

July 1999  
ECG Project No.: 0291-0799-A

Prepared by Storage Products  
Division

Compaq Computer Corporation

**Contents**

**Factors that Can Impact**

**Fault Tolerance.....3**

**RAID 5 Configuration Issues.....3**

Configuration Examples .....4

**Compaq Recommendation.....5**

Other Recommendations .....5

**Conclusion .....5**

**Appendix A.....6**

# Smart Array Controller Technology: RAID 5 Configuration Recommendations

*Abstract: This paper provides customers with the insight to Compaq's recommendations for the most advantageous RAID 5 fault tolerant configurations to ensure maximum uptime.*

Fault tolerant configurations for the Smart Array controllers play a vital role in reducing the possibility of data loss after an inevitable drive failure. Because not all RAID levels are created equal, there are both recommended and non-recommended configuration methods and a range of risks associated with each possible configuration.

The purpose of this guide is to communicate the advantages and disadvantages of several RAID configurations so that the risk of data loss can be significantly reduced by proper up-front planning and configuration design.

## Notice

The information in this publication is subject to change without notice and is provided "AS IS" WITHOUT WARRANTY OF ANY KIND. THE ENTIRE RISK ARISING OUT OF THE USE OF THIS INFORMATION REMAINS WITH RECIPIENT. IN NO EVENT SHALL COMPAQ BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, SPECIAL, PUNITIVE OR OTHER DAMAGES WHATSOEVER (INCLUDING WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION OR LOSS OF BUSINESS INFORMATION), EVEN IF COMPAQ HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

The limited warranties for Compaq products are exclusively set forth in the documentation accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

This publication does not constitute an endorsement of the product or products that were tested. The configuration or configurations tested or described may or may not be the only available solution. This test is not a determination of product quality or correctness, nor does it ensure compliance with any federal state or local requirements.

Product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

Compaq, Contura, Deskpro, Fastart, Compaq Insight Manager, LTE, PageMarq, Systempro, Systempro/LT, ProLiant, TwinTray, ROMPaq, LicensePaq, QVision, SLT, ProLinea, SmartStart, NetFlex, DirectPlus, QuickFind, RemotePaq, BackPaq, TechPaq, SpeedPaq, QuickBack, PaqFax, Presario, SilentCool, CompaqCare (design), Aero, SmartStation, MiniStation, and PaqRap, registered United States Patent and Trademark Office.

Netelligent, Armada, Cruiser, Concerto, QuickChoice, ProSignia, Systempro/XL, Net1, LTE Elite, Vocalyst, PageMate, SoftPaq, FirstPaq, SolutionPaq, EasyPoint, EZ Help, MaxLight, MultiLock, QuickBlank, QuickLock, UltraView, Innovate logo, Wonder Tools logo in black/white and color, and Compaq PC Card Solution logo are trademarks and/or service marks of Compaq Computer Corporation.

Microsoft, Windows, Windows NT, Windows NT Server and Workstation, Microsoft SQL Server for Windows NT are trademarks and/or registered trademarks of Microsoft Corporation.

NetWare and Novell are registered trademarks and intraNetWare, NDS, and Novell Directory Services are trademarks of Novell, Inc.

Pentium is a registered trademark of Intel Corporation.

Copyright ©1999 Compaq Computer Corporation. All rights reserved. Printed in the U.S.A.

Smart Array Controller Technology: RAID 5 Configuration Recommendations  
White Paper prepared by Storage Products Division

First Edition (July 1999)

Document Number ECG Project No.: 0291-0799-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. When configuring an array with these hardware RAID controllers, there are several considerations that may have an impact on data availability:

- **Number and logical grouping of drives into “logical arrays”<sup>1</sup>** – 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.
- **Online Spares** – An Online Spare drive (hot spare) not only automatically reduces maintenance time for replacing failed drives, but can also reduce the risk of data loss.

## RAID 5 Configuration Issues

In choosing a RAID level, four trade-off factors must be considered: Performance, Availability, Capacity, and Budget. While one or two of these factors may be more important to any given user, each factor plays a role in the future of your data. Your RAID level choice is one of the primary factors that impacts each of these four trade-off factors.

Because the need for larger and larger logical volumes is a growing necessity, the temptation to create a single logical RAID 5 volume is often great. RAID 5 provides the best price/capacity ratio, but may be a high risk to your data if you have too many physical drives in a single logical array (hereafter referred to as a “logical drive”). The following explains why:

A RAID 5 logical array takes a number of physical disk drives and groups these drives together into a single “logical drive.” To the server operating system, all of the physical drives look and act like a single drive. Because these drives are now “logically” a single drive which contains X number of disk drives, any single physical drive in the logical drive that fails can cause the entire logical drive to be at risk for data loss. Multiple physical disk drive failures will cause the entire logical drive to fail and data will be lost. The probability of multiple physical drive failures occurring is proportional to the number of physical drives in the logical drive.

---

<sup>1</sup> “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.

The chart below details the estimated risk factors involved associated with a given number of disk drives and any given RAID level.

**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

## Configuration Examples

The following is an example of how different configurations may impact fault tolerance:

**FACTS:** The Smart Array 4200 Controller is designed to support up to 56 drives, with a recommendation of up to 14 drives in each logical array. Fault tolerance methods supported by the Smart Array 4200 controller include RAID 0, RAID 1 or 0+1, RAID 4, and RAID 5.

**EXAMPLE SITUATION:** Company A has configured the following array:

- Smart Array 4200 Controller, which has four external ports, each of which is connected to one ProLiant Storage System Model UE and has twelve 9.1GB disk drives in each of the four cabinets—for a total of 48 total disk drives
- RAID level 5 configuration in a single logical array

**EVALUATION:** Using the information from the “Risk Matrix for Array Reliability” chart above, the following can be determined about the possible risk factors involved in RAID configurations for the Company A scenario:

- Forty-eight drives grouped together as one RAID 5 array is a high risk candidate when all 48 physical drives are configured as a single logical drive. This configuration increases the probability that 2 or more disk drives may fail before a failed drive can be replaced and the RAID 5 fault tolerance can be restored.
- By using more than the recommended 14 drives in a single logical drive, Company A has potentially increased the risk of multiple sequential drive failures—and subsequent data loss.
- Furthermore, the lack of an Online Spare in this configuration can further increase the risk that a second disk drive may fail before the failed drive can be replaced and the RAID 5 fault tolerance can be restored.

## Compaq Recommendation

Compaq recommends a lower risk configuration for Company A.<sup>2</sup>

### Recommended Configuration

There are two possible configuration 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. **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.
2. **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. (For more information on the StorageWorks Virtual Replicator software, please visit our web site at: [www.Compaq.com/StorageWorks](http://www.Compaq.com/StorageWorks).)

### Other Recommendations

**Online Spares:** 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. For more information on Online Spares, refer to the technical white paper, *Smart Array Controller Technology: Online Spares 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 can 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.

---

<sup>2</sup> 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.

## 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<sup>3</sup> - (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.

---

<sup>3</sup> Compaq intends to discontinue RAID 4 support for Smart Array Controllers released after July 1999.