications, you can use the new command **snapvault snap preserve**.

This feature exists only in 7-mode.

The **snapvault snap preserve** command enables you to preserve the required Snapshot copy. This command prevents Data ONTAP features (such as **snap autodelete**) from deleting the snapshot copy.

Example 8 shows the use of the command to preserve a Snapshot copy at the volume SnapMirror primary system.

Example 8 SnapVault snap preserve

```
TUCSON1> snapvault snap preserve
usage: snapvault snap preserve <vol> <snapname> [<tagname>]

TUCSON1>snapvault snap preserve vol1 snap1 cactus
TUCSON1>snap list vol1
Volume vol1
working...
%/used %/total date name
------ ------- --- -----
 29%(29%)0%( 0%) Feb 22 09:07 snap1(acs)
```

If the Snapshot copy that you have preserved should no longer be retained, you can unpreserve it.

Example 9 shows the preserved snapshot and how to delete a preserved snapshot copy.

If the **-all** option is specified, all preservations on a specified Snapshot are removed.

Example 9 SnapVault snap unpreserve

```
TUCSON1>snapvault snap preservations vol1
snapid snapname
---------- ----------
1 snap1
system> snapvault snap preservations primary
snap1 type tagname
---------- ------------
cli cactus
TUCSON1> snapvault snap unpreserve
usage: snapvault snap unpreserve <vol <snapname> {[<tagname>] | [-all]}
TUCSON1>snapvault snap unpreserve vol1 snap1 cactus
```

Boot menu

The boot menu changed in Data ONTAP 8 7-mode to provide increased options for administrators.

Example 10 shows the version of the bootloader we used in the lab and that the options at the firmware prompt have not been changed.

Example 10 Boot at firmware prompt

```
Boot Loader version 1.7
Copyright (C) 2000-2003 Broadcom Corporation.
Portions Copyright (C) 2002-2009 NetApp

CPU Type: Dual Core AMD Opteron(tm) Processor 265

Starting AUTOBOOT press Ctrl-C to abort...
Loading x86_64/freebsd/image1/kernel:....0x100000/3312752
0x529000/3201008 0x8367f0/564584
Autoboot of PRIMARY image aborted by user.

LOADER> ?
Invalid command: "?"
Available commands: version, update_flash, netboot, boot_diags,
boot_backup, boot_primary, boot_ontap, flash, bye, set, lsmod,
autoboot, go, boot, load, ndp, ping, arp, ifconfig, show, savenv,
saveenv, unsetenv, set-defaults, setenv, printenv, help
```

Archived

```
*** command status = -1
LOADER>
```

The new boot menu shown in Example 11 has been enhanced to provide administrators with the following capabilities:

Choice 4: Initialize all the disks and create a FlexVol root volume.

Choice 5: Go into maintenance mode, where some aggregate and disk operations can be performed.

Choice 6: Update the CompactFlash card from a backup configuration.

Choice 7: Install new software on a Gateway system.

Choice 8: Reboot the storage system.

Example 11 Enhanced boot menu

```
Loading x86_64/freebsd/image1/kernel:....
Starting program at 0x80144c00
IBM Data ONTAP 8.0RC3 7-Mode
ichpwr module loaded
vsensor: <vsensor>
Copyright (C) 1992-2009 NetApp.
All rights reserved.
*******************************
*                             *
* Press Ctrl-C for Boot Menu. *
*                             *
*******************************
^CBoot Menu will be available.
Restoring /var from /cfcard/x86_64/freebsd/varfs.tgz

Please choose one of the following:

(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
Selection (1-8)?
```

Archived

Upgrade to DOT8 7-mode

This section provides information about how to upgrade to Data ONTAP 8 7-mode. First, general steps are outlined and second, a practical step-by-step sample shows how to perform the upgrade task.

Pre-upgrade evaluation

Before upgrading to DOT8 7-mode you should carefully inspect your system, including installed hardware and software. Upgrade all software to the most current release.

Figure 8 shows possible migrations paths from earlier versions to Data ONTAP 8. Only migrations from 7.3.x to Data ONTAP 8 7-mode provide the possibility for a non disruptive upgrade (NDU). This upgrade path is the only one that can be reverted without data loss.

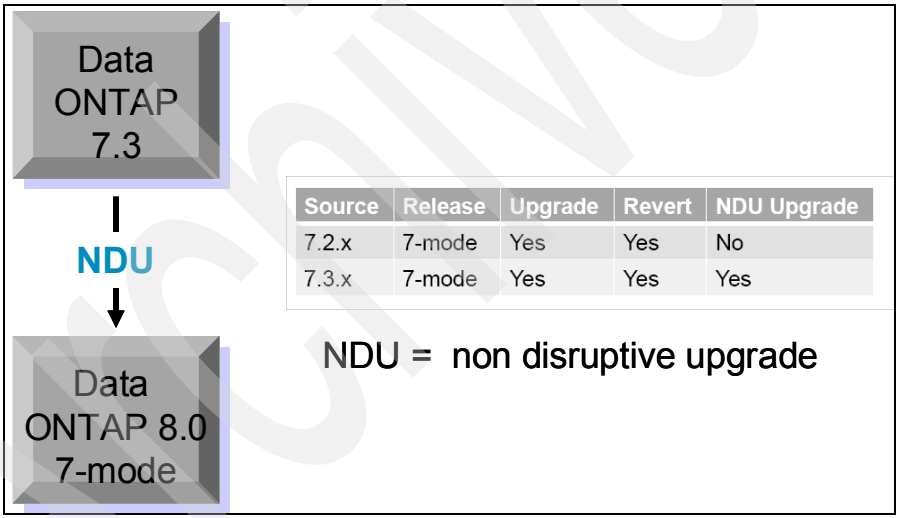

Figure 8 Migration paths towards Data ONTAP 8

Upgrade steps

At a high level the upgrade steps are as follows:

1. Review your current system hardware and licenses.
2. Review all necessary documentation.
3. Generate an AutoSupport e-mail.

Archived

4. Obtain the Data ONTAP upgrade image.
5. Install the software and download the new version to the CompactFlash card.
6. Reboot the system.
7. Verify the install.

In addition to the steps listed previously, you should perform the following steps *before* attempting to install a non disruptive upgrade:

- Validate the high-availability controller configuration.
- Remove all failed disks to allow giveback operations to succeed.
- Upgrade disk and shelf firmware.
- Verify that system loads are within the acceptable range; load should be less than 50% on each system.

Table 6 shows supported NDU upgrade paths.

Table 6 Supported high-availability configuration upgrade paths

Source	Release	Upgrade	Revert	NDU
7.2.x	7-mode	Yes	Yes	No
7.3.x	7-mode	Yes	Yes	Yes

Evaluate free space for LUNs

Before upgrading a storage system in a SAN environment, ensure that every volume containing LUNs has available at least 1 MB of free space. The space is needed to accommodate changes in the on-disk data structures used by the new version of Data ONTAP.

System requirements

Generally, Data ONTAP 8 requires you to use 64-bit hardware. Older 32-bit hardware is not supported. Currently, the supported systems and hardware are:

- N series: N7900, N7700, N6070, N6060, N6040, N5600, N5300, N3040
- Performance Acceleration Cards (PAM)

Archived

Revert considerations

Reverting from Data ONTAP 8 to an earlier version is a disruptive process. Furthermore, the following considerations and limitations apply to the revert process:

- User data will be temporarily offline and unavailable during the revert.
- Administrators should plan when the data will be offline to limit the unavailability window and make it fall within the time out window for the SAN attach kits on hosts.
- Data ONTAP 8.0 7-mode features must be disabled before reverting.
- 64-bit aggregates and 64-bit volumes cannot be reverted, thus the data must be migrated.
- You cannot revert while an upgrade is in progress.
- The **revert_to** command will remind administrators of the features that need to be disabled to complete the reversion.
- FlexVol volumes must be online during the reversion.
- Space guarantees should be checked after the reversion.
- You must delete any Snapshot copies made on Data ONTAP 8.0.
- You must reinitialize all SnapVault relationships after the revert because all snapshots associated with Data ONTAP 8.0 were deleted.
- SnapMirror sources must be reverted before SnapMirror destinations are reverted.

Example 12 shows details of the **revert_to** command.

Example 12 revert_to command

```
TUCSON1> revert_to
usage: revert_to [-f] 7.2 (for 7.2 and 7.2.x)
       revert_to [-f] 7.3 (for 7.3 and 7.3.x)

    -f   Attempt to force revert.
TUCSON1>
```

You cannot revert while the upgrade is still in progress. Example 13 shows how to use the *privileged* mode to check for upgrade processes that are still running.

Example 13 WAFL scan status

```
TUCSON1> priv set advanced
Warning: These advanced commands are potentially dangerous; use
```

```
         them only when directed to do so by IBM
         personnel.
TUCSON1*> wafl scan status
Volume vol0:
 Scan id                    Type of scan     progress
       1    active bitmap rearrangement     fbn 454 of 1494 w/
max_chain_len 7
...
```

Example 14 shows how the revert process was performed. The following steps summarize the process:

1. All 64-bit aggregates were removed.
2. All snapshots were deleted for all volumes and aggregates (the command in Example 14).
3. Snapshot schedules were disabled.
4. SnapMirror was disabled.
5. The **software upgrade** command was issued.
6. The **revert_to** command was issued.
7. The system rebooted to the firmware level prompt.

You are now able to perform a netboot or use the **autoboot** command.

Example 14 Revert to Data ONTAP 7.3

```
TUCSON1> snapmirror off
...
TUCSON1> snap delete -A -a aggr0
...
TUCSON1> software list
727_setup_q.exe
732_setup_q.exe
8.0RC3_q_image.zip
TUCSON1> software update 732_setup_q.exe
...
TUCSON1> revert_to 7.3
...
autoboot
...
TUCSON1> version
Data ONTAP Release 7.3.2: Thu Oct 15 04:39:55 PDT 2009 (IBM)
TUCSON1>
```

Archived

Netboot process for clean install

You can use the netboot option for a fresh install of the storage system, booting from a Data ONTAP version stored on a remote HTTP or TFTP (Trivial File Transfer Protocol) server.

Use the following steps for a netboot install:

1. Ensure that the hardware is fine, and includes a 1GB CompactFlash card, an RLM (Remote LAN Module) card, and a network interface card.
2. Upgrade the BIOS if necessary.

   ```
   ifconfig e0c -addr=10.10.123.??? -mask=255.255.255.0 -gw=10.10.123.1
   ping 10.10.123.45
   flash tftp://10.10.123.45/folder.(system_type).flash
   ```

3. Enter one of the following commands at the boot environment prompt:
   - If you are configuring DHCP, enter:

     ```
     ifconfig e0a -auto
     ```

   - If you are configuring manual connections, enter:

     ```
     ifconfig e0a -addr=filer_addr -mask=netmask -gw=gateway
     -dns=dns_addr -domain=dns_domain
     ```

     Specify the pertinent values:

`filer_addr`	Address of the storage system.
`netmask`	Network mask of the storage system.
`gateway`	Gateway for the storage system.
`dns_addr`	IP address of a name server on your network.
`dns_domain`	Domain Name Service (DNS) domain name. If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL; you need only the server's host name.

4. Set up the Boot environment.

   ```
   set-defaults
   setenv ONTAP_NG true
   setenv ntap.rlm.gdb 1
   setenv ntap.init.usebootp false
   setenv ntap.mgwd.autoconf.disable true
   ```

   Depending on N6xxx or N7xxx, set it to `e0c` for now. You can set it back to `e1a` later.

   ```
   setenv ntap.bsdportname e0f
   setenv ntap.bsdportname e0c
   ```

Archived

```
"a New variable for BR nay be needed."
setenv ntap.givebsdmgmtport true #before installing build
setenv ntap.givebsdmgmtport false #after installing build
"FOR 10-MODE"
setenv ntap.init.boot_clustered true
ifconfig e0c -addr=10.10.123.??? -mask=255.255.255.0 -gw=10.10.123.1
ping 10.10.123.45
```

5. Netboot from the loader prompt.

   ```
   netboot http://10.10.123.45/home/bootimage/kernel
   ```

6. Enter the NFS root path. The NFS root path is the IP address of an NFS server followed by the export path.

   ```
   10.10.123.45/vol/home/web/bootimage/rootfs.img
   ```

7. Press Ctrl-C for Boot Menu.
8. Select **Software Install** (option 7). Enter the URL to install the image.

   ```
   http://10.10.123.45/bootimage/image.tgz
   ```

Note: The provided URLs are examples *only* and should be replaced with the URLs for your environment.

A sample upgrade

Our test environment was composed of two N5600 systems, each with a designated EXN4000 shelf. This section describes an upgrade we performed from Data ONTAP 7.3.2. In case a clean install is required, Data ONTAP 8 7-mode also supports the netboot process.

We used the following steps to perform the upgrade:

1. We reviewed the current system configuration using the **`sysconfig -a`** command. The output is shown in Figure 9.

Archived

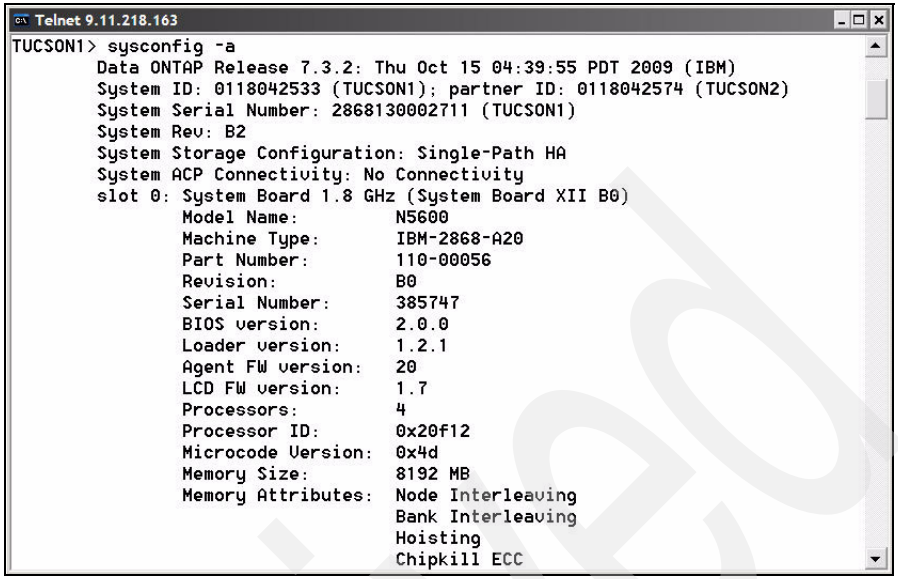

Figure 9 Sysconfig command

2. We verified the existing firmware level using the **`version -b`** command (Figure 10).

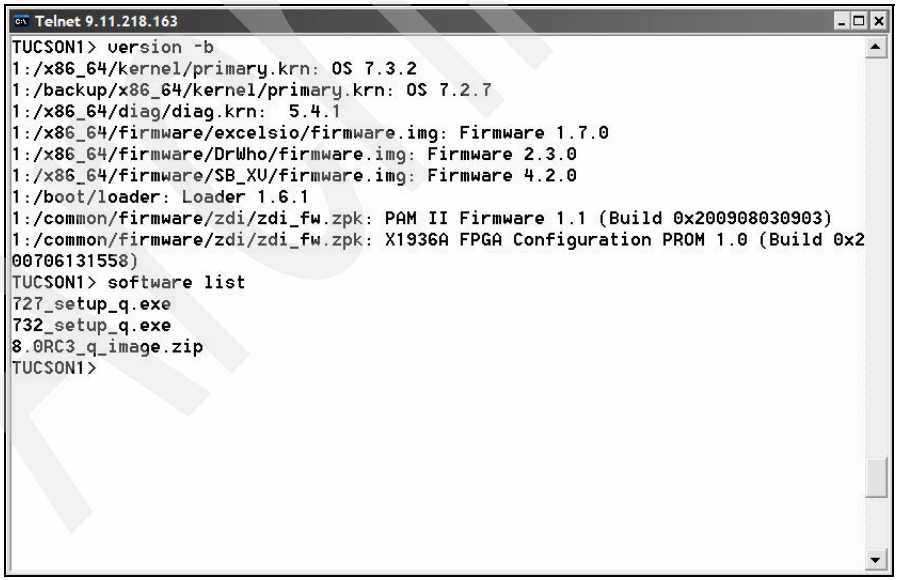

Figure 10 Version command

Archived

You could also use the **`license`** command to verify what software is licensed on the system (Confidentiality issues prevent us from including an illustration of this step).

3. Directory /etc/software was created to host installable ONTAP releases (Figure 11). The install images were copied from a Windows client using the administrative share \\filer_ip\c$.

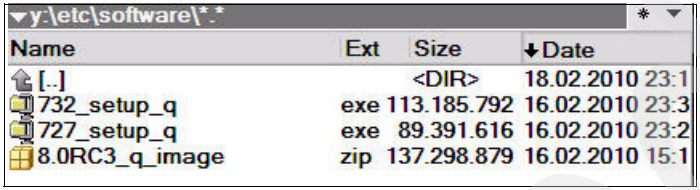

Figure 11 /etc/software directory

4. The **`software list`** command was used to display possible images (Figure 10.

   Starting with Data ONTAP 8, software images end with *.zip* and are no longer *.exe* or *.tar* files. The **`software`** command must be used to install or upgrade Data ONTAP 8 versions.

5. We used the **`software update`** command. At the time of this writing only Data ONTAP 8 7-mode Release Candidate 3 was available, so all tasks were performed using this software version. The update process is displayed in Figure 12 and takes much longer compared to previous Data ONTAP versions, so be patient.

Archive

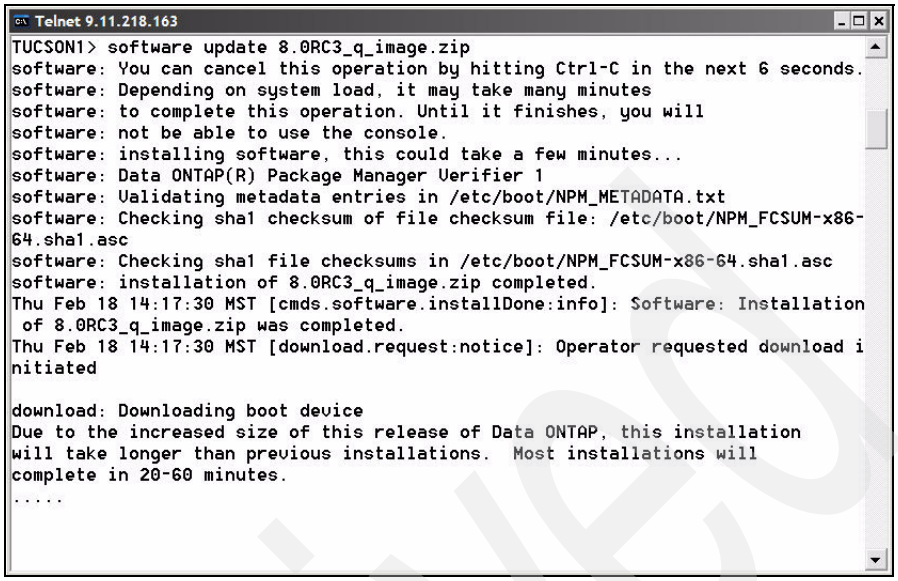

Figure 12 software update

6. When the system rebooted, we pressed **CTRL-C** to access the first bootmenu.

Important: The bootloader must be upgraded, otherwise Data ONTAP 8 will not load. Instead the previously installed version will continue to boot.

7. The bootloader of the system was upgraded using the **update_flash** command (Figure 13).

Important: Be sure that all firmware is up to date! However, if you are experiencing long boot times, you can disable the auto update of disk firmware prior to downloading Data ONTAP using this command:
options raid.background_disk_fw_update.enable off

Archive

```
Boot Loader version 1.2.1
Copyright (C) 2000,2001,2002,2003 Broadcom Corporation.
Portions Copyright (C) 2002-2005 Network Appliance Inc.

CPU Type: Dual Core AMD Opteron(tm) Processor 265
> setenv USE_SECONDARY false


Starting AUTOBOOT press Ctrl-C to abort...

Autoboot of PRIMARY image aborted by user.

LOADER> update_flash
New BIOS Version: 2.4.0
New Loader Version: 1.7
Saving Primary Image to Secondary
Programming .+.+.+.+.+.+.+.+.+.+.+.+.+.+.+.+.+.+.+.+.+.+.+.+.+.+.+.+.+.+.+.+.+ don
e. 2097152 bytes written
Updating Primary Boot Flash
Programming .+.+.+.+.+.+.+.+.+.+.+.+.+.+ done. 917504 bytes written
LOADER>
```

Figure 13 Flash update

8. Use the **autoboot** command and perform another **reboot** if Data ONTAP 8 did not load immediately after the flash update.
9. Once the boot process has completed, verify that the correct version is in place using the **version** and **sysconfig** commands (Figure 14).

```
Telnet 9.11.218.163

Data ONTAP (TUCSON1.)
login: root
Password:
Thu Feb 18 15:19:10 MST [telnet_0:info]: root logged in from host: 9.11.144.117
TUCSON1> version
Data ONTAP Release 8.0RC3 7-Mode
TUCSON1> sysconfig
        Data ONTAP Release 8.0RC3 7-Mode
        System ID: 0118042533 (TUCSON1); partner ID: <unknown> ()
        System Serial Number: 2868130002711 (TUCSON1)
        System Rev: B2
        System Storage Configuration: Single-Path HA
        System ACP Connectivity: NA
        slot 0: System Board
                Processors:         4
                Processor type:     Opteron
                Memory Size:        8192 MB
                Memory Attributes:  Node Interleaving
                                    Bank Interleaving
                                    Hoisting
                                    Chipkill ECC
        Remote LAN Module           Status: Online
        slot 0: BGE 10/100/1000 Ethernet Controller
                e0a MAC Address:    00:a0:98:06:cc:53 (auto-1000t-fd-up)
```

Figure 14 Software version verification

Archived

The team who wrote this IBM Redpaper

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization, San Jose Center.

Alex Osuna is a Project Leader at the International Technical Support Organization, Tucson Center. He writes extensively on all areas of storage. Alex has 32 years of experience in the I/T industry focused mainly on storage and working in areas of maintenance, planning, technical support, Early Ship Programs, Sales and publishing. Alex holds certifications from IBM, Microsoft, RedHat, and the Open Group.

Roland Tretau is an Information Systems professional with over 15 years of experience in the IT industry. He holds master's degrees in engineering and business, and is an author of many storage-related IBM redbooks. Roland has a background in project management, consulting, operating systems, storage solutions, enterprise search technologies, and data management.

Thanks to the following people for their contributions to this project:

Thorsten Busch, Irene Springer, Thomas Grüter, Guido Boessow, Falk Schneider
IBM Germany

Sandro De Santis
IBM Italy

Rucel Javier
IBM Philippines

Kyle Burrell, Mark Taylor
IBM US

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author - all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Archive

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

Archived

Archived

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:
This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

© Copyright International Business Machines Corporation 2010. All rights reserved.
Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

This document REDP-4671-00 was created or updated on July 8, 2010.

Send us your comments in one of the following ways:
- Use the online **Contact us** review Redbooks form found at:
  ibm.com/redbooks
- Send your comments in an email to:
  redbook@us.ibm.com
- Mail your comments to:
  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099, 2455 South Road
  Poughkeepsie, NY 12601-5400 U.S.A.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

IBM®
Redbooks (logo) ®
Redpaper™
System Storage®

The following terms are trademarks of other companies:

Snapshot, Network Appliance, FlexShare, FlexCache, WAFL, SyncMirror, SpinServer, SpinFS, SnapVault, SnapMirror, SnapLock, SnapDrive, MultiStore, FlexVol, FilerView, Data ONTAP, NetApp, and the NetApp logo are trademarks or registered trademarks of NetApp, Inc. in the U.S. and other countries.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Archived