

DRIVE SPARING IN EMC® SYMMETRIX® VMAX® FAMILY SYSTEMS

Applied technology

Abstract

Drive sparing significantly increases data protection and availability. EMC Symmetrix VMAX Family systems support Permanent and Direct Sparing. This white paper explains the benefits of Permanent Sparing and Direct Sparing and describes the sparing processes in VMAX systems running Enginuity™ 5874, 5875, and 5876 as well as VMAX All Flash and VMAX3 systems running HYPERMAX OS 5977. Configuration rules and considerations are also discussed.

April 2016

Copyright © 2016 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided “as is.” EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

Part Number h6541
Revision A16

Table of Contents

Executive summary	4
Audience	4
Introduction	4
Terminology	6
Replacement Sparing	7
Permanent Sparing Process	7
Direct Sparing	8
Sparing with a Preferred Spare	10
Sparing with a Regular Spare	10
Sparing with a Non-Preferred Spare	10
Benefits	12
Vault drives	12
VMAX 10K Series vault information	12
VMAX 20K Series vault information:	12
VMAX 40K Series vault information:	13
Sparing of Vault Drives	13
Spare Drive Replenishment Process	13
EMC Certified Data Erasure for Symmetrix Disks	14
Considerations	14
Permanent spare coverage	14
Configuration file drive map	15
Query Spare Pool Status	15
Full spare coverage	16
Spare rules	16
HDD spare rules	16
Enterprise Flash Drive spare rules	17
Sparing in VMAX All Flash and VMAX3 systems running HYPERMAX OS 5977	17
VMAX All Flash and VMAX3 spare drive requirements	19
VMAX All Flash and VMAX3 spare rules	19
Conclusion	19

Executive summary

As the demands of internal and external customers force IT organizations to deliver the very highest service levels, availability expectations have increased to where any interruption to the access of information is unacceptable. Binding service-level agreements (SLAs) commit IT organizations to deliver agreed-to measurable support metrics. Even in the face of a component failure, these organizations are required to provide continuous data availability without compromise.

IT executives around the globe have recognized that downtime is not only measured in minutes or hours, but is also calculated as potential revenue loss, missed opportunities, or dissatisfied customers. Any operational impacts resulting from a component failure must be completely transparent to the individual applications and users relying on information availability to drive the business.

EMC® Symmetrix® has many built-in design features to provide protection against data unavailability due to a single component or power failure. A redundant design allows Symmetrix systems to remain online and operational during service events and upgrades. Proactive error detection and remote support allow components to be replaced before actual failures occur.

EMC Symmetrix VMAX® Family systems implement Permanent and Direct Sparing to increase data availability without compromising performance or functionality. This white paper describes the functionality of Permanent and Direct Sparing in Symmetrix VMAX systems. Also discussed are the benefits and considerations of Permanent and Direct Sparing, as well as the spare drive configuration rules. This white paper addresses Symmetrix VMAX systems operating with Enginuity 5874, 5875, and 5876, as well as VMAX3 and VMAX All Flash arrays with HYPERMAX OS.

Audience

This white paper is intended for technology professionals who configure and support Symmetrix VMAX systems. It is also intended for operations and data-center personnel who plan, install, and configure Symmetrix VMAX storage arrays in their IT environments.

The information in this paper is based on current microcode behavior and is subject to change. EMC is continuously improving and modifying operational functionality of all features based on testing and events experienced in the field.

Introduction

Permanent Sparing is a process that replaces a failing drive with a spare drive. Permanent Sparing is initiated automatically upon detection of certain error conditions, reducing the amount of time that a failed or failing drive remains in the system. Permanent Sparing performs the replacement of the failing drive in the shortest time possible to help protect against multiple drive failure in a RAID group.

Direct Sparing, first introduced in Enginuity™ 5876, is an enhancement of Permanent Sparing. With Direct Sparing, the invoked spare drive is added as another member of the RAID group. During a drive rebuild, the option to directly copy the data from the failing drive to the invoked spare drive is supported. The failing drive is removed only when the copy process is finished, and like Permanent Sparing, Direct Sparing is initiated automatically upon detection of appropriate drive-error conditions.

All Symmetrix systems are self-diagnostic and proactive on parts replacement before an actual failure occurs. The sparing algorithms are designed to work in the same manner. The microcode determines when a physical volume is about to fail and begins the sparing operation. This decision is based on error statistics maintained within the array, the intelligent drive microprocessor self-testing information, and an active error-checking system. A calculated number of drives are reserved as standby spares to be used for Permanent and Direct Sparing. These drives are not user-addressable and are configured in various physical locations for maximum effectiveness.

Permanent Sparing is used in combination with all protection types, including RAID 1, RAID 5, and RAID 6. Volumes with SRDF® connectivity may also take advantage of Permanent Sparing. Permanent Sparing is used in Enginuity 5874, 5875 and, when necessary, 5876. Permanent Sparing offers the following benefits:

- Automatically initiates the replacement and rebuild upon error detection
- Reduces the time between drive failure and rebuild to new drives to reduce the risk of losing a second drive in the RAID group before rebuild is complete
- Ensures that the spare copy is identical to the original copy
- Increases data availability of all volumes in use without loss of any data capacity
- Is transparent to the host and requires no user intervention
- Does not require the use of an additional mirror position

Direct Sparing is used with protection types, RAID 1, RAID 5, RAID 6, except RAID 6 (14+2), which is not supported. RAID 6 (14+2) protection is achieved by way of the Permanent Sparing mechanism. Volumes with SRDF connectivity may also take advantage of Direct Sparing. Direct Sparing in Enginuity 5876 offers the following benefits:

- The Symmetrix system can copy data from the failing RAID member (if available), removing the need to read the data from all of the members and perform a parity rebuild. Copying to the new RAID member is less CPU intensive.
- If a failure occurs in another member, the Symmetrix system can still recover the data automatically from the failing member (if available).
- More than one spare for a RAID group is supported at the same time.

Terminology

The following are explanations of terms that are commonly used throughout the paper:

Data drive: A physical drive that has hypervolumes allocated on it that may be accessible to the host. Hypervolumes on two or more physical drives are used to form logical volumes.

Hypervolume: A single physical drive may be split into two or more hypervolumes. Each hypervolume is grouped into separate logical volumes. The logical volumes may also be of different protection schemes, including RAID 1, RAID 5, RAID 6, and unprotected volumes.

Note: Unprotected volumes have only one instance of the data. This protection method was typically reserved for short-term (for example, daily) backups or test data. Failure of an unprotected volume could cause significant delay in backups or other procedures, such as database consistency checks and quality assurance operations that leverage secondary copies of production data. EMC requires an approved RPQ for unprotected standard volumes.

Mirror positions: Symmetrix Enginuity allows four logical mirror positions for each logical volume. Beginning with Enginuity 5874, all local protection schemes occupy only one mirror position. This includes RAID 1, RAID 5, and RAID 6 volumes. Other Enginuity features may occupy additional mirror positions of a logical volume, such as SRDF.

For example, a RAID 1 volume with one SRDF mirror has two mirror positions occupied. There are two free mirror positions. If the RAID 1 volume had concurrent SRDF, then three mirror positions are occupied, and there is one mirror position free.

Not Ready spare: A spare drive that was previously a data drive and was replaced by Permanent Sparing. The drive has a Not Ready flag set that prevents it from being used until replaced with a new drive.

Permanent Sparing: An automated, self-healing process that runs on the service processor to permanently replace a failing drive with a spare drive through a configuration change. Permanent Sparing can also be thought of as a “permanent replacement.”

Direct Sparing: An automated, self-healing process that runs on the service processor that invokes a spare drive and adds it as another member of the RAID group. The failing drive is removed when the copy process is finished.

Spare drive: A physical drive that is not user-accessible and is reserved for use by Replacement sparing in the event of a drive failure.

Replacement Sparing

Replacement Sparing consists of two different technologies that VMAX systems use for sparing, Permanent and Direct Sparing. These are automated, self-healing tasks that run on the service processor to permanently replace a failing drive with a spare drive through a configuration change. The Symmetrix system calls home to notify the EMC Customer Support Center of this event. Replacement Sparing is enabled by default on all Symmetrix VMAX systems.

Permanent Sparing Process

When Symmetrix Enginuity detects that a drive is about to fail, the Permanent Sparing process begins. It looks for a spare drive of the same block size, capacity, and speed in a good location to permanently replace the failing drive by means of a configuration change. Permanent Sparing is used in:

- Enginuity 5874
- Enginuity 5875
- Enginuity 5876 with RAID 6 (14+2)

The Permanent Sparing process identifies a good spare location using the following rules:

- No RAID 1 mirror or RAID 5 group member on the same loop
- No more than two members of the same RAID 6 group on the same loop

Note: If the process cannot identify a spare in a good location, or if the Permanent Sparing process cannot begin for any reason, the system calls home to the EMC Customer Support Center to inform that immediate replacement is required.

When a suitable spare is identified, the Permanent Sparing process loads a new configuration file in which all the logical volumes initially configured on the failing drive are now configured on the selected spare (new) drive. The configuration change typically takes a few seconds to complete, during which time the configuration is locked. Data is then rebuilt onto the new drive. The Symmetrix system continues to process host I/O requests at the highest priority to minimize any effects on performance. The configuration is not locked during the rebuild process.

The failed drive becomes a Not Ready spare in the spare pool and can be replaced at a later date. Multiple spares are planned for each array to be available for subsequent drive failures. The Not Ready state prevents the drive from being introduced onto the back-end Fibre Channel loop in the event of system power cycle, system IML, or disk director IML. The Not Ready setting will be removed once the drive is physically replaced, at which time it will again become an available spare.

Figure 1 gives a high-level illustration of the Permanent Sparing process, using RAID 1 as the example. Permanent Sparing of RAID 5 and RAID 6 volumes is identical in that the configuration change moves the data volumes to the spare drive. The manner in which data is built onto the new location depends on the protection scheme of the

logical volume. A physical drive may be configured with logical volumes of different protection schemes.

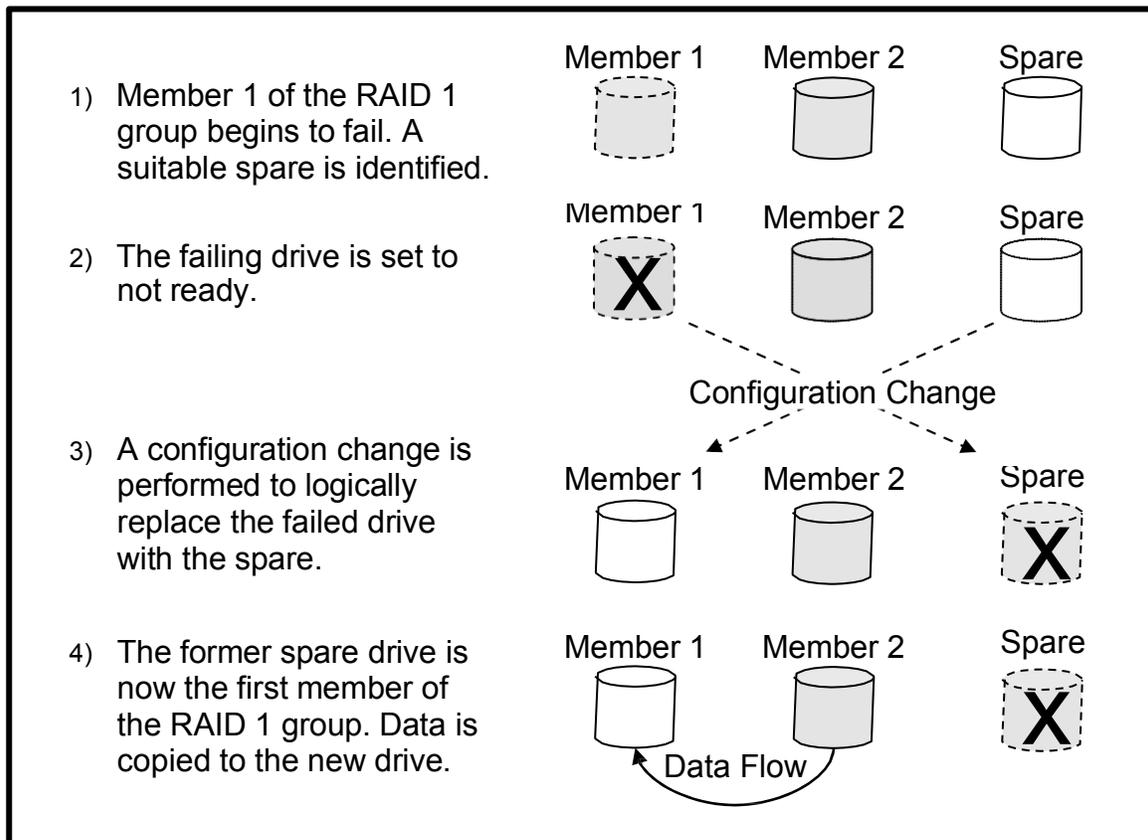


Figure 1. Permanent Sparring process

Direct Sparring

When Symmetrix Engineuity detects that a drive is about to fail or has failed, a Direct Sparring (DS) process is initiated. The DS process looks for available spares that are of the same block size, capacity and speed. The DS script always uses the best available spare. Direct Sparring is used in:

- Engineuity 5876

Spare drives are dynamically divided into three Spare Categories:

- Preferred Spare (Level 1)
- Regular Spare (Level 2)
- Non-Preferred Spare (Level 3)

The Direct Sparring process always picks a Preferred Spare when available. The DS script selects the best available drive by using the following rules as described in the table below.

Spare category	Description	Spare placement behavior		Script handling behavior	
		Same DA/loop distribution	Keeps high availability configuration rules	Failing drive gets restored to its original positions	Spare Drive Replenishment Process*
Preferr ed Spare (Level 1)	<p>The spare is configured on the same loop as the failing drive. The spare is mapped to the same DA Slice as the failing drive.</p> <p>Use of this spare will result in a valid RAID or Power Vault device configuration.</p> <p>This spare does not create a configuration with multiple RAID members on the same loop.</p>	Yes	YES	NO	YES
Regula r Spare (Level 2)	<p>The spare is not configured on the same loop as the failing drive.</p> <p>Use of this spare will result in a valid RAID or Power Vault device configuration.</p> <p>This spare doesn't create a configuration with multiple RAID members on the same loop.</p>	NO	YES	YES	YES
Non-Preferr ed Spare (Level 3)	<p>The spare is not configured on the same loop as the failing drive.</p> <p>Use of this spare results in a Non-Preferred RAID configuration (multiple RAID members configured in the same loop) and Power Vault devices may be moved to a different loop or DA.</p>	NO	NO	YES	NO

*See page 13 for info on the Spare Drive Replenishment Process

Sparing with a Preferred Spare

A configuration change is initiated that converts the spare drive into a new RAID member and converts the failing or failed drive to an auxiliary RAID member.

The DS script starts a copy operation to the new RAID member (the spare drive that was just converted as the new RAID member). The DS script creates a Symmwin scheduler task called Sync Split Erase and terminates.

The script Sync Split Erase monitors the copy process until it is 100 percent completed. The sync Split Erase script then performs a configuration change that converts the auxiliary RAID member (the failing or failed drive) to a bad or Not Ready spare. The VMAX system then calls home asking for the bad spare drive to be replaced.

At this point, the system is fully restored to its original state, so the replacement of the spare may be deferred. When the Not Ready spare is replaced the operation is complete.

Sparing with a Regular Spare

A configuration change is initiated that converts the spare drive into a new RAID member, and converts the failing or failed drive to an auxiliary RAID member. The DS script also records that the failing drive needs to be reconfigured to its original position.

The DS script starts a copy operation to the new RAID member (the spare drive that was just converted as the new RAID member). The DS script creates a Symmwin scheduler task called Sync Split Erase and terminates.

The script Sync Split Erase monitors the copy process until it is 100 percent complete. The Sync Split Erase script then performs a configuration change that converts the auxiliary RAID member (the failing or failed drive) to a bad or Not Ready spare. The VMAX system then calls home asking for the bad spare drive to be replaced.

At this point, the system is fully restored to a fully protected state, so the replacement of the spare may be deferred. When the Not Ready spare is replaced, the system initiates a spare replacement, and steps 1-3 of the process occur again, restoring the data drive and spare drive back to their original physical locations. This completes the operation.

Sparing with a Non-Preferred Spare

A configuration change is initiated that converts the spare drive into a new RAID member, and converts the failing or failed drive to an auxiliary RAID member. The DS script also records that the failing drive needs to be reconfigured to its original position.

The DS script starts a copy operation to the new RAID member (the spare drive that was just converted as the new RAID member). The DS script creates a Symmwin scheduler task called Sync Split Erase and terminates.

The script Sync Split Erase monitors the copy process until it is 100 percent complete. The sync Split Erase script then performs a configuration change that converts the auxiliary RAID member (the failing or failed drive) to a bad or Not Ready spare. The VMAX system then calls home asking for the bad spare drive to be replaced. A non-preferred spare does not comply with Raid configuration rules so the drive replacement will not be deferred.

When the Not Ready spare is replaced, the system initiates a spare replacement, and steps 1-3 of the process occur again restoring the data drive and spare drive back to their original physical locations. This completes the operation.

Figure 2 gives a high level overview of the DS process when a preferred spare is being used. RAID 5 (3+1) is being used in this example.

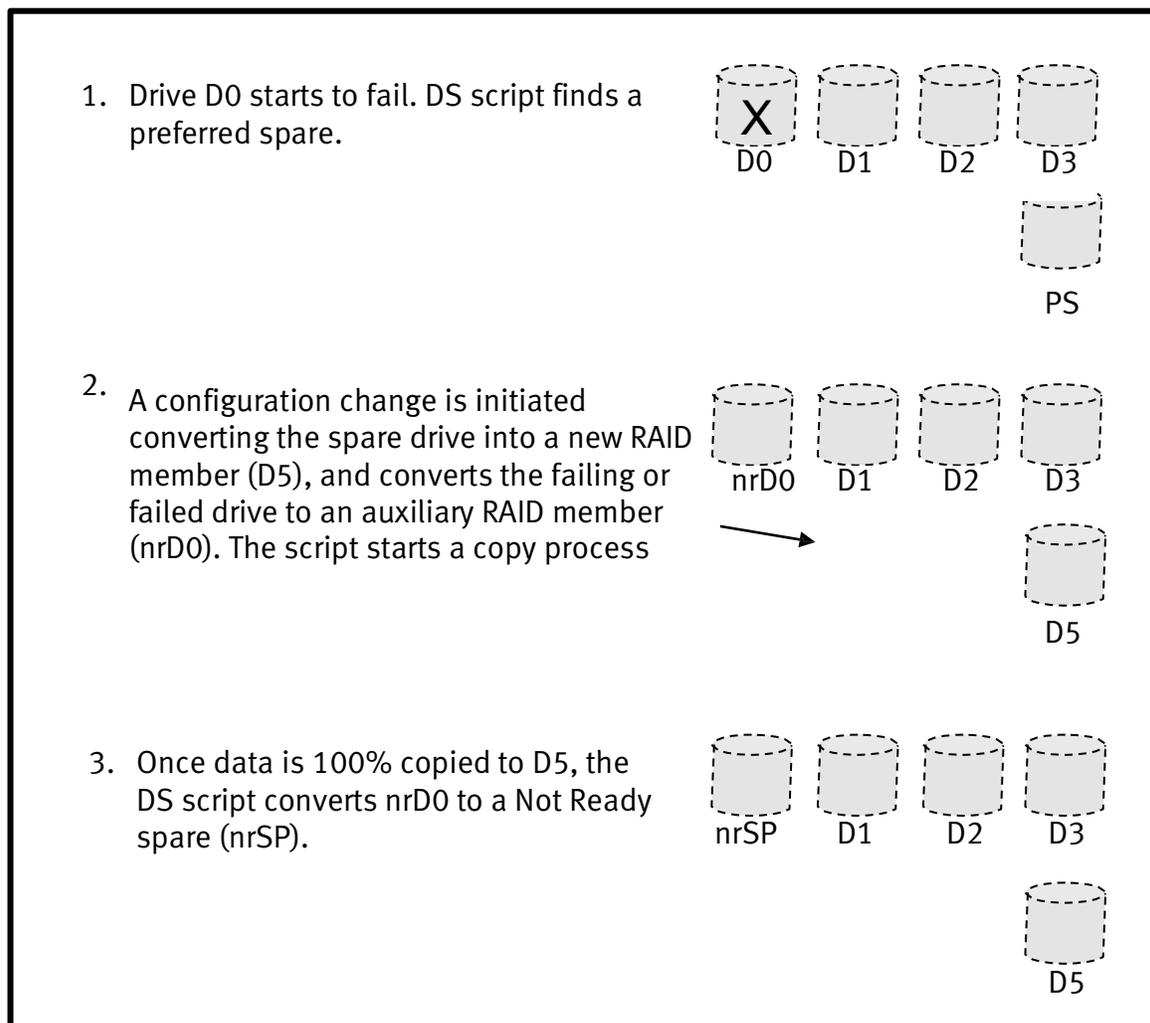


Figure 2. Direct Sparing process

Benefits

The benefits include:

- Permanent and Direct Sparing allows the system to get back to N+1 protection status in the shortest possible time frame. Both processes run unattended upon detection of a drive failure. The amount of time that a data drive remains not protected in the system is reduced.
- Only a single copy process is required, significantly reducing the amount of time required to rebuild and copy back the data from the failed drive.
- Permanent Sparing does not require an available mirror position. This avoids potential conflicts with other Enginuity features.
- Direct Sparing uses less DA resources.
- Direct Sparing keeps the failing drive as part of the RAID group until the new drive is rebuilt.
- In Direct Sparing, the data is copied from the failed or failing drive instead of being rebuilt.
- Drive rebuilds are completed faster in Direct Sparing.
- Direct Sparing also reduces the risk of a second RAID member failure, which can cause the customer a DU/DL event.

Vault drives

Symmetrix VMAX systems are configured with vault drives on back-end Fibre Channel loops. The total capacity of all vault volumes in the system is at least sufficient to keep two logical copies of the persistent portion of physical memory. The vault drive rules for an initial configuration on a newly installed VMAX are as follows:

VMAX 10K Series vault information

- 9 GB of space for vault data, including metadata, is reserved on dedicated vault devices.
- Each director pair requires 20 such devices, for a total of 180 GB of vault space per director pair.
- Five drives per loop containing vault devices are required on the first four loops of each engine.

VMAX 20K Series vault information:

- 5 GB of space for vault data, including metadata, is reserved on dedicated vault devices.
- Each director pair requires 40 such devices, for a total of 200 GB of vault space per director pair.
- Five drives per loop containing vault devices are required on all 8 loops of each engine.

VMAX 40K Series vault information:

- 9 GB of space for vault data, including metadata, is reserved on dedicated vault devices.
- Each director pair requires 40 such devices, for a total of 360 GB of vault space per director pair.
- Five drives per loop containing vault devices are required on all 8 loops of each engine.

Sparing of Vault Drives

Drives containing a Vault device are eligible for Permanent Sparing with Enginuity 5874 and 5875, and Direct Sparing with Enginuity 5876 (Preferred, Regular, and Non-Preferred). Sparing is allowed to relocate the contents of one vault drive on each fibre loop to a suitable spare drive on another disk director in the system. This ability greatly enhances the capabilities of Permanent Sparing by increasing its coverage for the vault drives in the system. Permanent Sparing can relocate the remaining vault drives across the loops on the same disk director if a suitable spare drive is available.

Spare Drive Replenishment Process

In addition to the benefits of Permanent and Direct Sparing, Symmetrix VMAX systems can take advantage of a default feature called Spare Drive Replenishment Process. This feature allows for spare drives that have been replaced by Permanent or Direct Sparing to not be called out for immediate replacement.

The Replenishment Process has the added benefit of allowing the actual physical replacement of the failed drive to be scheduled for a time slot that is convenient for the customer. Other spares will be available should another drive fail before the Not Ready spare is replaced. Furthermore, regardless as to whether or not the Spare Drive Replenishment Process is enabled or not, the actual replacement time of a spare drive is significantly less than that of a data drive since no copy or rebuild process is required depending on the sparing method used.

With the Replenishment Process enabled, a successful sparing operation generates a call home, creating a service request for the event. As long as all drives in the system still have spare coverage, the service request is automatically closed.

At the point when a Sparing operation causes any drive in the system to no longer have spare coverage, the system generates a log file named DDS.log that contains all of the specific details necessary for the replacement of all spare drives that were previously used by a Sparing operation.

The system calls home and generates a service request containing the log file, which is dispatched to an EMC Customer Engineer (CE). The CE contacts the customer to schedule a convenient time to conduct the replacement of multiple drives during a single trip to the customer's site.

If for any reason Sparing operation is unable to run on a failing drive, the system calls home to alert the EMC CE for same-day replacement.

To enable the Replenishment Process, an EMC representative must check the Enable Deferred Service flag in the Symmwin configuration program, residing on the service processor. This is a supported online change and does not require a configuration change or configuration lock.

EMC Certified Data Erasure for Symmetrix Disks

Data Erasure erases data per entire physical drive. The erasure can vary between a single pass write of data to the drive, all the way to a multi-pass erasure using complex algorithms as required in DoD 5220.22-M. The intent of secure disk erase is to remove all recoverable user data from a given physical drive, or report that the attempt to remove the data was unsuccessful.

If the EMC Certified Data Erasure for Symmetrix disks service license has been purchased, customers receive a certificate of completion for all disks successfully erased. Non-erasable drives are left with the customer for disposition.

Data Erasure is only able to run on a Not Ready spare upon successful completion of a Permanent Sparing operation. This prevents a drive that is visible to the host from being exposed to Data Erasure. Data Erasure is not currently supported for Enterprise Flash Drives.

More detailed information on EMC Certified Data Erasure for Symmetrix disks is available at EMC Online Support or through your local EMC Service Representative.

Considerations

Not all drive failures may be candidates for Permanent or Direct Sparing. One example is if a suitable spare cannot be identified. A failing drive cannot be permanently spared by a spare drive of a different speed, block size, or capacity. Further, neither Permanent nor Direct Sparing is able to run if the configuration is locked.

Permanent Sparing will not run if any unprotected volumes exist on the failing drive, as it would be considered a data-unavailable situation. In any data-unavailable situation, the system will call home to the EMC Customer Support Center and support personnel will investigate.

Direct Sparing will run in this situation if the auxiliary RAID member ensures that there is no data loss. If the failing drive is completely not ready, and a data-loss situation could occur, the system calls home to the EMC Customer Support Center and support personnel will investigate.

Both Permanent Sparing and Direct Sparing will not run if there are multiple drives not ready on a single DA processor. The cause of the failure may be something other than the drive itself. Again, the system calls home to the EMC Customer Support Center and support personnel will investigate.

Permanent spare coverage

With new system configurations, the required amount of spares is automatically taken from the drive resource pool. The spare drives are evenly distributed among the available disk directors. Local EMC Service Representatives can override this

algorithm and manually decide on the physical placement of the spare drives.

Sparing rules also apply when adding new drives (spare drives and data drives) to an existing system.

There are two ways to determine the Permanent Sparing coverage for each data drive. Both methods are currently available only to EMC personnel with access to the SymmWin configuration program residing on the service processor. Contact your local EMC Service Representative for assistance with determining Permanent spare coverage.

Configuration file drive map

The drive map in the Symmetrix configuration file has an option to view Permanent Sparing coverage. When selected, each drive in the map displays a number indicating how many spare drives it can use for Permanent Sparing. When a specific drive is selected, the acceptable spares are indicated. The drive map is based on the physical location of the spare drives and does not take into account if any spares are set to Not Ready. For Direct Sparing, this view capability only shows Preferred and Regular spares. There is no accommodation in this tool to view or predict the potential use of Non-Preferred spares.

Query Spare Pool Status

With Engenuity 5874 and 5875, the Query Spare Pool Status script queries the current status of the system to provide a detailed report of the status of the spare drives and the Permanent Sparing coverage of the data drives.

For each data drive configured in the system, the script reports which spare is the first choice, when available, for Permanent Sparing coverage. The script also reports which spare is the current choice. These two fields typically are the same, unless the first choice spare is currently unavailable.

The script also displays the following information:

- Total number of data drives
- Total number of spare drives
- Number of data drives that have Permanent Sparing coverage
- Number of data drives that do not have Permanent Sparing coverage
- Physically ready spare drives
- Physically not ready spare drives
- Spare drives that have the Not Ready bit set

All of this information is then stored in a log file on the service processor.

With Engenuity 5876, a Dynamic Report is used to query the status of spare drives. The Dynamic Report is available in the SymmWin configuration program residing on the service processor.

Full spare coverage

Starting with the Q4 2012 5876 Engenuity release, EMC offers full spare coverage. Full spare coverage is achieved by building initial configurations (new installations) of VMAX 10K, VMAX 20K, and VMAX 40K Series systems on that or higher 5876 Engenuity.

Spare coverage beyond previous sparing capabilities is achieved with the use of two Direct Sparing operations. The first DS operation allows for the replacement of the faulty drive, the second DS operation restores the data to the new drive in the original location.

When this technique is used, the first DS operation temporarily allows a spare drive to break RAID protection rules, such as creating a configuration with multiple RAID members on the same DA loop. The second DS operation returns the configuration to a fully redundant state.

Notes:

Full spare coverage is not offered with RAID 6 (14+2), as DS does not support RAID 6 (14+2) devices.

If the VMAX system was upgraded from 5875 or the initial release of 5876 to the Q4 2012 5876 release or a newer version of 5876 and spares have been invoked, EMC does not guarantee full spare coverage.

If the VMAX system was upgraded from 5875 or the initial release of 5876 to the Q4 2012 5876 release or a newer version of 5876 and spares have NOT been invoked, the default spare configuration is unchanged and full percent spare coverage is available going forward.

Spare rules

These rules for the numbers and types of spares required for each Symmetrix VMAX system are designed to minimize risk by providing multiple spares in various locations for each type of data drive configured. The system's current configuration as well as expansion potential is taken into account. Spares must be configured for each of the system's drive types (capacity, speed, and block size).

The spare rules are considered to be the minimum, and additional spares can be configured. Spare drives on every loop for every type of data drive on the loop gives the best possible Permanent Sparing coverage. Direct Sparing makes use of temporary relaxation of these stringent guidelines allowing full spare coverage with fewer spares.

HDD spare rules

Spare drives are required for every Symmetrix VMAX system. The configuration program residing on the service processor prevents the creation of Symmetrix VMAX configurations that do not contain the required numbers and types of spare drives. A valid pool of spares for each drive type must be available.

The amount and types of spares required for Symmetrix VMAX systems are calculated as follows:

- VMAX 20K and VMAX 40K Series systems running 5876 Enginuity, two spares are required for every 100 drives, with a minimum of eight spares.
- VMAX 10K Series systems running 5876 Enginuity, two spares are required for every 100 drives, with a minimum of four spares.

Enterprise Flash Drive spare rules

The system spare drive requirements discussed previously are not affected by the configuration of Enterprise Flash Drives or Flash drive spares. The system's required amount and types of spares are solely determined by the configured magnetic disk drives. Flash drive spare requirements are determined independently. Like the HDD spare rules, spares are required for each type of Flash drive configured in the system.

The amount of spares required for each type of Flash drive in a Symmetrix VMAX system is calculated as follows:

One Flash drive spare is required for the first 32 Flash drives configured in the system.

If more than 32 Flash drives are configured, then two Flash drive spares are required for every 100 physicals or portion thereof.

Sparing in VMAX All Flash and VMAX3 systems running HYPERMAX OS 5977

HYPERMAX OS provides Direct Sparing to automatically replace a failing drive with a spare drive. Direct Sparing is supported with all protection types, including RAID 6 (14+2).

Two traditional sparing factors, vault drives and drive location on the fibre loop, do not apply to VMAX All Flash and VMAX3 systems. These systems vault to flash I/O modules on the engine rather than vault drives, and HYPERMAX OS can dynamically relocate spare drives across backend directors within the same engine.

The major factor with Direct Sparing is the power zone within the disk enclosure where the spare drive, failing drive, and other RAID members are located. RAID 1 and RAID 5 are only allowed to have one member per power zone, and RAID 6 is allowed to have up to two members per power zone.

HYPERMAX OS uses the same categories for spare drives; Preferred, Regular, and Non-Preferred. The DS script selects the best available drive by using the rules in the following table.

Spare category	Description	Spare placement behavior		Script handling behavior	
		Same DAE power zone distribution	Keeps high availability configuration rules	Failing drive gets restored to its original positions	Spare Drive Replenishment Process*
Preferr ed Spare (Level 1)	<p>The spare is configured in either the same DAE power zone as the failing drive, or in a different DAE power zone where another RAID member does not already exist.</p> <p>Use of this spare results in a valid RAID configuration.</p> <p>This spare does not create a configuration with multiple RAID members in the same DAE power zone.</p>	Yes	YES	NO	YES
Regula r Spare (Level 2)	<p>The spare is not configured in the same DAE power zone as the failing drive.</p> <p>Use of this spare results in a valid RAID configuration.</p> <p>This creates a legal configuration with multiple RAID members in the same DAE power zone.</p>	NO	YES	YES	YES
Non-Preferr ed Spare (Level 3)	<p>The spare is not configured in the same DAE power zone as the failing drive.</p> <p>Use of this spare results in a configuration that breaks the rules for RAID member/DAE power zone distribution.</p>	NO	NO	YES	NO

VMAX All Flash and VMAX3 spare drive requirements

Spare drives are required for every disk group in every configuration. Direct Sparing cannot use a spare drive in one disk group to replace a failing drive in another disk group.

Local RAID configures all members of a RAID group behind a single engine. This improves backend performance and drive rebuild time while still supporting the dual-initiator failover/failback model. Local RAID is implemented in all systems.

In order to provide adequate spare coverage for all RAID members, a valid pool of spares is required for each disk group in each Engine.

VMAX All Flash and VMAX3 spare rules

The amount and types of spares required are calculated as follows:

- One spare drive is required for every 50 drives per disk group per engine
 - This calculation applies to both HDD and EFD drives.
 - There is no minimum spare count for the system as a whole. The required spares are calculated purely per disk group.

As with other VMAX systems, the spare rules are considered to be the minimum, and additional spares can be configured.

Conclusion

Today's mission-critical environments require a high-end storage solution that guarantees uncompromised levels of availability. This solution must include a robust architectural design to withstand any and all potential failures without impacting data availability. High-end availability is more than just redundancy; it means nondisruptive operations and being always online.

When a drive failure occurs, the respective RAID group is temporarily exposed. Sparing reduces the time that the RAID group is exposed. The sparing process is transparent to the host and requires no user intervention. While the sparing process takes place, host I/O requests continue to be processed at the highest priority to minimize the effect on performance.

Sparing is one of the advanced data-availability features that has elevated customer expectations for high-end storage to a new level, which makes EMC storage systems the ideal choice for critical applications and 24/7 environments that demand uninterrupted access to information.