ORACLE
FLASH STORAGE SYSTEMS

Implementing Boot From SAN on Oracle Flash Storage Systems

ORACLE WHITE PAPER | NOVEMBER 2015

ORACLE

Table of Contents

Introduction 3

Advantages of Booting From an FC SAN 4

- Decrease Administrative Overhead 4
- Facilitate Disaster and Server Failure Recovery 4
- Provide Significant Performance Improvements Over the Traditional Hard Drive 4

Overview of FC SAN Boot Strategy 6

FC SAN Boot Configuration 7

- Requirements 7
- Configure the FS System for FC SAN Boot 8
- Advantages of an Independent Boot LUN 9
- Set Up a Separate LUN for Applications 9
- Advantages of a Separate LUN for Applications 10
- Create an FC Boot LUN on the FS System 12

Configuring the Host Server for FC SAN Boot 15

- Set Boot Precedence in the System BIOS 15
- Configuring the HBA for FC Boot 16
  - Configure the QLogic HBA 16
  - Configure the Emulex HBA 17

Installing the Operating System on the Server 17

- Considerations for Installing Solaris 17
- Considerations for Installing Oracle Linux 18
- Considerations for Installing SUSE Linux Enterprise 11 or 12 18

Considerations for Installing Microsoft Windows 2008, Windows 2012 18

Best Practices for Configuring the FS System 18

Best Practices for Using Clones and LUN Copies 19

Managing a Dynamic Boot From SAN Environment 19

Best Practices for Managing a Boot From SAN Environment 19

Rebuild the Server Using the Boot LUN 20

Reconfigure the System in a New Environment 20

Use Case: Managing an Environment With a Boot LUN, Clones, and Copies 21

Conclusion 23

References 23

Introduction

Business organizations continually search for ways to simplify their IT management infrastructure, enhance scalability, and reduce costs while increasing reliability and performance in their data center. Traditionally, operating systems have been installed on local drives on individual servers or on direct attached storage (DAS). This approach presents various challenges to IT organizations:

- Dedicated internal boot devices cannot be shared with other servers and are therefore often underutilized.
- IT staff must perform management tasks on these systems locally rather than from a central management system, leading to increased administrative costs.
- For optimal redundancy and performance, additional RAID software or host bus adapters (HBAs) are required to manage these storage devices.

In addition, local drives on individual servers present particular challenges for multisite administration and disaster recovery site maintenance. Administrator tasks like the following can involve complex operations requiring specialized software:

- Creating clones of drive content to offsite hosts
- Replicating the operating systems and the data on a server to a disaster recovery site

Booting from a storage area network (SAN) provides solutions to these problems while offering additional advantages such as cost savings, enhanced security, ease of management, increased flexibility, and reduced downtime.

The difficult challenge of managing the servers for an entire enterprise can be simplified when administrators centrally manage all storage-related tasks, such as operating system (OS) maintenance, at the storage array level rather than at the individual server level. Server boot devices can be located on an Oracle Flash Storage (FS) System that is accessed by servers across a high-availability Fibre Channel (FC) SAN. This strategy enables increased efficiency, and even automation, of many administrative tasks, significantly reducing operating expenses.

This white paper describes how to implement an FC boot solution on an FS System with a high-availability SAN. The focus of this white paper is on booting from a SAN using Oracle Linux kernels such as the Red Hat Compatible Kernel (RHCK), the Unbreakable Enterprise Kernel (UEK), and similar kernels. And because Solaris and Windows operating systems support boot from SAN (BFS) solutions, the solutions described in this white paper can be adapted for those environments too.

The details that are involved in selecting, configuring, and troubleshooting the components used to build an FC SAN boot solution are beyond the scope of this white paper. For information about commands, implementation, specific features, and limitations, refer to the product documentation for the operating systems, servers, storage arrays, and FC switches in your environment. See "References" at the end of this white paper for a list of helpful external resources.

The BFS solutions described here are based on the FC protocol. Similar advantages are available for the Internet SCSI (iSCSI) and Fibre Channel Over Ethernet (FCoE) protocols.

NOTE: Oracle FS 1-2 Flash Storage Systems support multiple Storage Classes, while Oracle All Flash FS Storage Systems support only flash storage capability. In this white paper, the text and illustrations that describe multiple Storage Classes apply only to Oracle FS1-2 Flash Storage Systems.

Advantages of Booting From an FC SAN

The strategy of booting from an FC SAN provides some significant advantages.

Decrease Administrative Overhead

All OS storage is provisioned and managed from the FS System. A server can be easily replaced by remapping its corresponding boot LUN to a new server. If the new server has the same profile as the server being replaced, the new server will boot the OS from the SAN without additional reconfiguration. Snapshots and clones of OS images can be created and mapped to new servers on the SAN with just a few clicks of a mouse, simplifying migration and scalability tasks.

Facilitate Disaster and Server Failure Recovery

By installing operating systems on the FS System rather than on individual servers, administrators can take advantage of the data protection and redundancy features to help reduce downtime during maintenance and fault outages. OS images can be protected by using snapshots and clones, or can be backed up by using Oracle MaxRep for SAN or Network Data Management Protocol (NDMP).

Provide Significant Performance Improvements Over the Traditional Hard Drive

Booting from a LUN that is configured to use solid state drives (SSDs) provides a markedly superior performance over booting from a traditional hard drive. The data shown in Figure 1 illustrates how implementing a strategy using boot from SAN (BFS) can deliver a significant increase in I/O performance. An SSD disk can read and write three times more data than a standard disk for the same test dataset.

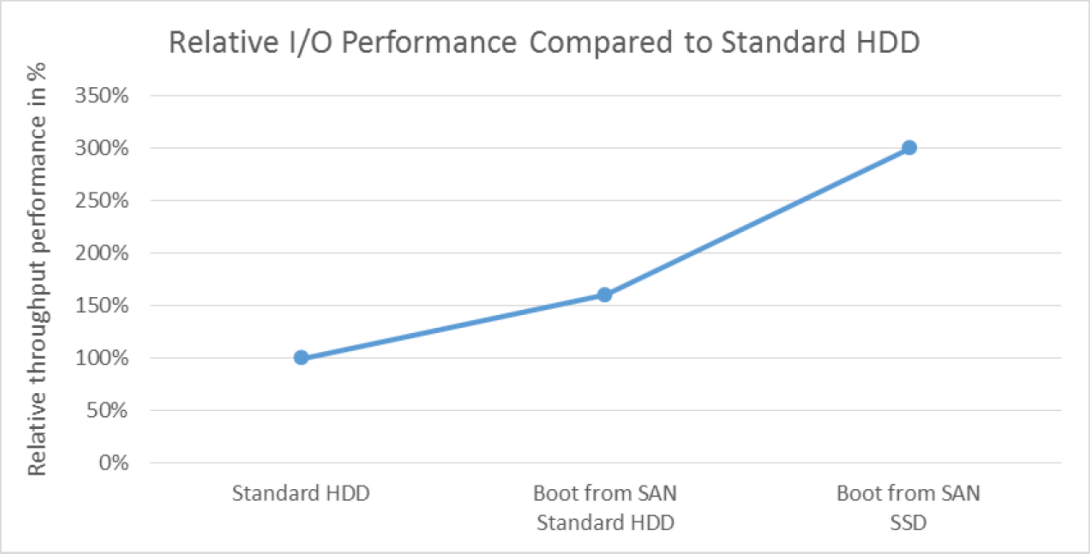

Figure 1. Comparison of relative I/O performance of various boot configurations

The test results in the graph above were compiled by running the Vdbench utility on the boot LUN for 2 hours[1]. The test involved writing 512 MB and reading 512 MB in cycles with 8 K block size for 2 hours on multiple threads. The Y axis represents the relative throughput of 50% random read and random write operations in units of data reads and writes per hour, compared to booting from a standard HDD. The graph shows the ratio of the total data transferred divided by the total data transferred by standard spindle disk. In the test results, the I/O on an SSD drive is faster than the I/O on a spindle disk by a factor of 3.

In addition to the advantages that are listed above, implementing a BFS solution offers other benefits.

Table 1. More Benefits of the BFS Strategy

Benefit	Description
Less space, reduced facility costs	Implement diskless servers and blade servers that use space more efficiently, consume less power, and require less cooling than other booting strategies.
Simplified upgrades	Just swap the boot disk to change to a new OS or to upgrade to a new version.
Standardize on OS and application images	Install consistent versions of OSs and applications across an organization.
Rapid deployment	Store and deploy numerous OS configurations, instances, variations, and proofs of concepts at will, without the bother of re-creating or reinstalling.
Enhanced security	Implement multiple encryption technologies for greater levels of system security.
Increased flexibility for IT departments	Try out patches and updates, and easily revert to previous configurations.
Virtualization features	Leverage features that are typically obtained only by implementing virtualization in the data center. Avoid additional overhead and cost by integrating BFS with cloning.
Scalability with high availability	Scale a data center effectively and quickly while providing high availability of the systems.
No need for RAID devices on servers	Because the FS System is RAID-protected shared storage array, there is no need to install RAID hardware or software on each server.

1 The results shown in Figure 1 are based on a specialized test setup and might not reflect typical results.

Overview of FC SAN Boot Strategy

Figure 2 shows an FC BFS implementation in an Oracle FS 1-2 Storage System.

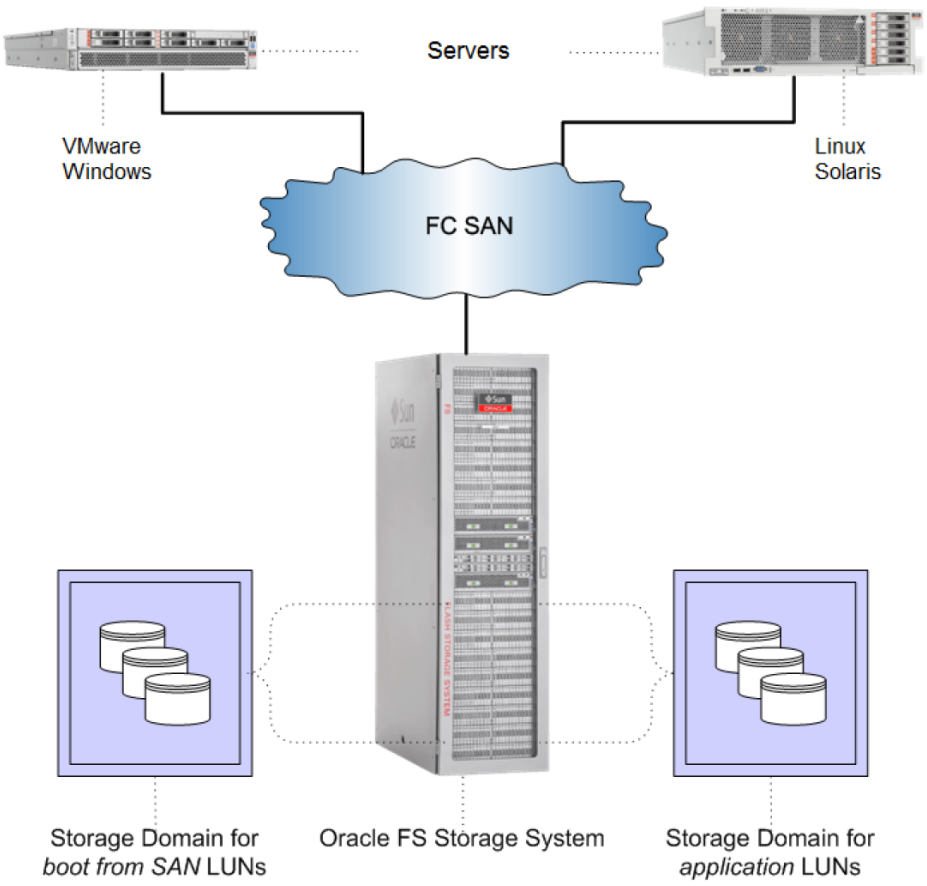

Figure 2. FC boot solution using an FS System

In this scenario, servers are booted from a dedicated Storage Domain of centrally managed volumes.

A Storage Domain is a virtual storage pool consisting of an assortment of drive groups, each containing drives of a particular Storage Class and of a particular capacity. Administrators can use Storage Domains to store data in logical groups to fulfill storage requirements.

Within the Storage Domain that contains BFS LUNs, each volume serves as boot LUN for a specific server. Figure 2 shows two types of servers used in the validation testing (x86 and Oracle Linux) with the operating systems that are validated on each server type.

When the array is also used for data storage, a dedicated Storage Domain and separate data paths should be used for boot devices. See "Best Practices for Configuring the Oracle Flash Storage" for more information.

Any server and HBA combination that supports BFS can be used to implement an FC boot solution using an FS System.

FC SAN Boot Configuration

Requirements

Configuring an FC BFS solution using an FS System requires the following:

- Zoning must be configured in the local SAN to enable the FC ports on the server to communicate with the FC target ports on the FS System.
- In the FS System, at least one FC boot LUN must be created and zoned with at least one port enabled. For details, see "Configure the Oracle FS System for FC SAN Boot."
- An HBA that supports FC BFS must be installed on each server that is provisioned from the SAN. The solution described in this white paper was tested with the following FC HBA driver and firmware versions:
  - QLogic QLE2562 (Firmware Version 4.03.02, BIOS Revision 2.02)
  - QLogic QLE8362 (Firmware Version 4.03.02, BIOS Revision 2.02)
  - Emulex LPe12002 (BIOS Version 2.11a2)

  The installation requirements for this configuration are described in *Sun Storage 16 Gb Fibre Channel PCIe Universal Host Bus Adapter, QLogic Installation Guide For HBA Model 7101674*.

NOTE: For more information about supported drivers and firmware, refer to *Oracle Flash Storage System Interoperability Guide* which is available from the SAN Storage – Oracle Flash Storage Systems documentation library.

- A system that is set up according to the following considerations:
  - The FC HBA on each server must be configured as the primary boot device.
  - A storage target LUN in the FS System must be provisioned with the appropriate operating system.
  - The LUN must be mapped to the server's initiator port.

  For more information, see "Configuring the Host Server for FC SAN Boot."

NOTE: For instructions about using the latest version of the Oracle FS System Manager (GUI), refer to *Oracle Flash Storage System Administrator's Guide* in the Oracle Flash Storage System Documentation Library on SAN Storage – Oracle Flash Storage Systems.

Configure the FS System for FC SAN Boot

Configuring the FS System to support FC SAN boot involves setting up an independent boot LUN. In addition, a Storage Domain and a volume group should be created for the new LUN.

The example in Figure 3 shows the dialog that the administrator can use to set up properties for a new SAN LUN that will become the boot LUN.

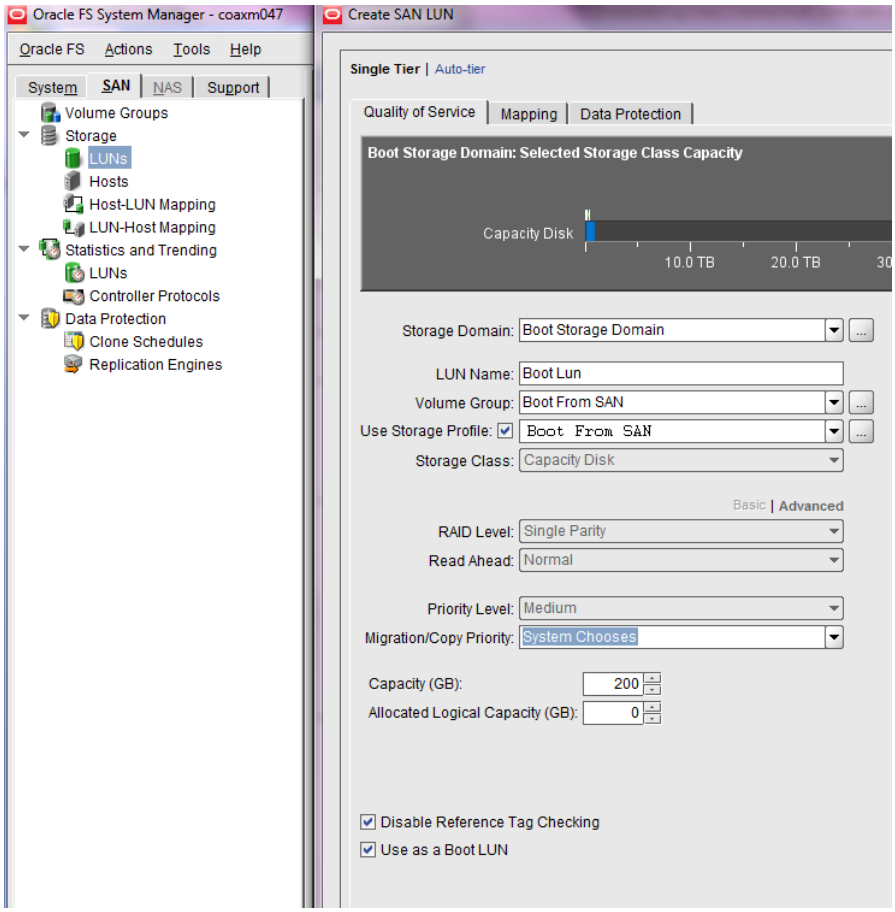

Figure 3. Configuring parameters for a boot LUN using Oracle FS System Manager (GUI)

For this example, the administrator used the following recommended settings for the new LUN:

- The LUN is created as a single-tier LUN.
- The Storage Domain is **Boot Storage Domain**. This is one of the Storage Domains that is supplied by the FS System software.
- The volume group is **Boot From SAN**. A *volume group* is an administrative system object that an administrator creates to group logical volumes into organizational units.
- The Storage Profile is **Boot From SAN**. FS System provides a collection of Storage Profiles that are optimized for specific uses, as shown in Table 2. The Boot From SAN profile is recommended for a boot LUN.
- The **Allocated Logical Capacity** is set to **0** to indicate that no blocks will be allocated to the LUN until they are written to the first time.

- The option to **Disable Reference Tag Checking** is selected to inhibit the SAN host from detecting a block of data that has not previously been written. Disabling reference tag checking helps prevent a slow-down in system performance or an inability to boot. (If this option is not selected, the system might trigger an event that advises the administrator to disable reference tag checking functionality for the boot LUN.)
- The option to **Use as a Boot LUN** is selected to indicate that the LUN will be used as a BFS LUN.

The example shown in Figure 3 is taken from version 6.2 of Oracle FS System Manager. The dialog might differ for other versions. For step-by-step procedures for creating LUNs, Storage Domains, and volume groups, refer to *Oracle Flash Storage System Administrator's Guide* in the Oracle Flash Storage System Documentation Library on SAN Storage – Oracle Flash Storage Systems.

NOTE: When the FS System is used for data storage, best practices recommend that a dedicated Storage Domain and separate data paths be used for boot devices. See "Best Practices for Configuring the Oracle Flash Storage" for more details.

Advantages of an Independent Boot LUN

The strategy of creating an independent boot LUN, and therefore keeping applications and data on separate LUNs, provides significant benefits:

- Requires little disk space, therefore the boot LUN can be allocated from high performance SSD disks
- Is an OS-only drive, so it can be migrated to other servers
- Can be cloned to replicate an organizational IT standard OS image, saving OS install and patch time while maximizing utility for storage
- Offers higher performance than traditional directly attached storage
- Provides the ability to apply policy-based replication uniformly to the IT environment, resulting in high availability

Set Up a Separate LUN for Applications

In addition to a BFS LUN, another helpful strategy is to set up additional LUNs, each using the Storage Profile that is appropriate for their particular applications.

A Storage Profile is a set of Quality of Service (QoS) attributes that can be used to configure a logical volume. Oracle provides a collection of preconfigured Storage Profiles that are optimized for specifics uses. Administrators can also create custom profiles to define specific properties for logical volumes.

FS System release 6.3 provides several predefined Storage Profiles which are listed in Table 2. Administrators running earlier releases of FS System can create similar Storage Profiles or additional customized profiles as needed. (For instructions on creating and managing Storage Profiles, refer to *Oracle Flash Storage System Administrator's Guide* in the Oracle Flash Storage System Documentation Library on SAN Storage – Oracle Flash Storage Systems.) The FS System supports high-capacity hard disk drives (HDDs), high-performance HDDs, and SSDs.

Table 2. Recommended Storage Profiles for BFS, application, and swap LUNs

Application	Storage Profile
BFS LUN	SSD premium boot from SAN SSD standard boot from SAN Standard HDD boot from SAN
Application LUN	General Purpose – Transaction Streaming Media (General - Read) Streaming Media (High Resolution - Read) Streaming Media (Render/Edit)
Swap LUN	Swap

The process involved in creating a LUN for an application involves the following tasks:

1. Scanning for the LUN in the OS
2. Ensuring that the multipath software is installed on the host and that the OS recognizes the FS LUN as an asymmetric logical unit access (ALUA) object
3. Formatting, labeling, and partitioning the new application LUN
4. Creating an appropriate OS file system on the LUN, for example, ext4 on Linux
5. Mounting the new partition on the OS boot LUN and ensuring that the mount points will be persistent across reboots

Once these tasks are complete, the LUN is ready for the application to be installed.

NOTE: Oracle FS Path Manager (FSPM) should be installed after OS installation. FSPM software and documentation are available for download from the Oracle Flash Storage System Downloads site.

Administrators might want to set up different Storage Domains for different applications. For more information about setting up Storage Domains for Oracle ASM, refer to "Best Practices for Optimizing Storage for Oracle Automatic Storage Management with Oracle FS1 Series Storage."

NOTE: When the FS System is used for data storage, best practices recommend that a dedicated Storage Domain and separate data paths be used for boot devices. See "Best Practices for Configuring the Oracle Flash Storage" for more details.

Advantages of a Separate LUN for Applications

Creating a separate LUN for applications ensures that the LUN offers the following functionality:

- Is replication ready and matches the performance requirements of the application.
- Is provisioned with storage capacity (or type) appropriate for the data that the LUN will be storing. For example, faster, more efficient storage for high I/O, frequently-accessed data.
- Provides rapid application recovery, in case of any problem, predictably and reliably.
- Includes a well-defined recovery point and recovery time for any application, without additional cost, by using the advanced features of the FS System.
- Brings the benefits typically available only with expensive virtualization technology, without incurring the price and performance overhead of virtualization. This refers to the portability of operating systems and applications across any servers.
- Supports thin provisioning capabilities, thereby optimizing storage usage and reducing costs.

» Scales up the disk size on demand when data or application requirements change, using the LUN expansion feature of the FS System.

» Reaps all the benefits of an elastic scalable Cloud infrastructure to allow for flexible deployment of IT resources on-demand.

» Allows the administrator to allocate swap space from lower-cost media and to prevent the swap data from being cloned.

In the example shown in Figure 4, LUNs are configured separately to maximize their performance and value on an FS System.

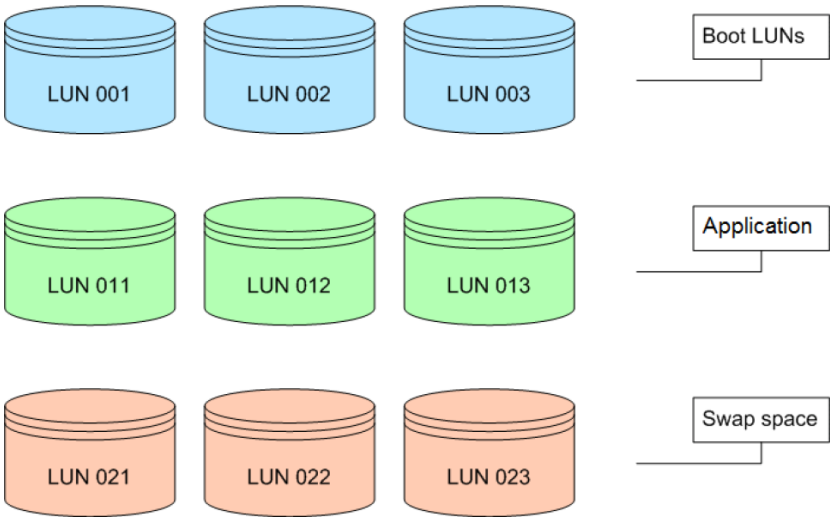

Figure 4. Allocating separate LUNs for boot, applications, and swap with different Storage Profiles

In this design, the administrator has configured three types of LUNs:

» Three LUNs are configured as BFS LUNs (LUNs 001, 002, and 003). The space allocated for the OS on these boot LUNs is minimal (approximately 10–60 GB).

» Three additional LUNs are configured to run applications such as Oracle Real Application Clusters (RAC) and Oracle Fusion. Each LUN has a different size and profile allocation to address the application requirements.

» The final three LUNs are allocated for swap drives based on the swap size requirements of the applications. The sizes and storage profiles of these swap LUNs are different from those used by the boot and the application LUNs.

Create an FC Boot LUN on the FS System

Using the Oracle FS System Manager (GUI), a system administrator can create and provision an independent boot LUN. The example screen in Figure 5 shows values for the Storage Domain, the volume group, the capacity of the drive, and so on.

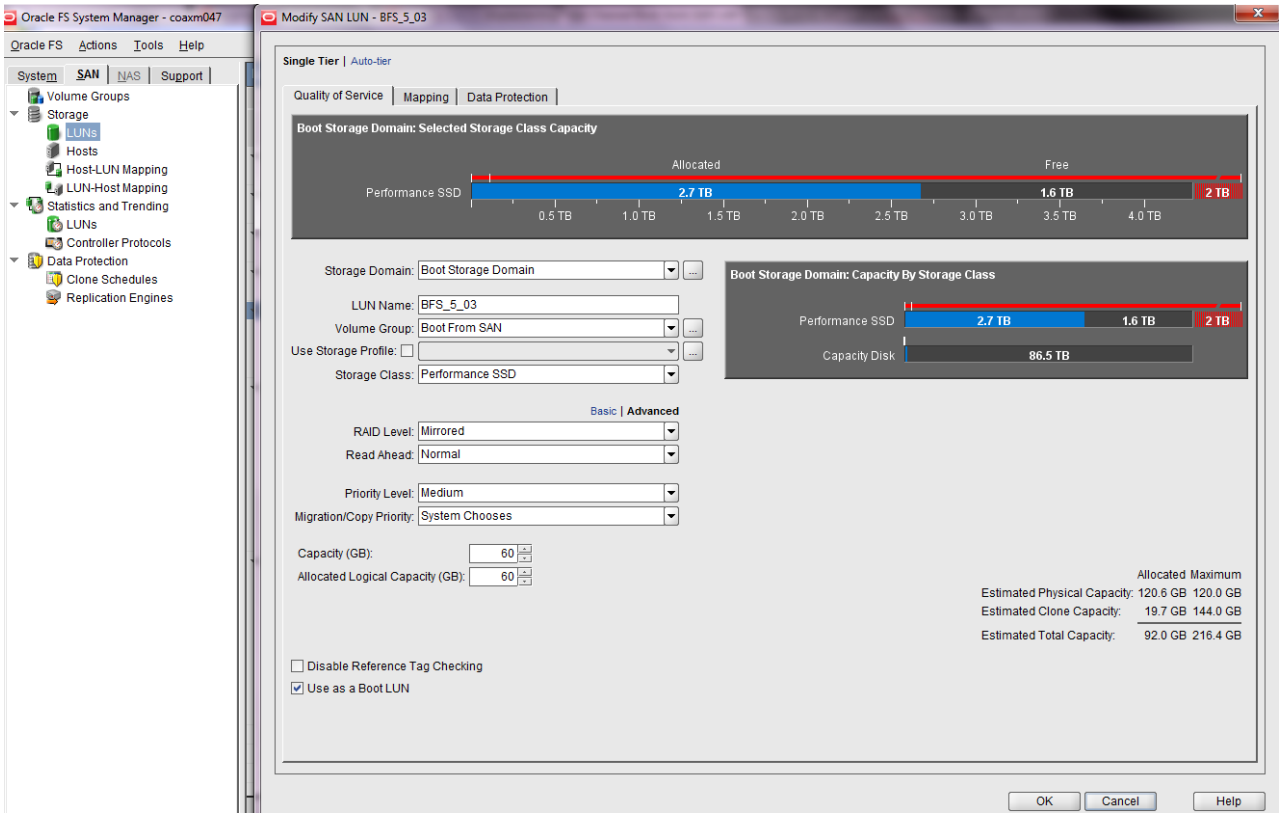

Figure 5: Creating a boot LUN using the Oracle FS System Manager (GUI)

Based on the criteria shown in Table 2, the administrator selects the appropriate storage profile of the drive from the Create SAN LUN screen shown in Figure 6.

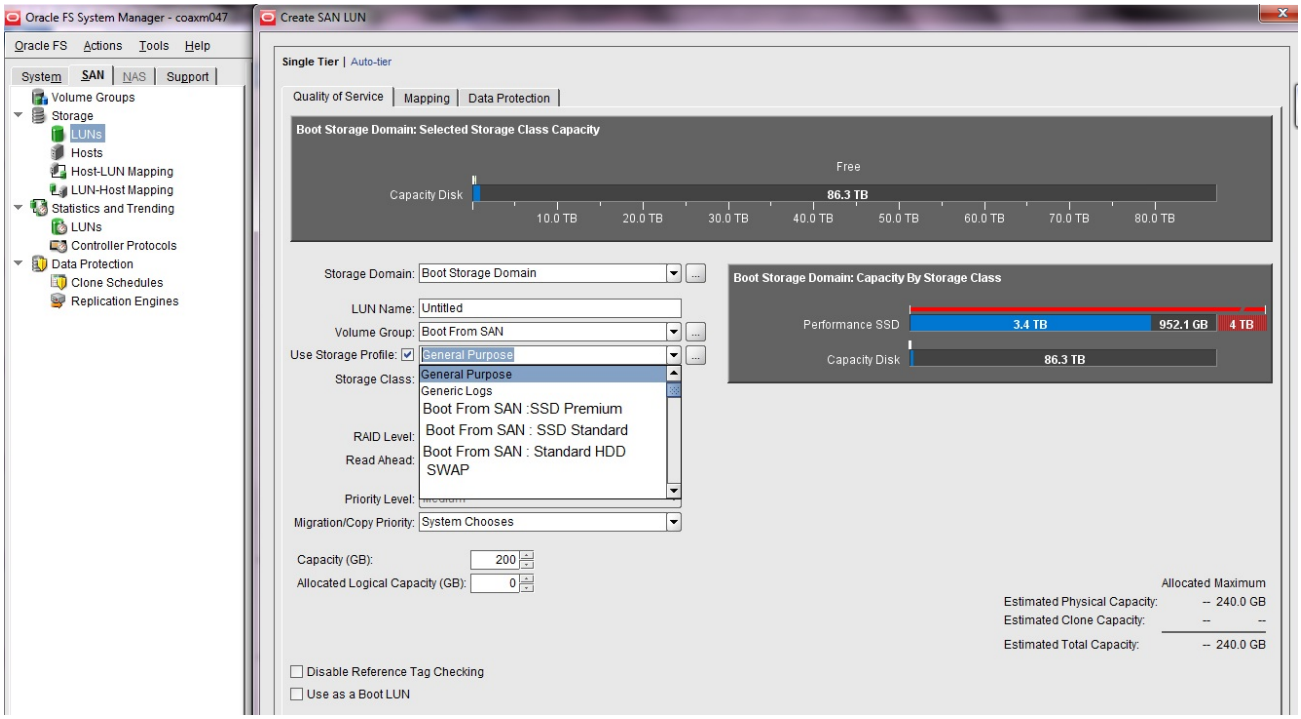

Figure 6. Selecting the Storage Profile for a BFS LUN, application LUN, or swap LUN

The initial capacity of the boot LUN can be set up based on an estimate of the size requirements, and then can be adjusted later if needed. The FS System supports thin provisioning, which allows the administrator to increase the capacity of the LUN and then make adjustments to the operating system on each server to enable access to the expanded LUN capacity.

The administrator selects **Use as a Boot LUN** in the lower left of the Create SAN LUN screen, as shown in Figure 7.

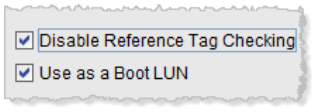

Figure 7. Selecting the options to Disable Reference Tag Checking and Use as a Boot LUN

In addition, the administrator might select **Disable Reference Tag Checking**.

NOTE: The Disable Reference Tag Checking option is required for Oracle Solaris 9, 10, 11, 12, and similar operating systems. This option is *not* required for Oracle Linux and other Linux-based operating systems.

Disabling reference tag checking on a LUN can improve performance by disabling some of the internal data integrity checks that are normally performed.

Figure 8 shows an example of how to map the independent boot LUN and separate LUNs to the appropriate servers. By selecting **Mapping**, choosing **Fibre Channel** as the Access Protocol, and then clicking **Create**, the administrator can map the LUNs to a host. If the host exists, the LUNs can be mapped to that host name. If the host does not exist, the World Wide Name (WWN) of the FC HBA of the host server can be located in the list as shown in Figure 11.

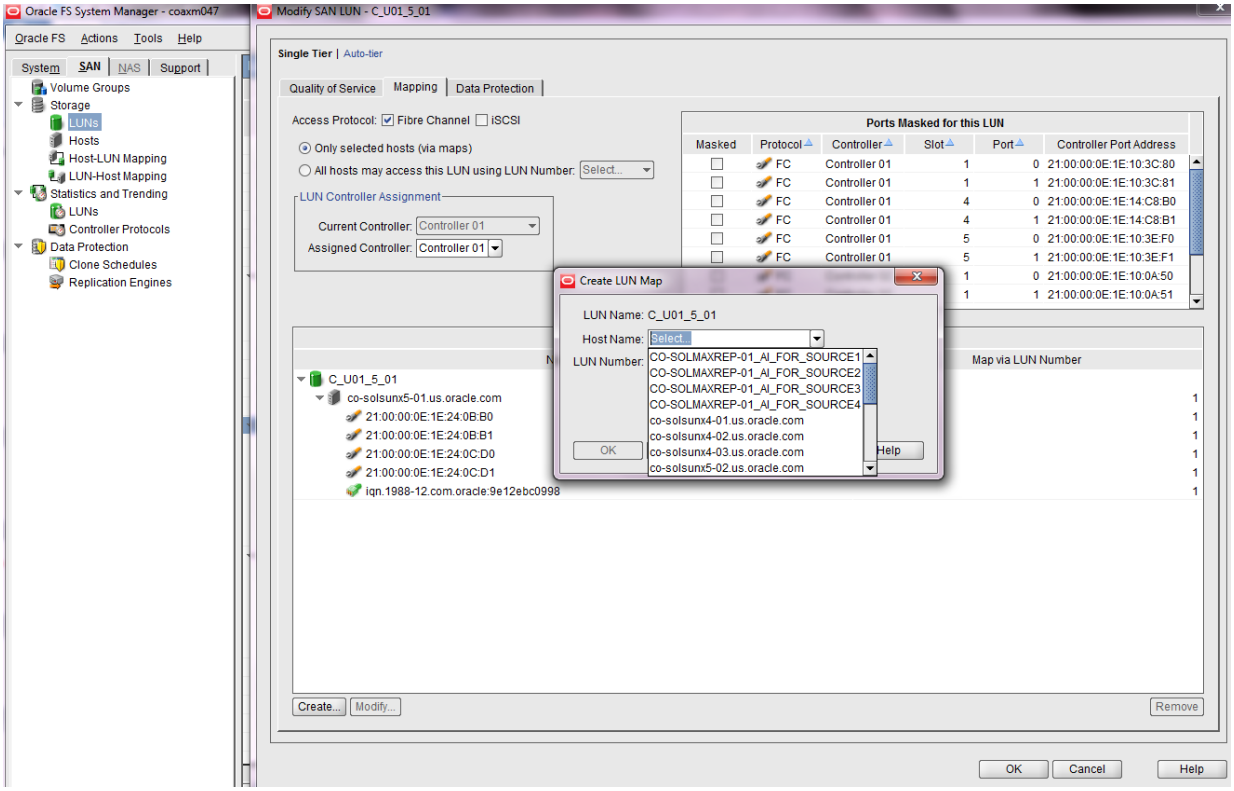

Figure 8. Mapping the BFS LUN and the application or swap LUNs to the correct servers on the FS System

NOTE: LUN 0 should always be used as the LUN mapping for the boot LUN.

As shown in Figure 9, the **Data Protection** tab provides options for automatically creating and managing clones for the LUN, including enabling Clone LUNs, and setting maximum capacity for clone storage. Based on the recovery point objectives, the administrator can also set the clone size and frequency of cloning.

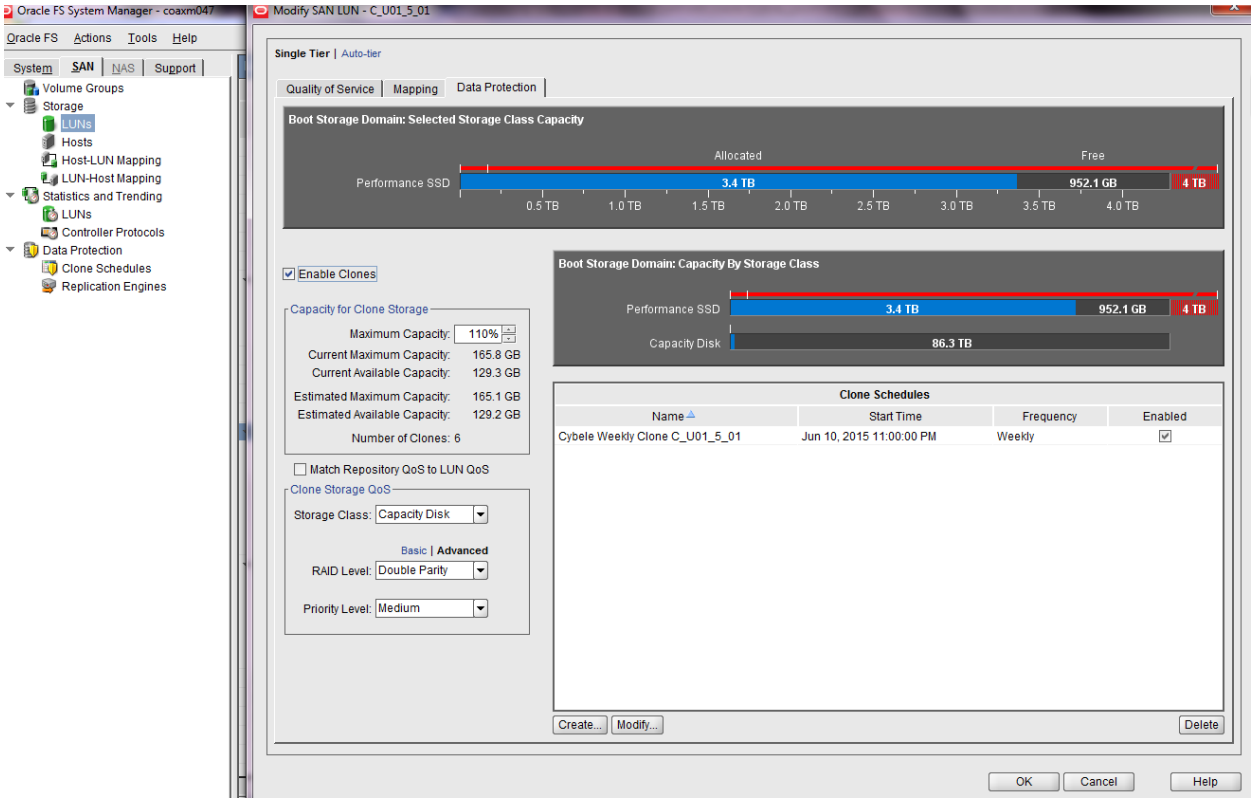

Figure 9. Setting data protection for the LUN by using the Oracle FS System Manager (GUI) to enable cloning

Configuring the Host Server for FC SAN Boot

Before setting up the operating system OS host server to support an FC SAN boot strategy, the administrator must confirm that an FC HBA is installed on the client and that the HBA supports SAN boot. Configuring the host server involves the following:

- Setting the boot precedence in the system BIOS to make the FC HBA card the highest priority boot device.
- Configuring the HBA to boot from the LUN on which the OS for that server has been installed in the FS System.

Set Boot Precedence in the System BIOS

The boot precedence in the system BIOS must be set so that the FC HBA card is the highest priority boot device.

NOTE: The following instructions apply to Oracle x86 servers. For other servers, refer to the manufacturer's documentation.

1. Reboot the server. While the server is initializing, press the **F2** key to display the system BIOS menu.
2. If the server has an LSI PCI card installed, disable the PCI slot in the system BIOS.
   In Oracle x86 servers, the LSI card takes higher boot precedence and will try to boot the local operating system. To prevent this from happening, complete the following steps:
   a. Select **Advanced** to display the PCI Configuration screen.
   b. Disable the PCI slot in which the LSI card is installed.
   c. Press **F10** to save the setting, exit the BIOS menu, and reboot the server.
3. Set the FC HBA card to be the highest priority boot device.
   a. From the BIOS menu, select **Boot** to display the Boot Device Priority screen.
   b. For the **1st Boot Device**, select the FC HBA.
   c. Press **F10** to save the setting, exit the BIOS menu, and reboot the server.

Configuring the HBA for FC Boot

One or more ports on the server's FC HBAs must be configured to boot from the LUN on which the OS for that server has been installed in the Oracle FS System. This involves making changes to BIOS settings as described briefly in the following procedure for configuring the QLogic HBA. Configuring an Emulex HBA involves making similar BIOS changes. For more information about supported FC HBA settings, refer to the Oracle FS Path Manager Installation guide for your operating system. See SAN Storage – Oracle Flash Storage Systems for the Oracle FS Path Manager Documentation Library.

Configure the QLogic HBA

1. Reboot the system. When the initialization screen appears, log in to the HBA BIOS menu.
2. Select one of the two Host Adapter HBA ports, then navigate to **Configuration Settings** > **Adapter Settings** to enable BFS.
3. In the Adapter Settings, enable **Host Adapter BIOS**.

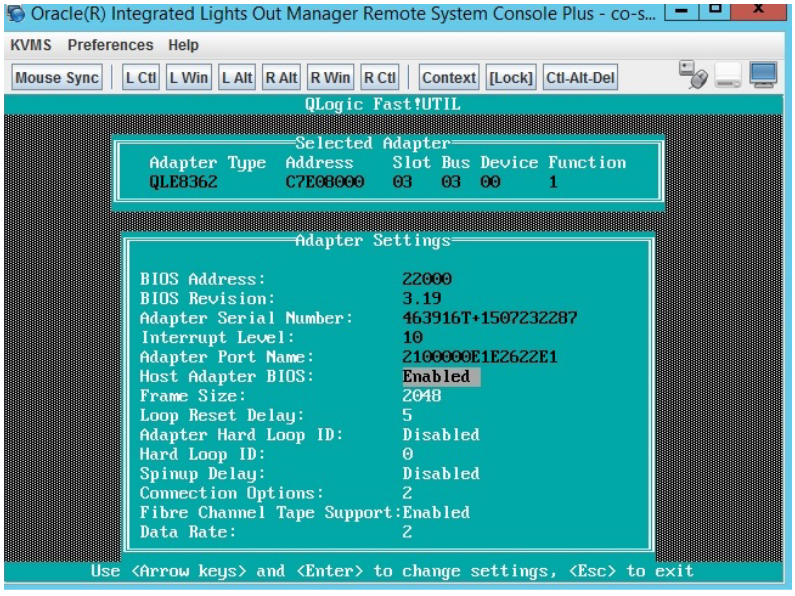

Figure 10. Adapter Settings screen for enabling BIOS settings

4. In the Configuration Settings, select **Selectable Boot Settings**, then enable **Selectable Boot**.
   A list of available FC devices is displayed.

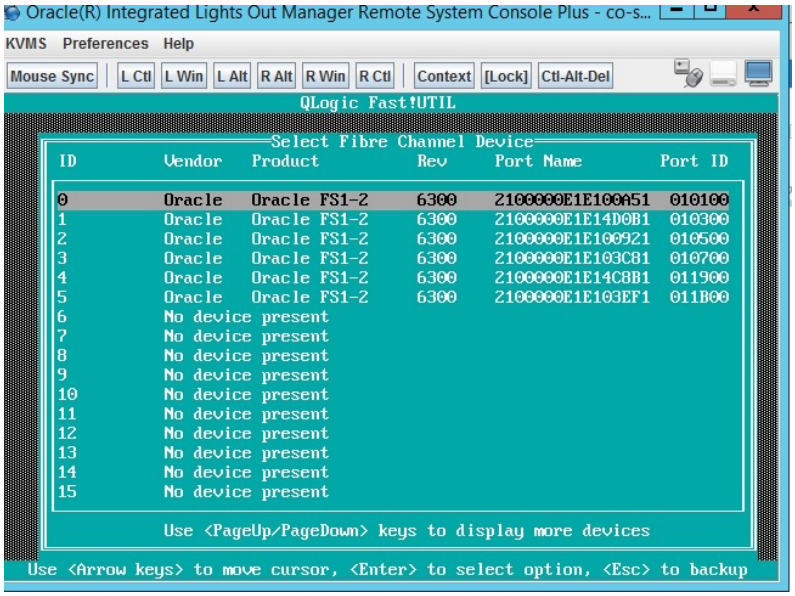

Figure 11. Select Fibre Channel Device screen for available target ports

5. Select the FC device to be used as the HBA target port on the FS System, then press **Enter**.
   A list of the available LUNs for the FC device is displayed.
6. Select the number of the LUN from which the server OS will boot.
7. From the BIOS menu, select **Boot**.
8. Select the FC HBA as the **1st Boot Device** with the highest priority.
9. Press **Esc**, then save the configuration settings.

The administrator can configure a second HBA port, repeating the same steps and using the same settings as for the first HBA port.

The new BIOS settings take effect when the server is rebooted. The server now uses the FC HBA as the primary boot device, and uses the primary boot setting in the HBA BIOS to select the LUN in the FS System from which to boot the OS.

Configure the Emulex HBA

Configuring the BIOS settings for an Emulex involves similar steps. However, an Emulex HBA offers the choice to boot the server using the World Wide Port Name (WWPN) or the Device ID (DID).

Installing the Operating System on the Server

Implementing a BFS strategy involves important considerations for each type of OS installation. The following sections contain instructions that apply to the installation of various OS versions. Additional best practices for configuring the storage array and implementing Clone LUNs and copies after installation are included.

Considerations for Installing Solaris

When installing Solaris on a server, the administrator must select the appropriate FC LUN from which to install the OS during the installation process.

The multipath I/O (MPIO) component must be configured as described in the instructions for installing the Solaris FS Path Manager (FSPM) software in *Oracle FS Path Manager Release 4 Linux Installation Guide*, *Oracle FS Path Manager 4 Release Notes*, and *Oracle FS Path Manager 4 Solaris Installation Guide*. See SAN Storage – Oracle Flash Storage Systems for the Oracle FS Path Manager Documentation Library.

Considerations for Installing Oracle Linux

When installing Oracle Linux on a server, during the installation process, the administrator must select **Advanced Configuration** to install the GRand Unified Boot Loader (GRUB) and the OS on the same FC LUN device. Otherwise, the GRUB master boot record (MBR) will be installed on the local drive and the OS will not boot from the primary boot FC LUN.

NOTE: The boot LUN must be a logical volume manager (LVM) Device Mapper LUN with ALUA enabled.

The MPIO component must be configured as described in the *Oracle FS Path Manager 4 Linux Installation Guide*.

Considerations for Installing SUSE Linux Enterprise 11 or 12

When installing SUSE Linux Enterprise 11 SP 4 or SUSE Linux Enterprise 12 on a server, during the installation process, the administrator must select **Advanced Configuration** to install the GRUB and the OS on the same FC LUN device. Otherwise, the GRUB MBR will be installed on the local drive and the OS will not boot from the primary boot FC LUN.

Considerations for Installing Microsoft Windows 2008, Windows 2012

When installing Microsoft Windows 2008 or Windows 2012 on a server, the administrator must complete the following steps:

1. Set the boot LUN as the primary boot device on the FC HBA on the server.
2. Proceed with the installation.
3. Install FSPM for each host as instructed in the *Oracle FS Path Manager 4 Windows Installation Guide*.

Best Practices for Configuring the FS System

To support the FC SAN boot solution after the appropriate LUNs have been set up, consider implementing the following recommendations:

- Configure storage on the FS System for the highest level of fault tolerance for FC SAN boot solutions. The system should use all the redundant paths that are available to avoid a single point of failure.
- Ensure Oracle FS high availability by including more than one of each component involved in the storage array. For example, use at least two Controllers, two switches, two HBAs, and multiple HBA ports from each server. Connect to both FS Controllers with a path to each from each of the two switches which connect to the two HBA's in each server. Add multiple servers that are clustered to allow for failover. The functioning of the storage array must not be effected by the failure of one or more components.
- Map storage LUNs to multiple HBA ports accessing the SAN. Identify multiple ports on the client server in the systems BIOS. Upon boot, the BIOS will go through the list of targets until the active path to the boot device is located.

  NOTE: ALUA is not supported for boot solutions, because most HBA firmware does not support ALUA.
- Separate the boot paths from the application data paths when configuring the FC boot device in an FS System when the storage array also hosts LUNS for other applications. Do not share FC target ports and Storage Domains that are designated for booting servers with FC target ports and Storage Domains that service applications. Doing so can adversely affect performance.
- Configure dedicated boot Storage Domains (mirrored) and separate application Storage Domains (variable).

- Map targets across alternate target ports on the FS System so that the targets are not shared with boot ports.
- Ensure that multiple HBAs can access the boot LUN through multiple paths and switches. This approach ensures high availability and fault tolerance, and helps prevent a single point of failure.

Best Practices for Using Clones and LUN Copies

Cloning or copying a LUN preserves a snapshot version of data from a point in time that can be restored when needed. To preserve reliable data, an administrator should clone or copy a LUN at specific stages before or after making changes.

For example, to preserve a reliable boot copy, the boot LUN should be cloned either after installation or prior to upgrade. This ensures that the Clone LUN can provide a reliable copy of the OS with a reliable recovery point for the boot LUN.

What is an Oracle FS Clone LUN?

An Oracle FS Clone LUN is a point-in-time, read-write partial-block snapshot of a LUN that can be accessed immediately. Changes made to the source LUN after the Clone LUN is created are not reflected in the snapshot. A Clone LUN shares space with the source LUN in the Clone LUN repository.

An Oracle FS Clone LUN uses copy-on-write technology and is therefore space efficient, because the Clone LUN only contains the changes from when the snapshot was created.

What is a snapshot?

A snapshot is a generic industry term for a point-in-time copy of a file, a LUN, a file system, or other type of container.

What is a LUN copy?

A LUN copy on an FS System is a Volume Copy (SAN and NAS), which is a full block-by-block copy of a LUN. This operation creates a new LUN with the same space requirement as the original LUN.

Changes made to the parent LUN after the LUN copy is created are not reflected in the LUN copy.

Managing a Dynamic Boot From SAN Environment

The following sections describe considerations for managing an environment that includes boot from SAN (BFS), and high-level procedures for rebuilding the server and reconfiguring the environment. A use case that provides an example implementation is included.

Best Practices for Managing a Boot From SAN Environment

- BFS LUNs should be mapped as LUN 0.
- LUN copies can be used to replicate OS images that can be used without a need for reinstallation.
- One or more Clone LUNs should be created before the OS or applications are updated.

  Setting up a policy to manage Clone LUNs to suit the requirements of the environment will greatly simplify management.
- In all cases, Oracle FS Path Manager for the appropriate OS should be downloaded from the Oracle Flash Storage System download site and installed as described in the Oracle FS Path Manager, Release 4 documentation.

» For environments running complex applications like Oracle Real Application Clusters (RAC) or Oracle Fusion, a separate LUN should be set up for boot only. The application should be installed on a different SAN LUN that is mounted separately.

Clone LUNs with the application installed provide flexibility, as does a BFS LUN. Applications can be shared across multiple hosts with no need for reinstallation, thereby saving time, maximizing storage usage, and optimizing space. Copies of these LUNs can also facilitate experimentation and testing of variations of the system.

Another time-saving consideration is the safe restore point that is always available. Also, numerous variations can be set up quickly by leveraging and reusing the time and space required for the initial configurations.

Centralized management of updates, snapshots, and clones can speed up the time required for backup to approach real time, enabling very high availability for applications and data.

Rebuild the Server Using the Boot LUN

Rebuilding a server from the boot LUN involves the following administrator actions:

1. Ensuring that UUID is not used in the `fstab/eth` scripts.
2. Removing the MAC address from the `eth` scripts.
3. Deleting the `/etc/udev/rules.d/70*net*` file.
4. Editing the `/etc/iscsi/initiatorname.iscsi` file to ensure that the "InitiatorName" on each server is unique.
5. Editing `the /etc/hosts` file and the `/etc/sysconfig/network` file to update the host name.
6. Editing `/etc/sysconfig/network-scripts/ifcfg-eth*` and `/etc/sysconfig/network-scripts/route-eth*` to update the MAC address and the IP address.
7. Editing `/etc/sysconfig/network-scripts/ifcfg-ib*` for the Infiniband address and the IPoIB address.
8. Modifying the IP address in `sshd_config`.
9. Running the following command:

   ```
   # sysctl kernel.hostname new_hostname
   ```

10. Running the `sys-unconfig` command as described in "Reconfigure the System in a New Environment," then rebooting.
11. Adding the new host name on the server and on BBSERVER.
12. Verifying `auto.home` and `resolv.conf` as required by the OS.
13. Verifying `RHEL.repo` or `OEL.repo` and any update patches that are required by the OS.

Reconfigure the System in a New Environment

The `sys-unconfig` command provides a simple way to reconfigure the system in a new environment.

As root user, type the following command when the system has booted up. The system will go down immediately as the destination disk is returned to an unconfigured state.

WARNING: Do not attempt to run this command on a production server, because the network settings will be disrupted. The system returns to an unconfigured state.

```
# sys-unconfig
```

The command creates a file named `/.unconfigured` which, at reboot, causes `/etc/rc.d/rc.sysinit` to reconfigure settings such as system authentication resources, system services, networking and routing, system runlevel, and other configuration programs. In addition, all persistent rules are deleted from `/etc/udev/rules.d/`.

To create a reusable boot LUN on the Solaris OS, refer to the following procedures:

» "Using the sysconfig Utility" from *Installing Oracle Solaris 11 Systems*

» "Reconfiguration Boot" from *Solaris Transition Guide*

For more information, refer to the administration guides from the vendor of your OS.

Use Case: Managing an Environment With a Boot LUN, Clones, and Copies

In this example, LUN BFS_5_01 has been configured to boot from SAN. Host 5_01 is running a Linux installation.

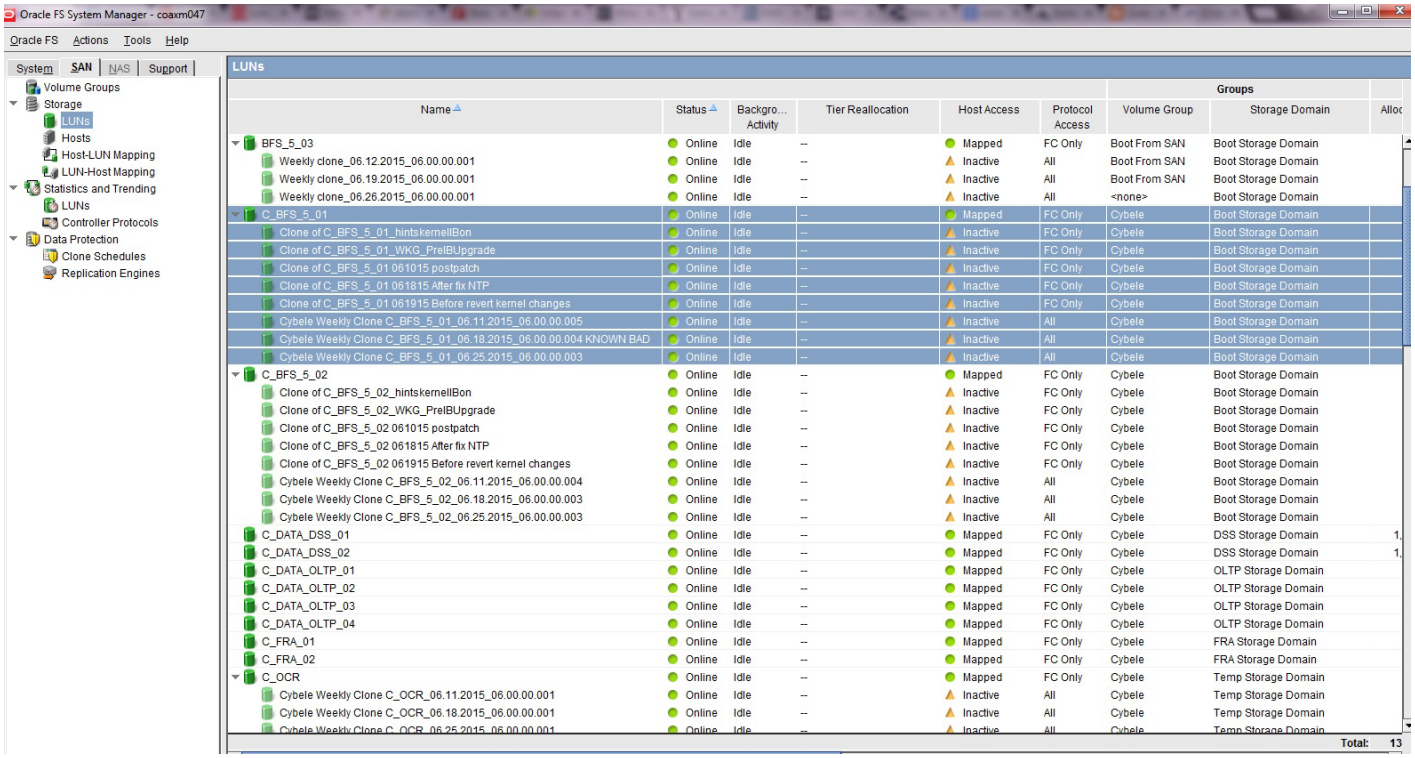

Figure 12. Oracle FS System Manager screen showing LUNs with their host mappings and Clone LUNs

In the Oracle FS System Manager (GUI) screen shown in Figure 12, LUN BFS_5_01 and its clones are highlighted. Before making changes to BFS_5_01, the administrator created the Clone LUNs to provide fail-safe backup in case future changes must be reverted back or the older configuration must be restored. The administrator can install different versions of the OS and applications on the Clone LUNs to ensure that variations of the system environment are available for use with minimal disruption.

For additional flexibility, the administrator exported LUN BFS_5_01 to a new host by mapping the LUN to the new host 5_02. Without doing a reinstall, the OS image from host 5_01, including all its changes, is preserved on the new host 5_02.

To do this, the administrator made a copy of LUN BFS_5_01 and named the copy BFS_5_02. Figure 13 shows the pop-up menu for cloning or copying LUN BFS_5_01.

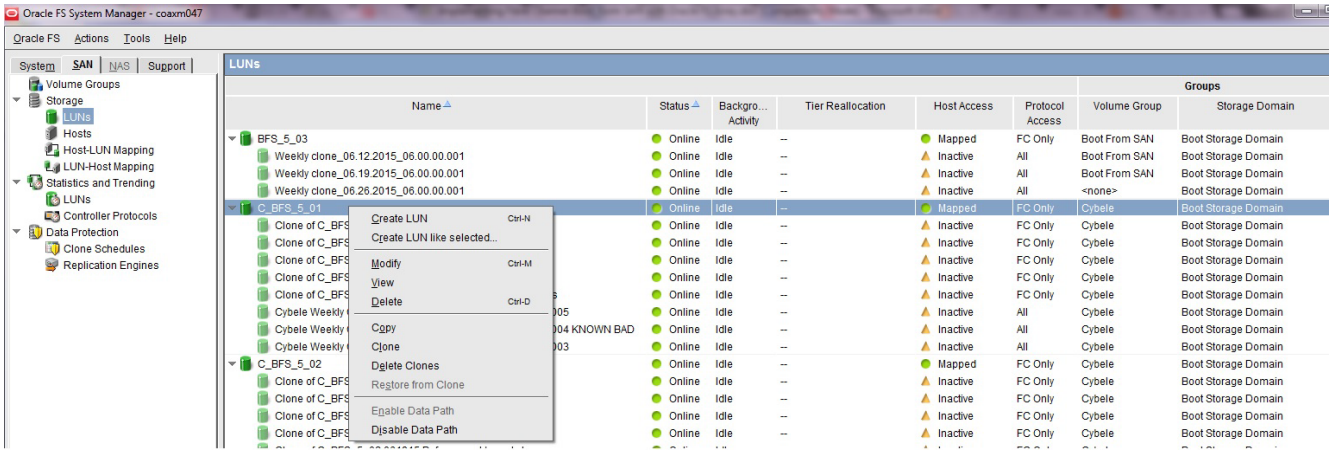

Figure 13. Oracle FS System Manager screen showing the right-click pop-up for copying or cloning a LUN

Eventually, the administrator needs to restore LUN BFS_5_01 to an earlier state. Figure 14 shows the pop-up menu that restores the LUN from a Clone LUN. Choosing this option returns LUN BFS_5_01 to its earlier state from when the clone was created.

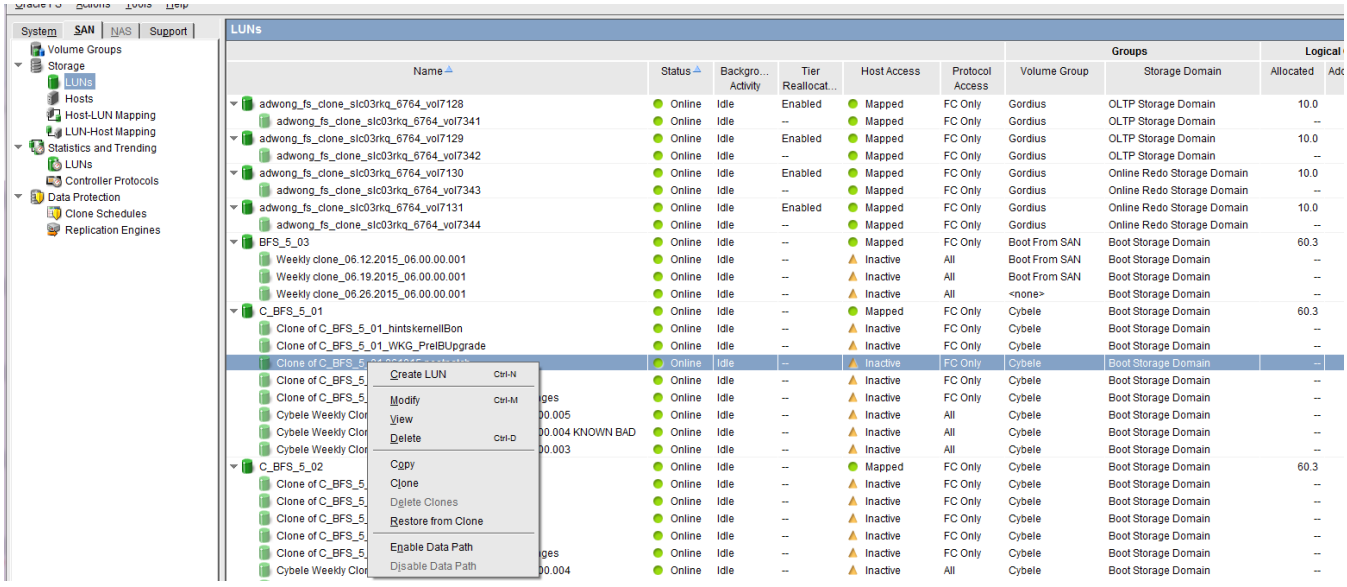

Figure 14. Oracle FS System Manager screen showing the right-click pop-up for restoring from a Clone LUN

Conclusion

This white paper showcased the advantages of implementing a BFS solution on an FS System. BFS enables IT organizations to simplify and centralize their IT management infrastructure, enhance scalability, and reduce costs while increasing reliability and performance. Other benefits of the BFS strategy include cost savings, enhanced security, ease of management, increased flexibility, and reduced downtime.

IT organizations that adopt BFS can support high availability enterprise applications like Oracle Real Application Clusters and Oracle Fusion with ease and reliability. The recommended implementations can achieve recovery point objectives and other features like server portability that are typically only available with virtualization, but without the additional overhead that virtualization requires.

This white paper recommended the Storage Profiles to use for a BFS strategy, and presented advanced BFS concepts, such as using Clone LUNs and LUN copies to increase productivity and scale an organization rapidly.

In addition to explaining how to implement BFS LUNs, this white paper also described the advantages of setting up application LUNs and swap LUNs. Best practices provide further recommendations and guidance for optimizing BFS capabilities. Finally, a use case provides a working example of managing a BFS environment, including sample screens from the Oracle FS System Manager (GUI).

References

Refer to SAN Storage – Oracle Flash Storage Systems for the latest versions of the following documents:

- Oracle Flash Storage System Documentation Library:
  - *Oracle Flash Storage System Administrator's Guide*
  - *Oracle Flash Storage System Glossary*
- Oracle Flash Storage System Utilities Documentation Library
- Oracle MaxRep for SAN Documentation Library
- Oracle FS Path Manager Documentation Library:
  - *Oracle FS Path Manager Release Notes*
  - *Oracle FS Path Manager Installation Guide for Linux*
  - *Oracle FS Path Manager Installation Guide for Solaris*
  - *Oracle FS Path Manager Installation Guide for Windows*
- *Oracle Flash Storage System Support and Interoperability Guide*

OS multipath configuration:

- *SUSE Linux Enterprise Server (SLES) 11 Storage Administration Guide*, "Planning for Multipathing"
- Microsoft Windows Server 2008 article, "Multipathing Support in Windows Server 2008"
- Oracle Linux Documentation, *Red Hat Enterprise Linux DM Multipath Configuration and Administration*

Oracle Linux 6 and 7 Operating Systems Documentation

Oracle Solaris documentation:

- Installing Oracle Solaris 11 Systems, "Using the sysconfig Utility"
- *Oracle Solaris Administration: Common Tasks*
- *Solaris Transition Guide*, "Reconfiguration Boot"
- Oracle Technology Network article, "How to Configure Oracle Solaris 11 Using the sysconfig Command"

» For troubleshooting information related to setting up the FC driver and ALUA, see *Solaris Express SAN Configuration and Multipath Guide*.

Other storage documents:

» *Sun Storage 16 Gb Fibre Channel PCIe Universal Host Bus Adapter, QLogic Installation Guide For HBA Model 7101674*

» White paper, "Best Practices for Optimizing Storage for Oracle Automatic Storage Management with Oracle FS1 Series Storage"

Oracle Corporation, World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries
Phone: +1.650.506.7000
Fax: +1.650.506.7200

CONNECT WITH US

blogs.oracle.com/oraclepartners

facebook.com/oraclepartners

twitter.com/oraclepartners

oracle.com/partners

Hardware and Software, Engineered to Work Together