



Sun StorEdge™ 3310 SCSI Array Best Practices Manual

Architectures and Tips for Optimum Usage

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054 U.S.A.
650-960-1300

Part No. 816-7293-11
October 2002, Revision A

Send comments about this document to: docfeedback@sun.com

Copyright © 2002 Dot Hill Systems Corporation, 6305 El Camino Real, Carlsbad, California 92009, USA. All rights reserved.

Sun Microsystems, Inc. and Dot Hill Corporation may have intellectual property rights relating to technology embodied in this product or document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and other countries.

This product or document is distributed under licenses restricting its use, copying distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Third-party software is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and in other countries.

Federal Acquisitions: Commercial Software - Government Users Subject to Standard License Terms and Conditions.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright © 2002 Dot Hill Systems Corporation, 6305 El Camino Real, Carlsbad, Californie 92009, USA. Tous droits réservés.

Sun Microsystems, Inc. et Dot Hill Systems Corporation peuvent avoir les droits de propriété intellectuels relatants à la technologie incorporée dans ce produit. En particulier, et sans la limitation, ces droits de propriété intellectuels peuvent inclure un ou plus des brevets américains énumérés à <http://www.sun.com/patents> et un ou les brevets plus supplémentaires ou les applications de brevet en attente dans les Etats - Unis et les autres pays.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a.

Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACHES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.



Contents

Entry-level Server Storage Requirements	2
Best Practices: Print Servers	4
Best Practices: File Servers	5
Best Practices: Application Servers	7
Best Practices: Mail Servers	9
Best Practices: Database Servers	11
Best Practices: Consolidated Servers	13
Summary	14

Abstract

This document outlines small and large storage solutions for six common entry-level server environments:

- Print
- File
- Application
- Email
- Database
- Consolidation

These solutions can optimally use the Sun StorEdge 3310 SCSI Array, a next-generation Ultra3 SCSI storage system designed to provide direct attached storage (DAS) to entry-level servers.

These solutions feature many of the performance and reliability, availability and serviceability (RAS) features using familiar SCSI technology, and can be used as-is or tailored to fit your exact needs.

Examples of customization opportunities include choosing RAID protection levels, selecting SCSI bus configurations, adding more disks and adding disk enclosures. A large number of potential storage solutions exist between the small and large recommendations. Choosing a configuration that fits within this range will provide the best results.

Entry-level Server Storage Requirements

Single Processor Server Applications

Print, file and application services are essential network requirements and have become among the most popular uses of entry-level servers. The servers used to provide these functions are generally very inexpensive, highly compact units that are often installed in racks for convenience.

One example of such a server is the Sun Fire V120, an expandable single-processor server that occupies just 1 unit (1U) of rack space. Network servers are often distributed throughout an enterprise, making the Sun StorEdge 3310 SCSI Array ideally suited for these applications due to its cost-effective direct attached storage design.

TABLE 1 Storage Requirements for Single Processor Servers

	Print Server	File Server	Application Server
Availability	Medium	Medium to High	Medium to High
Storage Capacity	Low	Low to High	Low to Medium
Special Needs	High Data Rates and Low Cost	High Data and Transaction Rates	Low Cost and High Transaction Rates
Access Pattern	Sequential	Sequential	Random

Multiple Processor Server Applications

Another popular use of entry-level servers is for providing mail, database and other services to workgroup, departmental and branch office users. The server and storage solutions used to provide these sophisticated services must be able to scale in application performance and storage capacity to keep pace with the needs of increasing number of network users.

The Sun Fire V480 Server and StorEdge 3310 SCSI Array are excellent examples of highly scalable solutions and combine to form a powerful solution. Performance scales by adding server memory and processors while capacity scales by adding enclosures and disks.

TABLE 2 Storage Requirements for Multiple Processor Servers

	Mail Server	Database Server	Consolidated Server
Availability	High	High	High
Storage Capacity	Medium to High	Medium to High	Medium to High
Special Needs	High Transaction and Data Rates	High Transaction and Data Rates	High Transaction and Data Rates
Access Pattern	Random	Random	Random

Best Practices: Print Servers

Architecture and Configuration

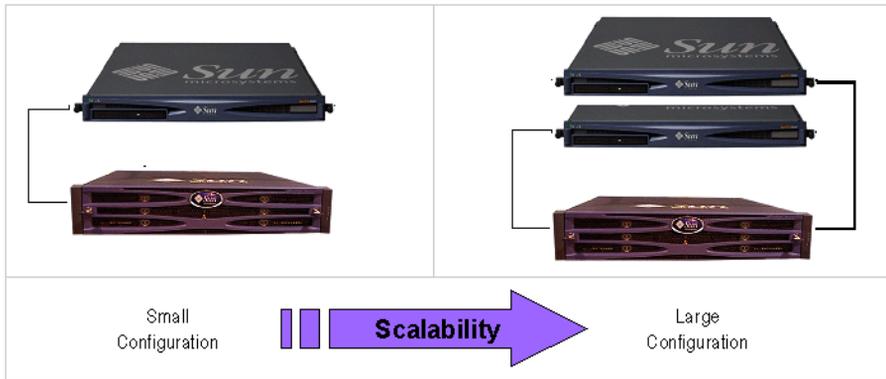


FIGURE 1 Optimized Architecture for Print Servers

TABLE 3 Configuration Details for Print Servers (JBOD Only)

	Small Configuration	Large Configuration
RAID Enclosures	0	0
JBOD Enclosures	1	1
Number of Controllers	Not Applicable	Not Applicable
Number of Disks	5	12
Bus Configuration	Dual Bus	Dual Bus
Cache Optimization	Not Applicable	Not Applicable
RAID Levels Used	Host Based	Host Based
Drive Configuration	Two LUNs One Spare	Two LUNs One Spare

Tips and Techniques

- Using a server's built-in SCSI port (if compatible) to minimize cost will result in satisfactory print server performance in most environments, even if the SCSI port does not operate at full Ultra160 speeds.
- Data protection can be provided via host-based software volume management software using the operating system's volume manager or a third-party volume manager. The recommended configurations will provide RAS similar to a single controller RAID array.
- Each server must be connected to a different SCSI bus when using the large recommended configuration.
- When adding a second server, one RAID 1 logical drive can be reassigned to the second server by moving its drives to the second SCSI bus if the server operating systems and volume managers are compatible.

Best Practices: File Servers

Architecture and Configuration

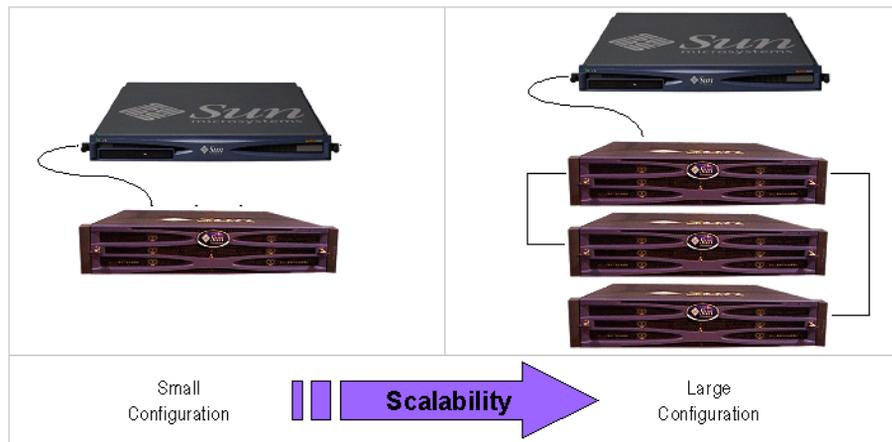


FIGURE 2 Optimized Architecture for File Servers

TABLE 4 Configuration Details for File Servers

	Small Configuration	Large Configuration
RAID Enclosures	1	1
JBOD Enclosures	0	2
Number of Controllers	1	2
Number of Disks	5	36
Bus Configuration	Single Bus	Single Bus
Cache Optimization	Sequential	Sequential
RAID Levels Used	RAID 3	RAID 5
Drive Configuration	One LUNs One Spare Drive	Two LUNs One Spare Drive

Tips and Techniques

- You can use an Ultra3 SCSI port to connect the RAID array whenever the server is connected to a LAN via gigabit Ethernet. Otherwise, a slower SCSI connection may create a performance bottleneck on the SCSI connection.
- If the server has room for only one host adapter and you must choose between Ultra160 SCSI or Gigabit Ethernet adapters, install the Gigabit Ethernet host adapter and use the server's built-in SCSI port to provide the most benefit to users. This will, however, reduce the I/O to a slower SCSI speed.
- A configuration with redundant RAID controllers and two logical drives provides very high transaction rates, even if there are a relatively small number of disk drives in use.
- Adding disk drives to logical drives will linearly improve transaction performance whenever cache misses occur, regardless of whether the additional storage capacity is utilized or not.
- Rapidly growing network response times as users are added is an indication the file server is limiting performance. If this occurs, examine the utilization of server memory, processors and network adapters and expand those with the highest usage.

Best Practices: Application Servers

Architecture and Configuration

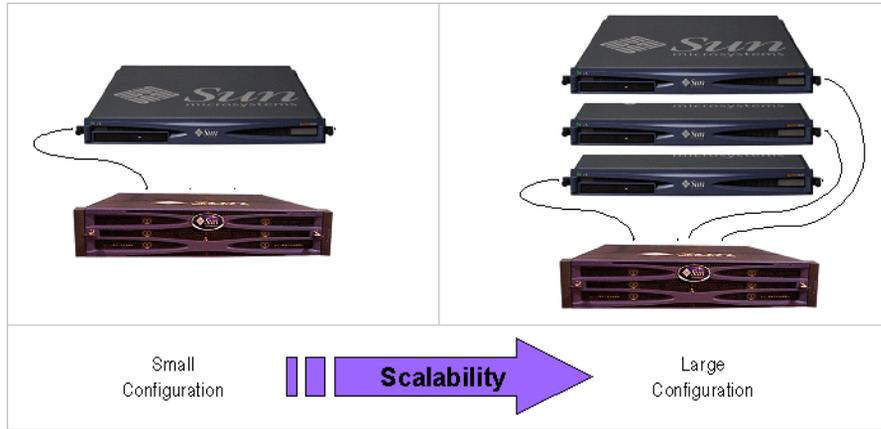


FIGURE 3 Optimized Architecture for Application Servers

TABLE 5 Configuration Details for Application Servers

	Small Configuration	Large Configuration
RAID Enclosures	1	1
JBOD Enclosures	0	0
Number of Controllers	1	2
Number of Disks	5	12
Bus Configuration	Dual Bus	Dual Bus
Cache Optimization	Random	Random
RAID Levels Used	RAID 5	RAID 5
Drive Configuration	One LUN One Spare Drive	Two LUNs One Spare Drive

Tips and Techniques

- A single RAID array providing storage for two servers reduces storage costs with little to no effect on application performance.
- Use the built-in SCSI ports of application servers to further minimize costs rather than adding a host adapter, particularly when not using Gigabit Ethernet to the LAN. Even a 40 MB/sec SCSI port is several times faster than a Fast Ethernet port.
- A second RAID controller can be added to enhance RAS without adding disks and creating a second Logical Drive. Redundant RAID controllers will operate in an active-standby mode when only one logical drive is available.
- Enhance application server availability by booting them from the RAID array rather than their internal drives. This also facilitates the rapid replacement of malfunctioning or failed servers.
- If attaching the RAID array to multiple hosts, you will want to assign a separate LUN to each server and SCSI bus.

Best Practices: Mail Servers

Architecture and Configuration

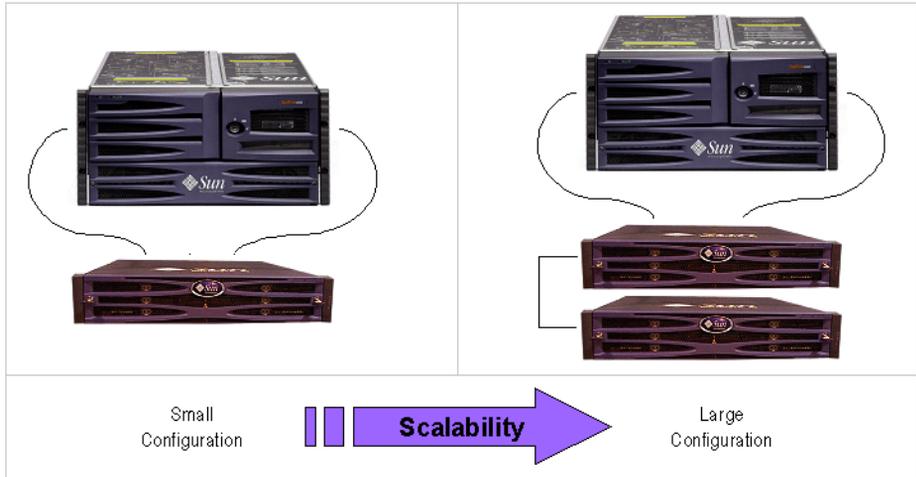


FIGURE 4 Optimized Architecture for Mail Servers

TABLE 6 Configuration Details for Mail Servers

	Small Configuration	Large Configuration
RAID Enclosures	1	1
JBOD Enclosures	0	1
Number of Controllers	2	2
Number of Disks	12	24
Bus Configuration	Dual Bus	Dual Bus
Cache Optimization	Random	Random
RAID Levels Used	RAID 3 and 5	RAID 3 and 5
Drive Configuration	Two LUNs One Spare Drive	Two LUNs One Spare Drive

Tips and Techniques

- Definitely use two Ultra160 SCSI ports when connecting the RAID array to the server for best performance.
- The recommended configuration for email servers consists of one RAID 3 and one RAID 5 logical drive. Use the RAID 3 logical drive to store attachments and other large files and the RAID 5 logical drive to store messages and small files.
- There is no need for the two logical drives to include similar numbers of disks. Assign disks as needed to establish the desired amounts of RAID 3 and RAID 5 storage capacity. Assign at least one disk as a spare.
- Scale storage capacity by adding disks to either the RAID 3 or RAID 5 logical drive, depending on need.
- Balance performance by assigning the RAID 3 logical drive to one RAID controller and the RAID 5 logical drive to the other RAID controller.
- When connecting a server to the array using two SCSI buses, you need to map each LUN to one SCSI bus so that they can both be active and have a dedicated path.

Best Practices: Database Servers

Architecture and Configuration

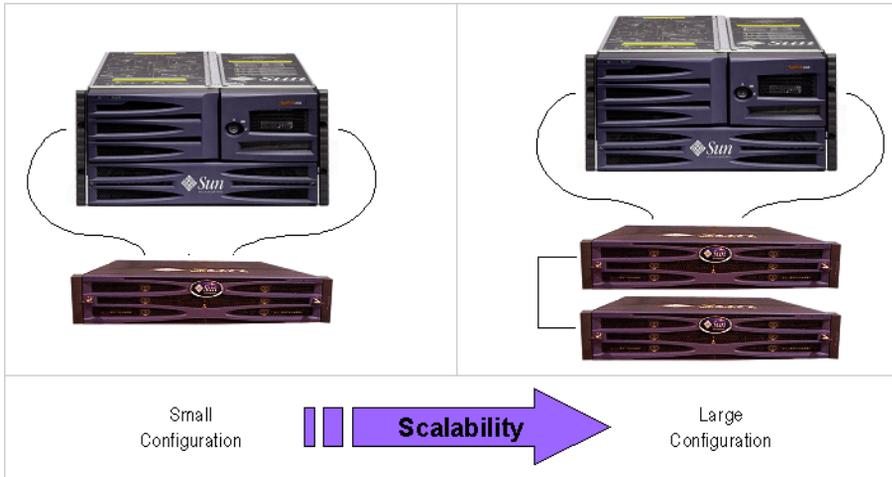


FIGURE 5 Optimized Architecture for Database Servers

TABLE 7 Configuration Details for Database Servers

	Small Configuration	Large Configuration
RAID Enclosures	1	1
JBOD Enclosures	0	1
Number of Controllers	2	2
Number of Disks	12	24
Bus Configuration	Dual Bus	Dual Bus
Cache Optimization	Random	Random
RAID Levels Used	RAID 1 and 5	RAID 1 and 5
Drive Configuration	Two LUNs One Spare Drive	Two LUNs One Spare Drive

Tips and Techniques

- Database servers make heavy use of storage resources. Always use Ultra160 SCSI ports for best performance. Further optimize performance by dedicating one RAID controller to each logical drive.
- The recommended configuration for database servers consists of one RAID 1 (0+1) and one RAID 5 logical drive. Use the RAID 5 logical drive to store data files and the RAID 1 logical drive to store tables and other performance sensitive files.
- Assign two drives as global spares, four drives to the RAID 1 logical drive (it automatically becomes RAID 0+1) and the remaining drives to the RAID 5 logical drive. Scale storage capacity by adding disks to the RAID 5 logical drive.
- Adding disk drives to RAID 5 logical drives will linearly improve transaction performance whenever cache misses occur, regardless of whether the additional storage capacity is utilized or not.
- For environments with extreme performance needs, use two of the small configuration arrays rather than one of the large configuration arrays. Doing so will double storage performance for less than twice the storage cost.
- When connecting a server to the array using two SCSI buses, you need to map each LUN to one SCSI bus so that they can both be active and have a dedicated path.

Best Practices: Consolidated Servers

Architecture and Configuration

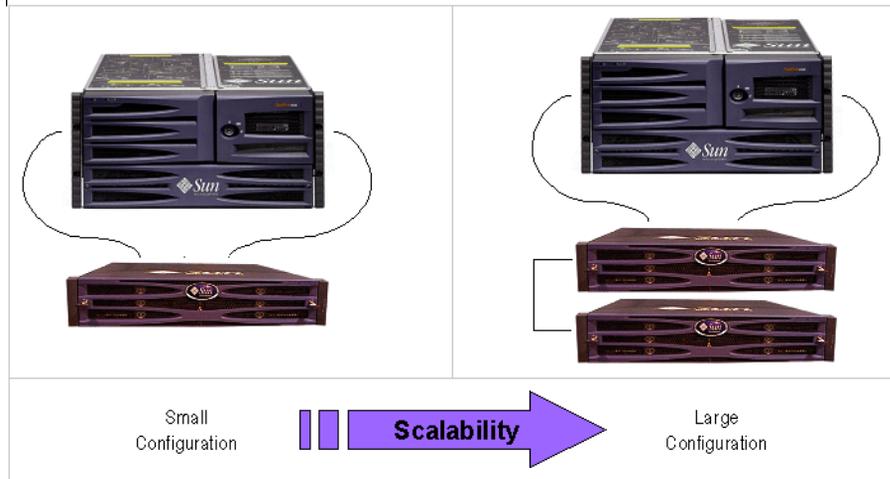


FIGURE 6 Optimized Architecture for Consolidated Servers

TABLE 8 Configuration Details for Consolidated Servers

	Small Configuration	Large Configuration
RAID Enclosures	1	1
JBOD Enclosures	0	1
Number of Controllers	2	2
Number of Disks	12	24
Bus Configuration	Dual Bus	Dual Bus
Cache Optimization	Random	Random
RAID Levels Used	RAID 1, 3 and 5	RAID 1, 3 and 5
Drive Configuration	Three LUNs One Spare Drive	Three LUNs One Spare Drive

Tips and Techniques

- Consolidated servers have very dynamic storage requirements. Use Ultra160 SCSI ports to eliminate potential bandwidth bottlenecks between the server and RAID array.
- If most of the server's resources are used for databases, create a RAID 1 logical drive using two disks, for log and transaction data. If more RAID 1 capacity is required in the future, create a new RAID 1 logical drive using two unassigned disks. Otherwise, use a portion of the RAID 3 logical drive for database storage as an alternative to a separate, dedicated RAID 1 logical drive.
- To avoid reduced performance, it is recommended that you add extra storage when a LUN reaches 80% full.
- Balance workloads by assigning the RAID 5 logical drive to one RAID controller and the RAID 3 logical drive to the other controller. If there is a RAID 1 logical drive created, assign it to the same RAID controller as the RAID 3 logical drive.
- When connecting a server to the array using two SCSI buses, you need to map each LUN to one SCSI bus so that they can both be active and have a dedicated path.

Summary

Entry-level servers are used for wide range of applications with distinct storage requirements, so the Sun StorEdge 3310 SCSI Array features a modular architecture with flexible configurations. For example, a storage solution can consist of a JBOD array, a RAID array or a combination of both.

Configuration preferences include user-selectable RAID protection levels, controller optimization and more. Modularity and flexibility allow the storage solution to be quickly and easily adapt to a particular environment.