COMPAQ
White Paper

July 1999
Project No. 0292-0799-A

Prepared by Storage Products Division

Compaq Computer Corporation

Contents

Smart Array Controller: Online Spare Recommendations

Abstract: *This guide provides customers with the insight to the advantages of Online Spares for the Compaq Smart Array Controllers to further improve system fault tolerance.*

Fault tolerant options for the Smart Array controllers play a vital role in reducing the possibility of data loss after an inevitable drive failure. One Compaq Smart Array feature, the utilization of Online Spares, improves the overall system fault tolerance by automatically rebuilding data lost on the failed drive and restoring the system to full RAID fault tolerance protection.

The purpose of this guide is to communicate the advantages of using Online Spares with the Compaq Smart Array Controller Family so that proper up-front planning can significantly reduce the risk of data loss.

Notice

Smart Array Controller: Online Spare Recommendations
White Paper prepared by Storage Products Division

First Edition (July 1999)
Document Number Project No. 0292-0799-A

Factors that Can Impact Fault Tolerance

With the introduction of new Smart Array Controllers, Compaq continues with its tradition of providing reliable arrays for high end workgroup servers' and low end departmental servers' storage needs. For greater fault tolerance, there are several conditions that can impact data availability including:

- **Number and logical grouping of drives into "logical arrays"**[1]– The greater the number of drives in a logical array, the higher the likelihood that one or more of those drives will fail (determined by mathematical probability).
- **RAID level selection -** RAID provides data redundancy to increase system reliability and performance. The user selection of one configuration method over another may impact data availability, performance and/or cost—and the potential for data loss.
- In addition, the option of utilizing Online Spares (hot spare) not only automatically reduces maintenance time for replacing failed drives, but can also reduce the risk of data loss.

Fault tolerant issues

In configuring fault tolerance, the method customers select can affect the amount of available disk storage capacity and performance of the drive array. The purpose of fault tolerance is to reduce the possibility of having to recover from data loss. A system can continue to function after a disk failure, and the failed drive can be replaced when it is convenient. In the event of a drive failure, the condition of the logical drive varies depending upon the fault tolerance method used. Therefore, it is advantageous to configure to the recommended method.

Table 1 illustrates the risk factor involved using different fault tolerant solutions.

Table 1. Risk matrix for array reliability using various fault tolerant methods

Number of drives	RAID 1	RAID 1 w/ Online Spare	RAID 4 & 5	RAID 4 & 5 w/ Online Spare
22-56	Low	Low	High	Moderate to High
15-22	Low	Low	Moderate	Low to Moderate
<14	Low	Low	Low	Low

Fault Tolerant Configuration Example

The following is an example of how an Online Spare can impact fault tolerance:

FACTS: The Smart Array 3200 Controller is designed to support up to 28 drives, with a recommendation of up to 14 drives in each logical array. Fault tolerant configurations supported by the Smart Array 3200 controller include RAID 0, RAID 1 or 0+1, RAID 4, and RAID 5.

[1] "Logical Array" is not a term typically found in Compaq documentation. Logical Array refers to a group of physical drives that have been grouped together (logically) into a single array set while configuring the array.

EXAMPLE SITUATION: Company A has a server deployed two years ago sitting in a remote workgroup office that serves 50 people. The server has (5) 4.3GB hard drives in a RAID 5 configuration. Office hours are Monday – Friday, 8 a.m. – 6 p.m.

Three days previously, the office had a lightening strike that brought down all the servers for a day. There were difficulties getting the server to get back up. That Friday, there was a big drive failure at 11 p.m. after all office workers had already gone home for the weekend.

EVALUATION: Using the information from "Array Reliability", the following can be determined about the possible risk factor involved in fault tolerance for Company A's scenario:

- If there was a failure at 11 p.m. Friday, right after the previous power failure from which there was difficulty in recovering, chances are that before the weekend is over, another failure could occur and cause data to be lost.
- Because no one is in the office at Friday at 11 p.m., the problem would not discovered until Monday morning. Technicians would have to be called and valuable server uptime would be consumed while the technicians work on the server. With no redundant drives to take over in case of failure, the chances for data loss would also be increased. In case of multiple drive failure, the storage system would become inoperative.

Compaq recommendation

Compaq recommends a lower risk configuration and Online Spares to be used in all RAID configurations to lower the risk of data loss.[2]

Recommended Configuration

There are three possible recommendations that could significantly lower the risk of a second drive failure related to the number of drives that Company A has purchased for this server:

1. **Online Spares** – At least one drive could be assigned as an Online Spare. Having an Online Spare is critical to reducing the risk of multiple drive failures. In the event a second drive failure that occurred close in time to the first drive failure, having an Online Spare could mean the difference between data availability and data loss.

   If Company A had installed an Online Spare, then the controller would rebuild the data that was on the failed drive to the Online Spare. The controller would send data that it would normally store on the failed drive directly to the Online Spare, restoring fault tolerance.

2. **Configure the array to RAID 1:** Because RAID 1 is built on "mirrored pairs" of drives, multiple physical drives (of the 48 disk drives) can fail without losing user data, as long as both drives in any one "mirrored pair" do not fail. One of the drives in each pair must survive until it can be rebuilt to restore fault tolerance.

3. **Configure smaller logical drives in RAID 5:** Compaq recommends no more than 14 physical disk drives for each logical drive using RAID 5. If a larger logical drive is needed for higher capacity under a single logical volume for special application needs, you can use Compaq StorageWorks Virtual Replicator software (optional purchase) to tie multiple logical drives together into a storage pool and present a single virtual volume, up to 1 terabyte of capacity, to the operating system. By decreasing the number of drives per array, multiple failures become less likely.

[2] Because many other factors outside the control of Compaq can impact data availability, no recommendation, including the Compaq recommendation, guarantees that data will be maintained.

For more information on the StorageWorks Virtual Replicator software, please visit our web site at www.Compaq.com/StorageWorks. For more information on RAID configurations, refer to the technical white paper, Smart Array Controller Technology: RAID 5 Configuration Recommendations and Smart Array Controller Technology: RAID 1 Configuration Recommendations

Other Recommendations

Rebuild Priority Setting: Inside the Array Configuration Utility (ACU), select the Rebuild Priority setting and set to "High." This will insure that as soon as a drive fails, rebuilding the failed drive onto an Online Spare or replaced disk drive can take place at the fastest possible rate.

Conclusion

Proper array configuration is critical to maintaining needed availability, performance, and capacity. Recognizing the advantages and disadvantages of specific RAID configurations and using the Compaq recommended configurations could minimize unplanned and costly downtime. Limiting the number of drives per logical array, assigning Online Spares, and using RAID 1 for the lowest risk array configuration are all recommended for higher reliability.

Appendix A

Definitions:

Online Spare - An Online Spare is a drive that has been identified in the configuration as a "stand-by" drive in the event of a drive failure in a logical array. When the failure occurs, the Online Spare can be automatically activated to begin an immediate rebuild of the logical array, thus reducing the risk of data loss.

RAID 5 - (a.k.a., Distributed Data Guarding), stores parity data across all the drives in the array. If a drive fails, the controller uses the parity data and the data on the remaining drives to reconstruct data from the failed drive, allowing the system to continue operating with a slightly reduced performance until the failed drive is replaced. RAID 5 is the most cost effective of the fault tolerant RAID levels, but is slower than RAID 1 implementations.

RAID 4 [3]- (a.k.a., Data Guarding), stores parity in a single dedicated drive. If a disk drive fails, the controller uses the data on the parity drive and the data on the remaining drives to reconstruct data from the failed drive. This allows the system to continue operating with slightly reduced performance until you replace the failed drive. RAID 4 is as cost effective as RAID 5, but is slower than RAID 1 implementations.

RAID 1, 0+1 & 10 - (a.k.a., Data Mirroring), is the highest performance and highest fault tolerant RAID method. Drive mirroring creates fault tolerance by storing two sets of duplicate data on a pair of disk drives. Beyond the use of just two drives, RAID 1 is the most expensive fault tolerance RAID method, because 50 percent of the drive capacity is used to store the redundant data. To improve performance in configurations with two or more drive pairs, the data is striped across the drives and then mirrored. This is referred to as RAID 0 + 1 or RAID 10.

RAID 0 - (a.k.a., Striping), does not provide fault tolerance. This level of RAID stripes data across all of the drives of the array, but does not incorporate a method to generate redundant data. If one physical drive fails, data loss will occur for the entire logical drive. Because none of the capacity of the logical drive is used for redundant data, RAID 0 offers the best processing speed and price/capacity ratio—but no fault tolerance or data protection.

[3] Compaq intended to discontinue RAID 4 support for Smart Array Controllers released after July 1999.