

October 2001
155A-1001A-WWEN

Prepared by Industry Standard
Storage Group

Compaq Computer Corporation

Contents

Overview	3
Defining RAID levels	3
Evaluating RAID levels.....	3
Choosing a RAID level.....	4
Assessing performance testing	5
Logical Drive Failures.....	7
RAID ADG Features and Benefits	8
When should customers use RAID ADG?.....	11
Summary	11

Assessing RAID ADG vs. RAID 5 vs. RAID 1+0

***Abstract:** This paper will serve as a guide in pointing out features, advantages, and trade-offs between RAID ADG, RAID 5, and RAID 1+0.*

Because not all RAID configurations are created equal, there are both advantages and a range of risks associated with each possible configuration. The purpose of this paper is to communicate the advantages and risks of RAID ADG, RAID 5, and RAID 1+0 so that the threat of data loss can be significantly reduced by proper up-front planning and configuration design.

Notice

Compaq and the Compaq logo, and SmartStart are trademarks of Compaq Information Technologies Group, L.P. in the United States and other countries. Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation. Intel and Pentium are trademarks of Intel Corporation. All other product names mentioned herein may be trademarks of their respective companies.

Compaq shall not be liable for technical or editorial errors or omissions contained herein. The information in this document is provided “as is” without warranty of any kind and is subject to change without notice. The warranties for Compaq products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

Copyright ©2001 Compaq. All rights reserved. Printed in the U.S.A.

First Edition (October 2001)

Document Number

Overview

With a growing need for larger storage capacity, customers are demanding reliable ways to protect large volumes of data stored across an increasing number of disk drives. RAID technology allows a group of disk drives to be “tied” together to act like a single logical disk drive from the operating system perspective, providing increased performance and fault tolerance. Compaq continues with its tradition of providing reliable arrays for high-end workgroup servers’ and low-end departmental servers’ storage needs. Currently, Compaq offers very reliable RAID systems with the family of [Smart Array Controllers](#) that support RAID ADG, RAID 5, and RAID 1+0.

Defining RAID levels

RAID ADG, (a.k.a. Advanced Data Guarding), a Compaq innovation, is the newest RAID level providing unparalleled fault tolerance, greater than RAID 1+0 or RAID5 and at a lower cost than RAID 1+0. RAID ADG creates additional sets of parity striped data across the disks in order to withstand multiple disk failures without data loss. Designed for support on the [Smart Array 5300 Controller](#), RAID ADG is ideal for applications requiring large volumes. With the Smart Array 5300 Controller, RAID ADG can safely protect an array up to 56 total drives, ensuring customers of a powerful solution with high fault tolerance.

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 it is slower than the RAID 1+0 configuration.

RAID 1+0, (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+0 is the most expensive fault tolerant 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. RAID 1+0 is sometimes referred to as simple RAID 1.

Evaluating RAID levels

Customers who wish to create large logical drives with a high number of disk drives or high capacity disk drives need to consider the limitation of current RAID schemes. With RAID 1+0, it is expensive to create large volumes based upon the consumption of disk drives for mirroring. With RAID 5, it is not recommended to configure a logical drive array with greater than 14 drives due to risk of data loss.

In addition, today’s larger disk drive capacities may leave an array exposed to a considerable amount of risk following a single disk drive failure due to the time required to rebuild the failed drive. A second drive failure in the same mirrored pair in RAID 1+0 and any second drive failure in RAID 5 would cause an array to fail, causing data loss and downtime. With this in mind, larger capacity disk drives present new risks due to longer drive rebuild times. RAID ADG reduces the risk of an array failure and requires far less storage capacity overhead than RAID 1+0.

The following table illustrates the current RAID levels used by customers, their purposes, and limitations:

Table 1. RAID levels and functions.

RAID LEVEL	PURPOSE	LIMITATION
RAID 1+0 ¹	Mirroring: Identical data stored on multiple drives, high fault tolerance, and improved performance.	Requires 50% of capacity to be dedicated to fault protection.
RAID 5	Distributed Data Guarding: Parity data is distributed across all drives. Protects against the failure of any one drive in an array. Provides improved performance at a minimum cost.	Limited to 14 drive volumes and provides for continuous availability in the event of the failure of one drive only.
RAID ADG	Advanced Data Guarding: Two sets of parity data distributed across all drives. Provides for continuous availability in the event of the simultaneous failure of any two drives in an array. Provides high fault tolerance at a minimum implementation cost.	Lower performance than other RAID levels.

Choosing a RAID level

Customers should consider a variety of factors when choosing which RAID level best suited for their needs. These factors include:

1. Capacity utilization and budget;
2. Availability and Uptime requirements; and
3. Performance requirements.

Table 2 can help determine which RAID level best suits particular requirements.

Table 2. Choosing a RAID method.

MOST IMPORTANT	SECONDARY IMPORTANCE	RAID LEVEL CHOICE
Cost Effectiveness (cost per usable capacity)	Fault Tolerance →	RAID ADG
	Performance →	RAID 5 (RAID 0 if fault tolerance is not needed)
Fault Tolerance	Cost Effectiveness →	RAID ADG
	Performance →	RAID 1+0
Performance	Cost Effectiveness →	RAID 5 (RAID 0 if fault tolerance is not needed)
	Fault Tolerance →	RAID 1+0

¹ In the past, RAID 1+0 has sometimes been referred to as RAID 0+1 by Compaq.

Assessing performance testing

The following graphs illustrate the online transaction processing, and read/write performance comparison between RAID 5 and RAID ADG utilizing a Compaq ProLiant 8000 server with Microsoft Windows 2000:

Figure 1. RAID 5 vs. RAID ADG – Online Transaction Processing




Figure 2. RAID 5 vs. RAID ADG – 64K Sequential Read

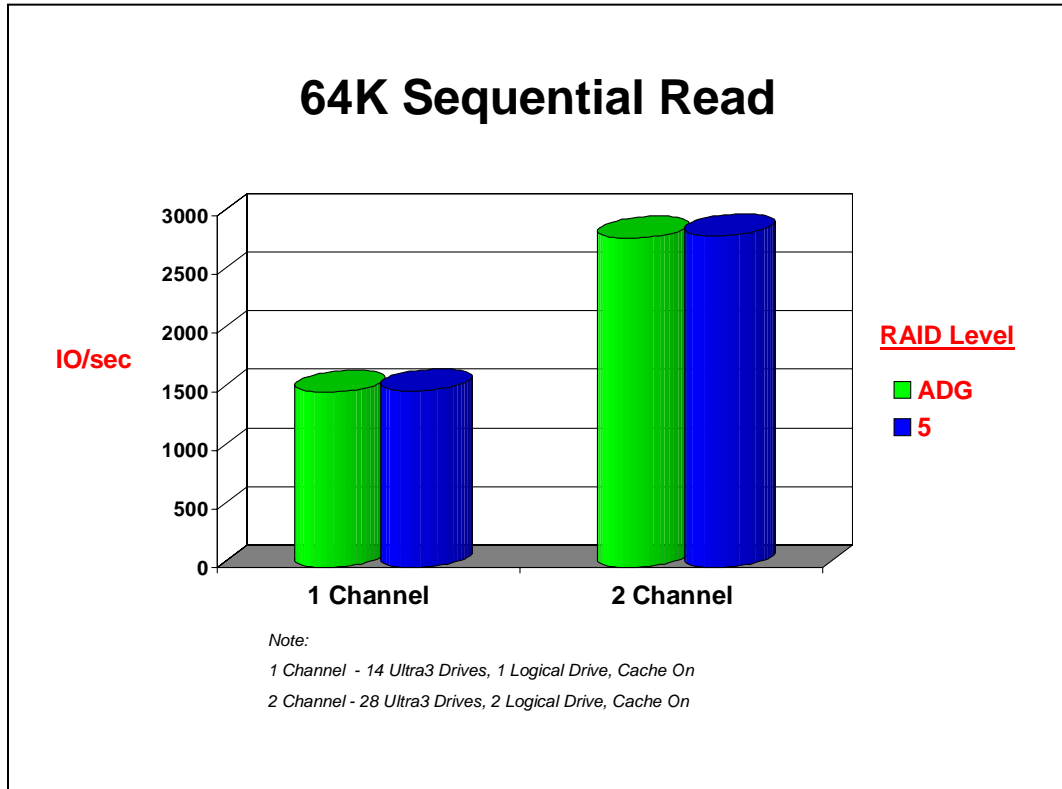


Figure 3. RAID 5 vs. RAID ADG – 64K Sequential Write

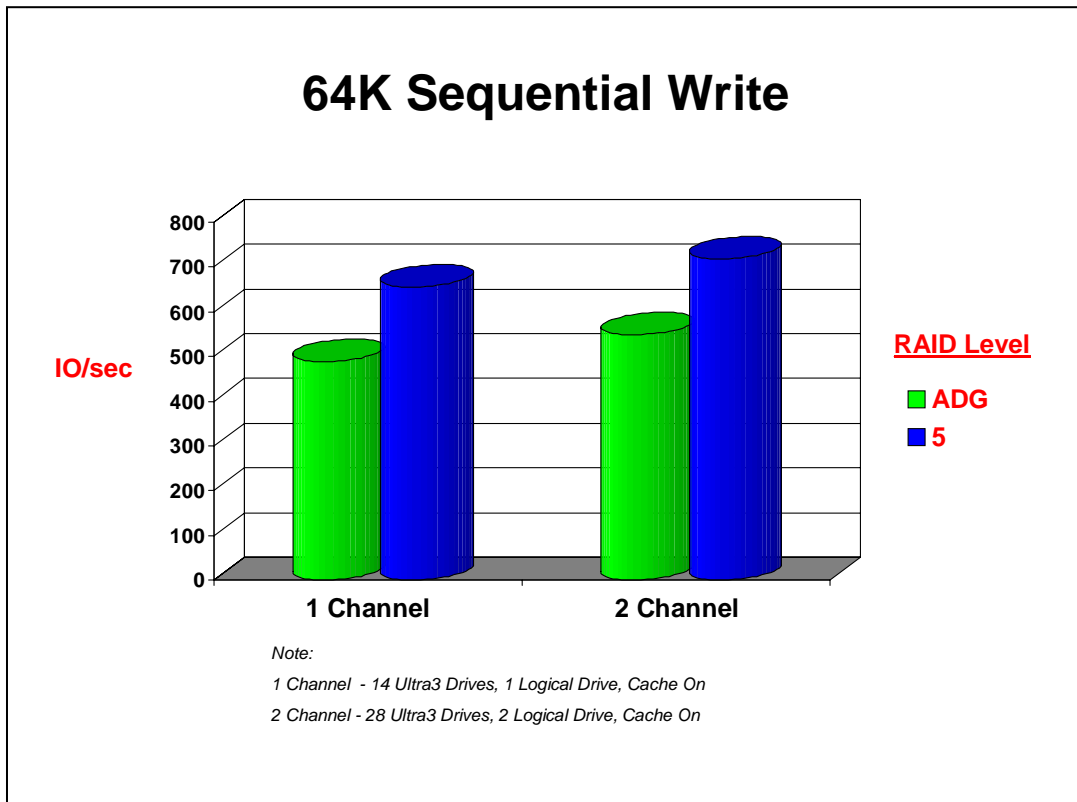


Table 3. Test Configuration

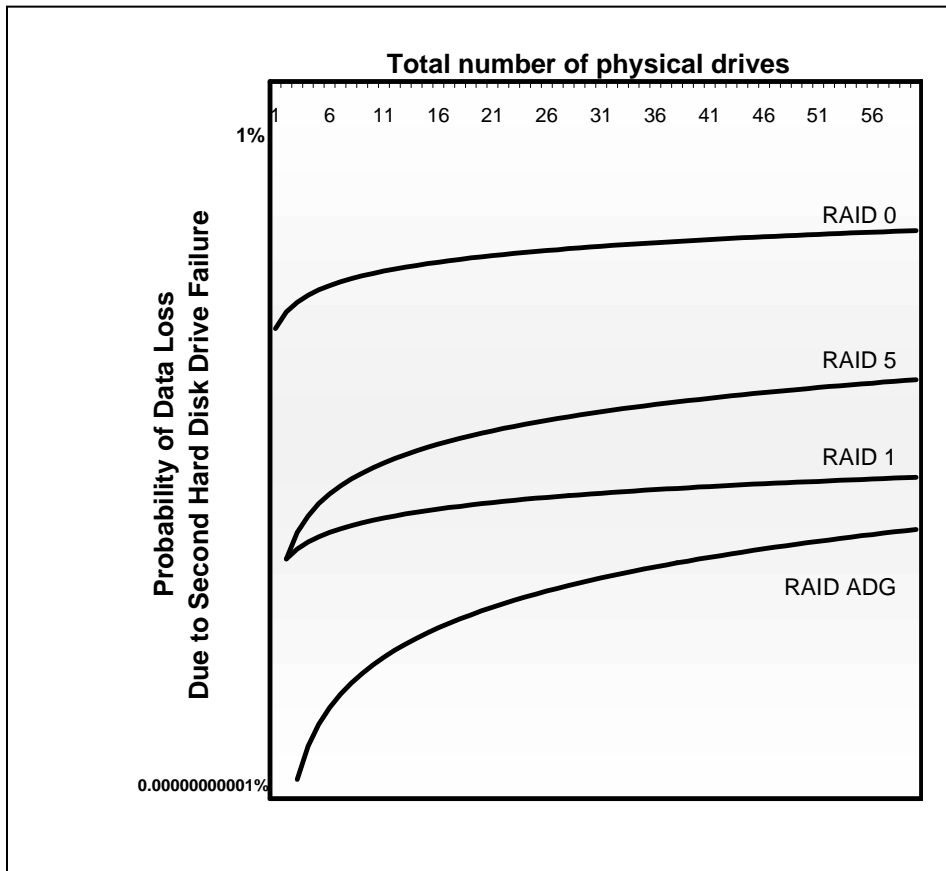
Testing Configuration	
Host System	Compaq ProLiant 8000
Host Operating System	Microsoft Windows 2000
Host CPU	8 x 550MHz Pentium Xeon with 1MB Cache
Host Memory	1024MB
Storage Configuration	(up to 2) Compaq StorageWorks Enclosure 4300 Family, (up to) 28 x 9.1GB Ultra3, 10K HDDs
Controller, Cache Size, Firmware	Compaq Smart Array Controller 5304, 128MB Cache On, Firmware v 1.62

Logical Drive Failures

In the figure below, RAID ADG shows a greater fault tolerance when compared to RAID 5 or RAID 1+0. Figure 1 shows the probability of logical drive failure for various RAID levels and different physical drive counts:

- With RAID 1+0, the maximum number of hard drives that can fail without failure of the logical drive is $n/2$. However, a RAID 1+0 logical drive will fail if only two hard drives fail, if they are mirrored to each other.
- With RAID 5, the logical drive will fail if two physical drives fail.
- With RAID ADG, three hard drives must fail before data loss is incurred.

Figure 4. Fault tolerance comparison



RAID ADG Features and Benefits

RAID ADG is advantageous for customers who require a fault tolerance scheme that would provide higher reliability than RAID 5 at a lower cost than RAID 1+0. RAID ADG benefits include:

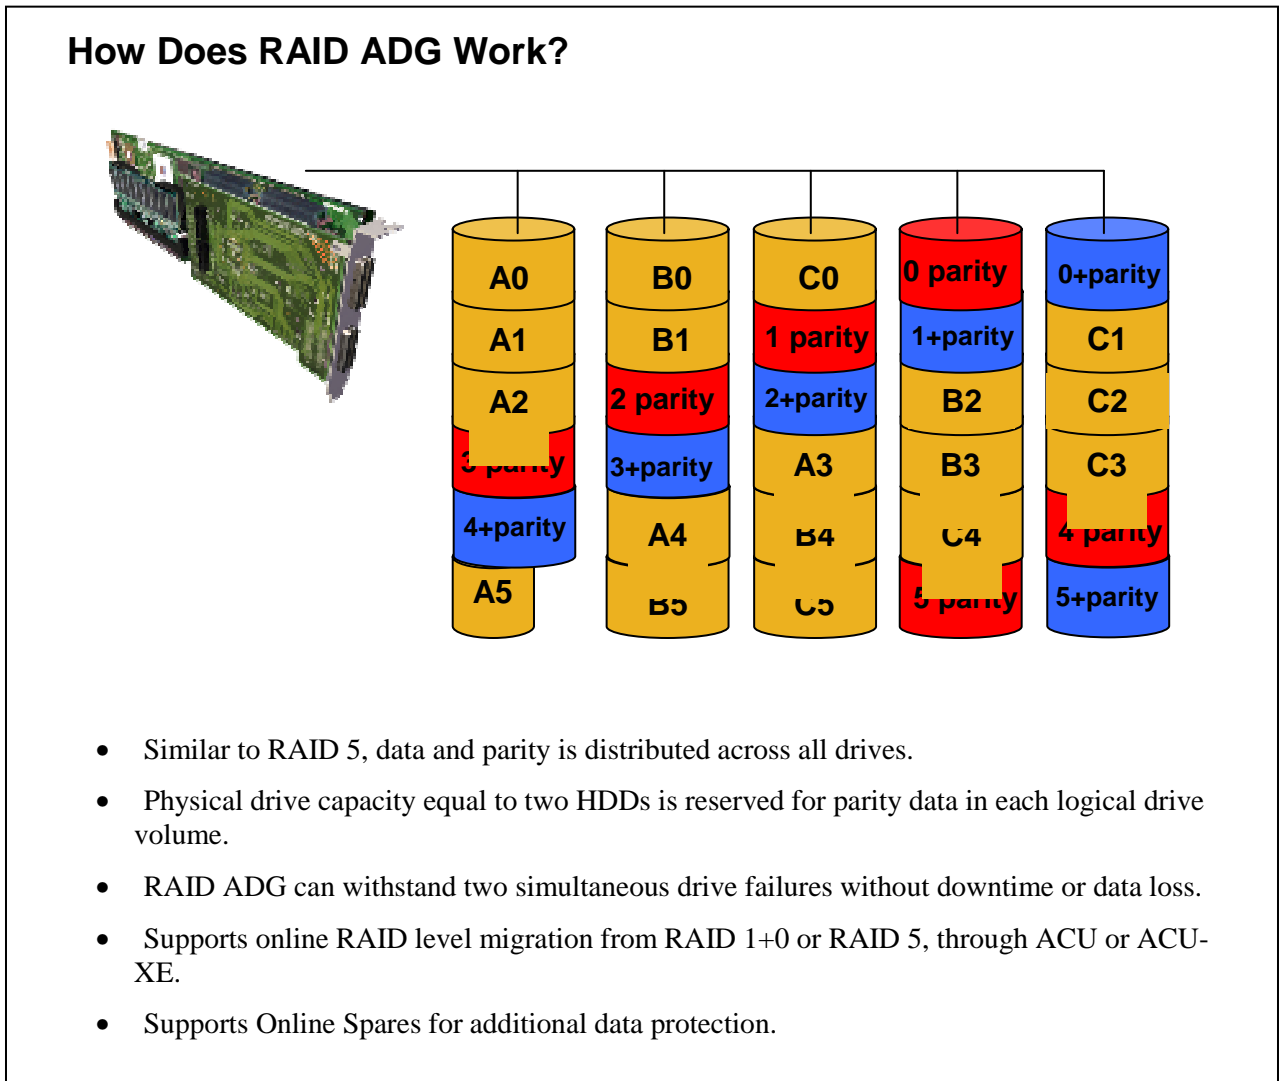
- RAID ADG can tolerate multiple simultaneous drive failures without downtime or data loss.²
- Ideal for applications requiring large logical drives
- Can safely protect an array up to 56 total drives
- Greater fault tolerance than RAID 1+0 or RAID 5
- Lower implementation cost than RAID 1+0
- Supports Online Spare Drive
- Supports Online RAID Level Migration from RAID 1+0 or RAID 5

² A RAID ADG Enabler Module must be attached to the controller to enable RAID ADG configuration via Array Configuration Utility (ACU-XE). The module is shipped with the 128MB cache board, a minimum of 64MB cache is required.

Greater fault tolerance with RAID ADG



With growing numbers of individual disk drives needed in a single logical volume on a single controller, RAID ADG provides higher data reliability than previously available from a RAID storage provider. As illustrated below, RAID ADG has a superior fault tolerance when compared to RAID 5 or RAID 1+0.

Figure 5. RAID ADG can withstand two simultaneous drive failures without downtime or data loss.



RAID ADG also improves density when compared to RAID 1+0 as illustrated in the figure below.

Figure 6. Density comparison with RAID 1+0 vs. RAID ADG

<h3>RAID 1+0 Solution</h3>  <ul style="list-style-type: none">▪ 504 Usable Capacity▪ RAID 1+0 = 28 drives (mirrored)▪ 6U Form Factor▪ Rough deployment cost \$40K 2 Enclosures, 28 36GB drives	<h3>ADG Solution</h3>  <ul style="list-style-type: none">▪ 432GB Usable Capacity▪ RAID ADG = 14 drives▪ 3U Form Factor▪ Rough deployment cost \$20K 1 Enclosure, 14 36GB drives▪ Half the power consumption▪ Half the heat emission
---	---

When should customers use RAID ADG?

Advanced Data Guarding is ideal for applications requiring larger storage volumes. RAID ADG is best implemented when customers have the following needs:

- For greater capacity utilization
- For mission-critical data
- When large logical drives are required
- When better capacity utilization than RAID 1 is needed
- When higher fault tolerance than RAID 1 or RAID 5 is necessary

In addition, customers can migrate from their existing RAID levels to RAID ADG and from RAID ADG to other RAID levels as well.

Note: If a RAID volume does not have available capacity equal to one disk drive, it may require adding a disk drive and expanding the RAID volume. Refer to the Compaq User Manual for more detailed information.

Summary

Proper array configuration is critical to maintaining needed availability, performance, and capacity. Compaq provides reliable arrays for high-end workgroup servers' and low-end departmental servers' storage needs that include RAID ADG, RAID 5, and RAID 1+0. For customers, recognizing the advantages and risks of specific RAID configurations can minimize unplanned and costly downtime. Customers should consider these factors when choosing a RAID level in order to meet their particular requirements.