

# HP StorageWorks 4x00/6x00/8x00 Enterprise Virtual Array configuration best practices to optimize availability white paper



Abstract.....	2
Background .....	2
Overview.....	2
Best practices to optimize availability .....	2
Optimizations “inside the box” .....	2
Use Vraid1 .....	2
Do not use Vraid0 .....	3
Do not use Vraid0 with HP StorageWorks Continuous Access EVA.....	4
Use single protection level .....	4
Keep a multiple of eight disks in each disk group .....	5
Utilize multiple disk groups .....	6
Maintain sufficient free space .....	7
Follow the proper procedure to replace failed disks.....	9
Follow the proper procedure to add new disks.....	10
Maintain all components at current revision levels .....	11
Use the correct Fibre Channel HDD type for its intended purpose .....	11
Optimizations “outside the box” .....	11
Maintain operational environmental specifications .....	11
Implement pro-active remote fault monitoring (ISEE/HPCC).....	12
Run database integrity checks .....	12
Maintain proven backup and recovery processes .....	12
Implement fully redundant fabric configuration .....	13
Implement remote mirroring—HP StorageWorks Continuous Access EVA .....	13
Implement disaster recovery solutions.....	13
Deploy the highest service level.....	14
Sign up for Subscriber’s Choice .....	14
Glossary.....	15
For more information.....	16

## Abstract

When considering data availability, several factors must be taken into account. Factors to consider include the physical configuration of the storage system and the storage area network (SAN), the availability and usage of supporting software and toolsets, and ensuring that proper procedures to maintain and support the environment are followed.

This white paper highlights these key areas as they apply to the HP StorageWorks Enterprise Virtual Array (EVA) and provide best practice recommendations to optimize availability.

## Background

Choices made to increase availability may have an impact on the overall cost of the solution as well as performance. This document focuses strictly on availability. Much of this information can also be found in the "HP StorageWorks 4x00/6x00/8x00 Enterprise Virtual Array configuration best practices white paper," where it is presented in context with other best practices focused on reducing cost or optimizing performance. The priority of implementation of the recommendations presented in this document is best determined by the customer after considering the potential impact on cost and performance, if any does exist.

## Overview

The HP StorageWorks EVA is designed for high data availability (business-critical) applications. Redundancy features enable the array to continue operation after a wide variety of failures. The EVA is also designed to be flexible. Some configurations allow higher availability by limiting exposure to failures that exceed the fault-tolerant design of the array. The goal is to create configurations that have the most independent protection domains so that a failure in one domain does not reduce the resiliency of another domain.

When discussing the various factors affecting availability, it can be useful to consider these factors as being either "inside the box" or "outside of the box" optimizations. For purposes of discussion, those optimizations performed directly on the array including specific configuration recommendations or best practice procedures for maintaining the array are considered "in the box" optimizations. Those optimizations performed in the environment surrounding the array are considered "out of the box" optimizations.

The following guidelines address the "inside the box" and the "outside of the box" optimizations to improve availability in sequential and multiple simultaneous failure scenarios.

## Best practices to optimize availability

### Optimizations "inside the box"

#### Use Vraid1

While Vraid5 provides availability and data protection features sufficient for most high-availability applications, some applications may require the additional availability and data protection features of Vraid1. Vraid1 configurations can continue operation in failure scenarios where Vraid5 cannot. A statistical model of the EVA shows that, for an equivalent usable capacity, Vraid1 provides over four times the data protection of Vraid5<sup>1</sup>.

---

<sup>1</sup> A statistical analysis is a prediction of behavior for a large population; it is not a guarantee of operation for a single unit. Any single unit may experience significantly different results.

This additional redundancy comes with additional cost caused by additional storage overhead and capacity requirements. Nevertheless, some applications or file sets within an application warrant this additional protection.

If cost constraints do not allow a total Vraid1 configuration, consider using Vraid1 for critical files or data sets. For example:

- In database applications, select Vraid1 for log files.
- In snapclone/snapshot applications, select Vraid1 for active data sets and Vraid5 for snapclones and snapshots. Another option available is mirrorclone. For performance reasons, it is best that mirrorclones use the same RAID level as the source vdisk.
- For Continuous Access Write History Logs use Vraid1.

---

**Best practice for highest availability**

Vraid1 provides the highest levels of availability and data protection.

---

---

**Note**

The higher availability and data protection capabilities of Vraid1 or Vraid5 should not be considered a replacement for good backup and disaster recovery processes. The best practices for business-critical applications always include frequent data backup to other near-line or offline media or devices.

---

**Do not use Vraid0**

Unlike Vraid1 or Vraid5, Vraid0 provides no data redundancy. Vraid0 is optimized for applications where data protection is not a requirement. Because Vraid0 has no redundancy, data in Vraid0 requires less physical capacity, and performance is not affected by additional operations required to write redundancy information. Thus, while Vraid0 provides the best performance for write-intensive workloads and the lowest cost of storage, it also provides the least availability.

---

**Vraid0 best practice to optimize availability**

Vraid0 is not advised for availability. Vraid0 provides no disk failure protection.

---

---

**Note**

For Vraid0, increasing the protection level of the disk group does not increase the availability of the virtual disk. A single disk failure renders a Vraid0 unit inoperable even when single or double protection is enabled within the disk group.

---

## Do not use Vraid0 with HP StorageWorks Continuous Access EVA

When HP StorageWorks Continuous Access EVA is used, LUNs can be logically associated into data replication (DR) groups. DR groups allow a database to recover a remote copy with transaction integrity.

Two characteristics of remote mirroring are consistency and currency. Currency refers to the time difference between the remote and local copies. Consistency refers to the content difference between the mirrored LUNs and the local LUNs. A consistent remote LUN is either equal to the content of the local LUN or equal to the content of the local LUNs at a past point in time. Synchronous replication by way of Continuous Access EVA provides mirror copies that are both current and consistent.

Asynchronous replication provides mirror copies that are consistent, but may not be current (they may be slightly delayed). To accomplish asynchronous consistency, Continuous Access EVA maintains the remote write ordering within each DR group. The order of the remote writes is identical to the order of the local writes. The writes may be delayed by the transmission distance, but the content of the LUNs in a DR group matches the current or a previous state of the associated local LUNs.

If a remote LUN becomes unavailable, Continuous Access EVA cannot continue to write to the remaining LUNs in the DR group without losing DR group consistency. If the EVA continued to write to the remaining LUNs and the unavailable LUN became available, its contents would not be consistent with the other LUNs of the DR group and a database recovery at this point would lose transaction integrity. Thus, the availability of a remote DR group is tied to the availability of the individual LUNs. If the remote DR group contains a Vraid0 LUN, mirroring to the entire DR group is tied to the availability of that Vraid0 LUN. While it may be desirable to use a Vraid0 LUN as the target of a remote mirror to save costs, the solution must be able to tolerate the resulting loss of availability due to its inclusion.

---

### Best practice to optimize availability

Do not use Vraid0 as a target for Continuous Access EVA mirror.

---

## Use single protection level

The disk group protection level defines the number of disk failure–autoreconstruction cycles that the array can accomplish without replacement of a failed disk. Following a disk failure, the controller re-creates the missing data from the parity information. The data is still available after the disk failure, but it is not protected from another disk failure until the reconstruction operation completes.

Single protection reserves 2x the disk size because of the reconstruction mechanism within the array. For example, in a failure scenario involving VRaid1, data from the remaining or “widowed” disk is moved to two other distinct locations during the reconstruct, resulting in a “fail safe” movement or copy of the VRaid1 data. The “widow” is then available to be paired with another disk after reconstruction completes.

For Vraid1 and Vraid5, protection level “none” provides resilience to a single disk failure. However, this is not a best practice configuration. **Vraid0 provides no protection to any disk failure. The failure of any disk in a disk group with “single” or “double” protection will render any Vraid0 vdisk in that disk group inoperative.**

Conversely, the statistical availability of disks and the typical service time to replace a failed disk (MTTR<sup>2</sup>) indicate that “double” protection level is unnecessary in disk groups of fewer than 168 disks in all but the most conservative installations. A mitigating condition would be a service time (MTTR) that exceeds seven days. Then a protection level of “double” might be considered for groups of less than 168 disks.

---

**Best practice to optimize availability**

Use “single” protection level as a minimum.

---

**Note**

Protection level reserved capacity is not associated with the occupancy alarm setting. These are independent controls.

---

**Keep a multiple of eight disks in each disk group**

Within a disk group, the EVA creates multiple subgroups of disks called enhanced data protection domains. Each enhanced data protection domain contains sufficient redundancy information to continue operation in the event of a disk failure within that enhanced data protection domain. The EVA can thus sustain multiple simultaneous disk failures while not losing user data, as long as no more than one disk per enhanced data protection domain fails. Enhanced data protection domains are created when a disk group is created, and additional sets are created as necessary when disks are added to the disk group. Enhanced data protection domains are created and optimally managed by the EVA controllers, with no user intervention required. The structure of the enhanced data protection domains is not configurable or viewable by the user. Using multiple disk groups is preferable as a user configurable mechanism for managing fault domains (see the discussion in the following section regarding multiple disk groups).

The target size of each enhanced data protection domain is eight disks, with a minimum of six and a maximum of 11. As disks are added to a disk group, the enhanced data protection domain automatically expands until it reaches 12 disks. At that point, it splits into two sets of six disks each. As more disks are added, one set increases from six to eight (the target size), and then the remaining set increases. After all disks have been added to a disk group, each enhanced data protection domain contains eight disks, with the possible exception of the last set, which contains between six and 11 disks. This is why it is a best practice to add disks in groups of eight.

---

**Best practices to optimize availability and capacity utilization**

- Keep the total number of disks in the disk group to a multiple of eight.
  - When creating a disk group, let the EVA choose which disks to place in the group.
- 

---

<sup>2</sup> Mean Time to Repair

## Utilize multiple disk groups

Although the EVA offers numerous levels of data protection and redundancy, a catastrophic failure<sup>3</sup> can result in loss of a disk group. An example would be the failure of a second disk in an enhanced data protection domain before the reconstruction operation is complete. The probability of these events is low. However, installations requiring the highest levels of data availability may require creating multiple disk groups for independent failure domains<sup>4</sup>. Multiple groups result in a slightly higher cost of ownership and potentially lower performance, but may be justified by the increased availability.

The strategy for multiple disk groups is to keep recovery data in a separate disk group from the source data. The typical use is to keep either a snapclone or mirrorclone of the database or the database log files in a separate disk group from the application. If the primary disk group fails, the recovery disk group may remain available. Additional work is required to restore the application, but the recovery data is online, thus reducing the recovery time.

---

### Note

Like snapclone, a mirrorclone can be created in a different disk group with a different raid type. A mirrorclone also has additional features such as delta-resync, instant restore, and the ability to have snapshots of the mirrorclone. Mirrorclone is the preferred option.

---

For two disk groups to prevent data loss, each disk group must contain sufficient independent information to reconstruct the entire application data set from the last backup. A practical example of this is a database that contains data files, configuration files, and log files. In this instance, placing the data files in one group and duplexing the log files and control files (duplexing or mirroring is a feature of some databases) to both the data file disk group and another group ensure that loss of a single disk group does not prevent recovering the data.

For example:

- Disk Group 1 contains data files, a copy of online redo logs, a copy of the control file, and an optional copy of archived logs (if supported by either the database or OS).
- Disk Group 2 contains a copy of online redo logs, a copy of the control file, the primary archive log directory, and an optional snapclone or mirrorclone of the data files for Disk Group 1.

If Disk Group 2 does not contain the snapclone or mirrorclone of the data files, the number of disks in Disk Group 2 should be determined by the sequential performance demand of the log workload. Typically, this results in more usable capacity than is required for the logs. In this case, choose Vraid1 for the log disks. Vraid1 offers the highest availability and, in this case, does not affect the cost of the solution.

A variation on this configuration is two separate disk groups for the log and control files, and a third for the data files. This configuration has a slightly higher cost but appeals to those looking for symmetry in the configuration. In this configuration:

- Disk Group 1 contains database data files.
- Disk Group 2 contains the database log, control file, and archived log files.
- Disk Group 3 contains a database log copy, control file copy, and archived log files copy (if supported).

---

<sup>3</sup> Defined as multiple, simultaneous failures that exceed the architectural redundancy

<sup>4</sup> A failure in one domain does not affect the availability characteristics of the other domain

Disk groups can be shared with multiple databases. It is not a best practice to create a separate disk group for each database.

Disk groups can be created for each application. By doing so disk group failures will only affect the application within that disk group.

---

**Best practice to optimize availability**

For critical database applications, consider placing data files and recovery files in separate disk groups.

---

**Best practice to optimize availability**

Consider creating a disk group for each critical application.

---

**Best practice to optimize availability**

Assign snapclones and mirrorclones to a separate disk group.

---

**Note**

Creating multiple disk groups for redundancy and then using a host application such as a volume manager to stripe data from a single application across both disk groups defeats the availability value of multiple disk groups. If, for cost or capacity reasons, multiple disk groups are not implemented, the next best practice is to store the database log, control file, and log archives in Vraid1 LUNs. Vraid1 provides greater protection to disk failures than Vraid5 or Vraid0.

---

It is important to understand that optimizing for availability sometimes is a trade off against optimizing for performance.

Creating multiple disk groups to optimize for availability is not considered an optimization for performance.

**Maintain sufficient free space**

Free space, the capacity that is not allocated to a virtual disk, is used by the EVA controller for multiple purposes. Although the array is designed to operate fully allocated, functions like snapshot, reconstruction, leveling, remote replication, and disk management either require or work more efficiently with additional free space.

Three controls manage free space in the EVA: the protection level, the capacity occupancy alarm, and the capacity reported as available for virtual disk creation. Successful capacity planning requires understanding specific requirements for availability and cost, and setting the appropriate protection level, capacity occupancy alarm, and total virtual disk capacity.

Set the protection level for the disk group. For information regarding setting the protection level of the disk group, see the previous discussion of protection level and availability.

Additional reserved free space—as managed by the occupancy alarm and the total virtual disk capacity—affects leveling, remote replication, local replication, and proactive disk management. The following best practice addresses the occupancy alarm setting and the total virtual disk capacity.

**The occupancy alarm is set for each disk group as a percentage of the raw capacity.** Base the occupancy alarm setting on the unique installation requirements for proactive disk management, remote replication, and leveling.

Proactive disk management (PDM) is a request by a customer or HP Services to ungroup a disk, or it is a predictive disk failure request by the EVA to ungroup a disk. In either case, the array migrates the contents of the disk to free space before removing the disk from use. PDM can occur only if sufficient free space is available. PDM operation capacity is independent of protection level capacity and this reserved protection level capacity is not usable for this purpose. Customers who desire the highest levels of availability elect to reserve additional free space for PDM.

The capacity used for PDM is twice the largest disk in the disk group for each PDM event anticipated. Typical choices are none, one, or two events. The greater the disk count in a disk group, the greater the opportunity for a PDM event. As previously noted in the example for protection levels involving a VRaid1 failure scenario, data is moved to two other distinct locations during the reconstruct, resulting in a “fail safe” movement or copy of the VRaid1 data.

---

**Best practice to optimize availability**

Set the occupancy alarm to allow space for one or two PDM events per disk group.

---

HP StorageWorks Continuous Access EVA uses free space for the DR group log (also known as the write history log). The DR group log is activated when the inter-site link to the remote array fails or is suspended. Until the link is reestablished, the EVA controller records changes locally in the DR group log. For free space management, allow for the maximum size of the write history logs in each disk group. The size of the write history log is specified when the DR group is created. The default size is the combined size of the DR group’s virtual disks.

Leveling and reconstruction performance require a minimum of 5 GB of free space per disk group.

---

**Best practice to optimize availability**

Set the occupancy alarm to the larger of the capacity required for PDM or the total Continuous Access EVA write history log capacity, plus 5 GB. This capacity is converted into a percentage of the raw capacity and then rounded to the next largest whole number. The pseudo-Excel formula would be (see footnotes for description of functions):

$$\text{Occupancy\_Alarm} = 100 - \text{ceiling}^5 \left[ \frac{(\max^6 (\text{PDM\_capacity}, \text{HP Continuous Access\_log\_sum}) + 5 \text{ GB})}{\text{total\_disk-group\_raw\_capacity}} \right]$$

---

**Remaining free space is managed by the creation of virtual disks,** and is measured by the capacity available to create additional virtual disks. This free-space capacity is used by space-efficient snapshots. It is critical that sufficient free-space is available for the space-efficient snapshot copies; else all snapshot copies in the disk group become inoperative<sup>7</sup>. (Fully allocated snapshot vdisks and snapclone or mirrorclone vdisks will continue to be available.)

---

<sup>5</sup> Ceiling function—compute the next largest whole number

<sup>6</sup> Choose the maximum of PDM\_capacity or HP Continuous Access\_log\_sum

<sup>7</sup> Space-efficient snapshot vdisks become inoperative individually as each attempts to allocate additional free-space. In practice the effect is that all become inoperative together.

Snapshots use copy-on-write technology. Copy-on-write occurs only when either the original virtual disk or the snapshot virtual disk is modified (a write), and then an image of the associated blocks is duplicated into free space, which is allocated as needed. For any given block, the data is copied only once. As snapshot copies diverge, the capacity available to create virtual disks decreases as this allocation occurs.

The actual capacity required for a space-efficient snapshot depends on the divergence of the original virtual disk and the snapshot. This value is unique for each application, but can range from 0 percent to 100 percent. The suggestion for the initial usage of snapshot (that is, when you do not know the actual physical capacity required by the snapshot) is to reserve (do not allocate to virtual disks) 10 percent of the capacity of the parent virtual disks times the number of snapshots per parent virtual disk. For example, if you need to create two space-efficient snapshot vdisks of a 500-GB virtual disk, you need to ensure that 100 GB (500 GB\* 10 percent\*2) of usable capacity is available. Compute the usable capacity using the RAID level selected for the snapshot vdisk.

---

**Best practice to optimize availability**

Leave unallocated virtual disk capacity, the capacity available to create virtual disks, equal to the sum of the capacities required for all space-efficient snapshot copies within a disk group.

---

---

**Best practice to optimize availability**

Respond to an occupancy alarm; evaluate what changed, replace failed disks, add disks, or reevaluate space-efficient snapshot vdisk usage (delete snapshot vdisks). Extended operation of the array in an occupancy alarm condition is not a best practice.

---

See examples for free space management in Appendix A of the “HP StorageWorks 4x00/6x00/8x00 Enterprise Virtual Array configuration best practices white paper.”

**Follow the proper procedure to replace failed disks**

Following the rules for shelf and disk organization is the best protection against potential data loss and loss of availability due to disk failure. However, when a disk fails, additional steps should be followed to minimize the risk of data loss or unavailability.

HP service engineers are trained on the proper EVA repair procedures and are alert to abnormal conditions that warrant additional steps to ensure continued operation. The best practice to maximize availability is to call for HP service. If HP service is not an option or is unavailable, use the following rules.

When a disk fails, the EVA rebuilds the failed disk data through a process known as reconstruction. Reconstruction restores the disk group resiliency to avoid disruption from another disk failure. After reconstruction or after a new disk is added to a disk group, the EVA redistributes the data proportionately and reorganizes enhanced data protection domains to the active disks.

---

**Best practice to optimize availability**

Use the following procedure for disk replacement:

Wait for the reconstruction to complete before removing the failed disks. This is signaled by an entry in the event log. If reconstruction does not complete, the EVA may internally assign the disk to a different enhanced data protection domain, and later cause data migration operations in the background, adversely impacting performance. Replacing the disk early provides no benefit, and may actually cause harm.

1. Use HP StorageWorks Command View EVA to ungroup or remove the disk. This assures the disk is not a member of a disk group.
  2. Replace the failed disk. The new disk must be inserted into the same slot as the failed disk.
  3. Ensure the disk addition policy is set to manual mode. Manually add the new disk into the original disk group.
- 

**Follow the proper procedure to add new disks**

To minimize false indications of excessive errors, insert multiple disks carefully and slowly, pausing between disks and checking each disk is visible in command view before adding another disk. This careful sequencing allows the initial bus interruption from the insertion and the disk power-on communication with the controller to occur without the potential interruption from other disks. In addition, this process sequences leveling to not start until all the new disks are ready.

Although the array supports replacing existing smaller disks with larger disks, this process is time consuming and disruptive and can result in a non-optimum configuration. Do this only if the option to build new disk groups and move existing data to the new disks is unavailable.

When replacing disks of a larger size, it is recommended that all disks of that size get added at once and to the same disk group, so that new enhanced data protection domains may be created to all use the new disks, rather than spreading the larger capacity disks across existing enhanced data protection domains.

---

**Best practice to optimize availability when adding disks to an array**

- Set the add disk option to manual.
  - Add disks one at a time, waiting at least 60 seconds between disks.
  - Distribute disks vertically and as evenly as possible to all the shelves.
  - Unless otherwise indicated, add new disks to existing disk groups using the HP StorageWorks SSSU add multiple disks command.
  - Add disks in groups of eight.
  - For growing existing applications, if the operating system supports virtual disk growth, increase virtual disk size. Otherwise, use a software volume manager to add new virtual disks to applications.
-

### **Maintain all components at current revision levels**

HP invests a considerable amount of engineering resources to continually improve the quality of the EVA product line. This investment is realized by HP customers by maintaining the component software and firmware at the latest active versions. Active versions of EVA firmware allow online firmware updates, which allow updating the EVA controller and drive firmware with zero downtime.

---

#### **Best practice to optimize for the highest levels of data availability**

Maintain array software and firmware at active revision levels.

---

---

#### **Best practice to optimize for the highest levels of data availability**

Maintain drive firmware at active revision levels.

---

### **Use the correct Fibre Channel HDD type for its intended purpose**

FATA drives are designed for lower duty cycle applications such as near online data replication for backup. These drives should not be used as a replacement for the EVA's high-performance, standard duty cycle, Fibre Channel drives. Using FATA drives for continuous duty cycle applications could shorten the life of the drive.

---

#### **Best practice to optimize for the highest levels of data availability**

Use only EVA's high-performance, standard duty cycle Fibre Channel drives for continuous duty cycle applications.

---

## Optimizations "outside the box"

### **Maintain operational environmental specifications**

To maintain optimum product operation, maintain the operational environment to meet required specifications. The ambient temperature (the enclosure air intake or room temperature) is especially critical. Providing separate power feeds/circuits into the EVA as well as the usage of an Uninterruptible Power Source (UPS) can help to minimize downtime due to unexpected power fluctuations.

---

#### **Best practice to optimize availability**

Maintain a clean and cool operating environment.

---

---

#### **Best practice to optimize availability**

Provide separate power feeds into the EVA.

---

---

#### **Best practice to optimize availability**

Utilize a UPS.

---

## **Implement pro-active remote fault monitoring (ISEE/HPCC)**

The best protection from downtime is to avoid the failure. Good planning and early warning of problems can minimize or eliminate many issues. For the EVA, Instant Support Enterprise Edition (ISEE) is a free service that forwards EVA faults and warnings directly to HP Services through a secure virtual private network. HP can evaluate and diagnose problems remotely, possibly avoiding more serious issues. If an array requires service, ISEE greatly increases the probability that the HP Service engineer arrives on-site with the correct replacement parts to minimize the time to repair. HP Configuration Collector (HPCC) is a component of ISEE that enables the capture of information about the EVA configuration. Capturing this information provides HP with the opportunity to address problems before a system outage occurs and to more quickly diagnose and resolve complex problems.

If site security policies exclude direct communication to HP Services, ISEE can be set up to report warnings to a local customer contact. If ISEE is not used, a custom automated alert process based on HP Web-based Enterprise Services (WEBES) or similar tool can be developed. It is critical that alarms from the EVA are not ignored.

---

### **Best practice to optimize availability**

Install and use ISEE/HPCC or equivalent tools to monitor and alert administrators to changes in EVA health status.

---

## **Run database integrity checks**

Exchange, Oracle®, and other databases include tools to verify the integrity of the database. These tools check the database for consistency between records and data structures. Use these tools as part of the ongoing data center processes, as you would data backup and recovery testing. Proper use of these tools can detect database inconsistency early, when repair or recovery options are still available.

---

### **Best practice to optimize availability**

Run database integrity checks as an ongoing data center process.

---

## **Maintain proven backup and recovery processes**

Independent of the storage device, include a proven backup and recovery process in array and data center management procedures. The EVA is supported by numerous backup applications. HP Data Protector and similar third-party backup software are supported on a variety of popular operating systems. They support the EVA directly or are integrated through Oracle, Microsoft® SQL, or other databases to provide zero-downtime backup.

Along with the backup data, save a copy of the EVA configuration files. An exact description of the array configuration greatly reduces recovery time should the need arise. These should be stored on media not associated with the array.

Do not consider array-based copies the only form of backup. In the rare case that a catastrophic failure occurs, non-array-based copies of the data may be the only restore option. Snapshot, snapclone, and mirrorclone copies complement a backup strategy that includes full copies to offline or near-line storage. In this application, online backups such as snapshots, snapclones, and mirrorclones can provide alternatives for reducing recovery time by providing the first option for recovery.

Perform regular backups and be sure to test the restore process twice a year. The EVA greatly simplifies the testing process by providing a simple process to create and delete disk groups or virtual disks. Capacity used for testing can be easily reused for other purposes.

---

**Best practice to maintain data recovery**

Perform regular data backup and biannual recovery tests.

---

**Best practice to maintain data recovery**

Include a copy of the EVA configuration with the backup data. This can be accomplished with the SSSU Capture Configuration utility.

---

**Implement fully redundant fabric configuration**

For mission-critical applications, HP recommends that you implement a level 4, fully redundant fabric configuration. You can justify the additional cost if you consider the cost of losing access to critical data. Refer to the HP StorageWorks SAN design reference guide for details.

---

**Best practice to optimize for the highest levels of data availability**

Implement a level 4 fully redundant fabric configuration.

---

**Implement remote mirroring—HP StorageWorks Continuous Access EVA**

HP StorageWorks Continuous Access EVA is an optional feature of the array that enables real-time remote data replication. Continuous Access EVA protects against catastrophic EVA or site failures by keeping simultaneous copies of selected LUNs at local and remote EVA sites. This feature can work standalone or in combination with system clustering software.

---

**Best practice to optimize for the highest levels of data availability**

Consider array-based Continuous Access EVA or operating system-based replication software to provide real-time mirroring to a second EVA.

---

**Implement disaster recovery solutions**

Many operating systems support optional fault-tolerant configurations. In these configurations, multiple servers and arrays are grouped together with appropriate software to enable continued application operation in the event of system component failure. These solutions span the replication continuum from simple local mirroring to complex disaster recovery solutions.

Continuous Access EVA can be used for simple mirroring to increase data availability or as a component of a complete disaster recovery solution. A complete disaster recovery solution automatically (or manually) transfers the applications to the remaining functional systems in the DR group in the event of a site or component failure.

The HP StorageWorks EVA is supported in many disaster recovery solutions. HP provides two disaster recovery products: Cluster Extensions for Windows and Linux, and Metrocluster for HP-UX. For more details, contact HP or your operating system vendor. For applications requiring the highest availability, these system-level solutions should be considered.

---

**Best practice to optimize application availability**

Consider disaster recovery solutions to provide continuous application operation.

---

**Deploy the highest service level**

By deploying the highest level of service offering, unplanned downtime can be significantly reduced. HP provides a comprehensive Mission Critical Support portfolio that addresses the full range of requirements.

---

**Best practice to optimize for the highest levels of data availability**

Deploy the highest level of service.

---

**Sign up for Subscriber's Choice**

To avoid downtime all customers, regardless of service level, can get proactive notification of array advisories by signing up for Subscriber's Choice. Sign up for Subscriber's Choice at:

<http://www.hp.com/go/myadvisory>

---

**Best practice to optimize for the highest levels of data availability**

To receive the latest advisories, sign up for Subscriber's Choice.

---

# Glossary

<b>data availability</b>	The ability to have access to data.
<b>data protection</b>	The ability to protect the data from loss or corruption.
<b>disk group</b>	A collection of disks within the array. Virtual disks (LUNs) are created from a single disk group. Data from a single virtual disk is striped across all disks in the disk group.
<b>free space</b>	Capacity within a disk group not allocated to a LUN.
<b>leveling</b>	The process of redistributing data to existing disks. Adding, removing, or reconstruction initiates leveling.
<b>LUN</b>	Logical unit number. An addressable storage collection. Also known as a virtual disk (Vdisk).
<b>occupancy</b>	The ratio of the used physical capacity to the total available physical capacity of a disk group.
<b>physical space</b>	The total raw capacity of the number of disks installed in the EVA. This capacity includes protected space and spare capacity (usable capacity).
<b>protection level</b>	The protection level defines the reserved space used to rebuild the data after a disk failure. A protection level of none, single, or double is assigned for each disk group at the time the disk group is created.
<b>protection space</b>	The capacity that is reserved based on the protection level.
<b>reconstruction, rebuild, sparing</b>	Terms used to describe the process of recreating the data on a failed disk. The data is recreated on spare disk space.
<b>reserved space</b>	Same as protected space.
<b>enhanced data protection domain</b>	A group of disks within a disk group that contains a complete set of parity information.
<b>usable capacity</b>	The capacity that is usable for customer data under normal operation.
<b>virtual disk</b>	A LUN. The logical entity created from a disk group and made available to the server and application.
<b>workload</b>	The characteristics of the host I/Os presented to the array. Described by transfer size, read/write ratios, randomness, arrival rate, and other metrics.

## For more information

This paper provides technical details to optimize the HP StorageWorks Enterprise Virtual Array (EVA) for availability. It is not intended to be a general-purpose tutorial on EVA operation. HP and its partners provide additional information and services to help you optimize your EVA. For more information on the HP StorageWorks EVA, go to <http://www.hp.com> or contact your local HP sales representative.

© 2007 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft is a U.S. registered trademark of Microsoft Corporation. Oracle is a registered US trademark of Oracle Corporation, Redwood City, California.

4AA1-4202ENW, July 2007

