

April 2002

Part Number:  
16JV-0402A-WWEN  
Edition 1.1

Prepared by:  
Global SAP Solutions  
Compaq Computer Corporation

**Contents**

**Introduction..... 3**  
  Overview ..... 3  
  Design Goals..... 4  
**Background..... 5**  
  I/O Subsystem Infrastructure .... 5  
  Disk I/O ..... 5  
  Disk Speed..... 6  
  Caching Data ..... 7  
  Backup and Restore  
  Considerations ..... 8  
  Enhanced Availability and  
  Multiple Data Pathing ..... 9  
**General Recommendations ..... 11**  
  Database Layout..... 11  
**Sample Configurations ..... 15**  
  Entry-level Configuration ..... 15  
  Medium Configuration ..... 17  
  High-end Configuration ..... 18

# StorageWorks Modular Array and Modular SAN Array Configuration Guide for mySAP.com

*Abstract:* This document contains information on how to configure Compaq StorageWorks Modular Array and Modular SAN Array disk subsystems for mySAP.com.

## 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.

Compaq, NonStop, Deskpro, Compaq Insight Manager, Systempro, Systempro/LT, ProLiant, ROMPaq, QVision, SmartStart, NetFlex, QuickFind, PaqFax, and Prosignia are registered with the United States Patent and Trademark Office.

ActiveAnswers, Netelligent, Systempro/XL, SoftPaq, Fastart, QuickBlank, QuickLock are trademarks and/or service marks of Compaq Computer Corporation.

Microsoft, Windows and Windows NT are trademarks and/or registered trademarks of Microsoft Corporation.

The following are trademarks or registered trademarks of SAP AG; ABAP/4, InterSAP, RIVA, R/2, R/3, R/3 Retail, SAP (Word), SAPaccess, SAPfile, SAPfind, SAPmail, SAPoffice, SAPscript, SAPtime, SAPtronic, SAP-EDI, SAP EarlyWatch, SAP ArchiveLink, SAP Business Workflow, and ALE/WEB. The SAP logo and all other SAP products, services, logos, or brand names included herein are also trademarks or registered trademarks of SAP AG.

Intel, Pentium and Xeon are trademarks and/or registered trademarks of Intel Corporation.

Oracle is a registered trademark of Oracle Corporation.

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

©2002 Compaq Computer Corporation. All rights reserved. Printed in the U.S.A.

StorageWorks Modular Array and Modular SAN Array Configuration Guide for mySAP.com  
White Paper prepared by Global SAP Solutions

Edition 1.1 (April 2002)

# Introduction

## Overview

SAP AG's mySAP.com is one of the most, if not the most, CPU, memory and I/O intensive application in the industry. Because the disk subsystem is the slowest part of a mySAP.com system it is extremely important to design it for optimal performance. Taking into account that the major number of companies using mySAP.com are very dependent on the service from the system, it is also very important for the mySAP.com system to have a high degree of availability.

This document concentrates on how to design the layout of a disk subsystem for a server running the database of a mySAP.com installation. Because the majority of large mySAP.com systems use Oracle on Unix or Oracle / Microsoft SQL Server in a Windows environment, the primary focus is on Oracle with specific comments on SQL Server when relevant. The sample configurations in this document focus on the StorageWorks HSG80-based RAID Array systems for the Open System Environment and the Modular Storage Array 1000 that is available for Intel-based servers today. The Enterprise Virtual Array is not covered by this document. Please see <http://emea-apps.eur.compaq.com/icc-sap/technical/ESGEnterprise.asp#2> for further information.

## Compaq Links

For further information on mySAP please refer to the following Compaq publications:

Compaq AlphaServer and Tru64 UNIX Performance Guide for mySAP.com

Enterprise Storage Group:

[http://www.compaq.com/storage/san\\_index.html](http://www.compaq.com/storage/san_index.html).

## SAP Links

Please refer to the following SAP publications for further technical information on mySAP.com:

SAP Advanced Planner & Optimizer – Installation on UNIX: Oracle Database

## Design Goals

Disk configuration is driven by three fundamental needs:

- Durability      Ensuring continued operation in the event of a single component failure
- Recoverability    Providing a robust system that can be rapidly restored in the event of serious hardware or software failure
- Performance      Balancing the I/O requirements of the database, application and operating system

The organization of database components on the disk subsystem and the provision of hardware redundancy determine the reliability of the database. All databases consist of data containers and transaction logs, which maintain a record of all changes that have been made to the database. In the event of system failure, successful recovery of the database depends on the availability of these transaction logs. It is therefore essential that the log files and data files are stored on physically separated disks. In the case of Oracle the “on-line” logs, once full, are archived to a separate area, which must also reside on a separate physical disk.

Equally important is the provision of hardware redundancy to ensure that the failure of a single hard disk drive does not result in total system failure. This is provided using a RAID-1, RAID-0/1 or RAID-5 disk array, supported by all Compaq StorageWorks solutions. The overall performance of the disk subsystem is very much influenced by the number of disks, the number of array controllers and, in the case of external storage, the number of host bus adapters plus the manner in which the database tables and logs are distributed across the available disks.

The following table lists the I/O characteristics of the key components in a mySAP.com system:

Object	Type of I/O	Access Characteristics
OS paging/swap file	Random	Continually accessed
R/3 log files	Sequential	Continually accessed, but low rate
Online log files	Sequential	Continually accessed, medium to high rate
Archived log files	Sequential	Occasionally, but high rate
Database tables	Random	Continually accessed

**Table 1, mySAP.com I/O Characteristics**

For performance reasons, it is critical to avoid mixing components with random and sequential I/O on the same volumes. As a minimum, the OS paging/swap file(s), online log and database tables have to be implemented on separate volumes or filesets. The executables, R/3 log files, and in the case of Oracle, the Oracle archive, backup, reorg, stat and trace files can be placed on a single volume or fileset. Further Oracle recommendations will be discussed later in this document.

## Background

### I/O Subsystem Infrastructure

In general, the advantages and benefits of external storage area network (SAN) based I/O storage subsystems should be known. However, we want to highlight the reasons why the sample configurations in this document are configured using Compaq Fibre Channel products.

**Performance** Compaq offers 1-gigabit and 2-gigabit Fibre Channel products, host bus adapters, FC switches and storage arrays. 1-gigabit and 2-gigabit products can coexist in the same SAN. See the *Heterogeneous Open SAN Design Reference Guide* for actual specifications at [http://www.compaq.com/storage/san\\_index.html](http://www.compaq.com/storage/san_index.html). Using available PCI slots efficiently, 2-gigabit FC host bus adapters offer a maximum transfer rate of 400 MB/s and are faster than a WideUltra3 SCSI channel by a factor of 2 to 3.

1-gigabit-based HSG80 products (MA8000, EMA12000, EMA16000) for Tru64 and Microsoft Windows environments provide a bandwidth of 70 MB/s per HSG80 pair.

The 2-gigabit-based MSA1000 product for a Windows environment provide a bandwidth of 160 MB/s in a 4-channel active configuration.

The actual numbers in terms of I/O in a typical or standard mySAP.com environment are generally not available for the specific products.

**Distance** The distance Compaq supports today in a FC-based SAN is 500 meters between FC switches and HBAs or storage subsystems using short-wave multi-mode cables. Distances beyond this limit up to 100 kilometers are spanned using single mode fiber between FC switches. The actual limits and capabilities are specified in the *Heterogeneous Open SAN Design Reference Guide*.

**Consolidation / Flexibility** As space requirements may grow rapidly in a mySAP.com environment, providing additional capacity without degrading performance is a major advantage of a SAN-based infrastructure. In addition, the server does not have to be touched if additional space requirements need to be satisfied, as long as the HBA or the PCI bus are not identified as the bottleneck in the I/O path of a mySAP.com configuration.

Experience has shown that choosing the fastest and most flexible I/O subsystem in an mySAP.com configuration is a wise decision. This is especially so when considering the hardware costs in relation to the overall costs of an SAP customer project and the performance problems generated by a poorly configured I/O subsystems.

### Disk I/O

The overall performance of the disk subsystem depends heavily on the number of disk drives installed. Therefore, it is essential to understand that large capacity disk drives are not preferable as storage media for a mySAP.com system. The type of disks used depends on the amount of data to be stored and the amount of I/O from/to the data to be accessed on the disk drives. In OLTP environments, where the majority of access to the disk subsystem is random or direct, the responsiveness of the database server depends largely on the amount of I/O the disk subsystem is able to deliver.

The more drives are configured in a volume (RAID-0, RAID-0/1 or RAID-5), the more I/Os per second and the higher the transfer rate the volume is able to provide.

Most hard disks used in disk subsystems configured for mySAP.com provide a similar number of I/Os per second, no matter the speed at which they rotate. They are all able to deliver a maximum of around 120-150 random I/Os per second. The capacity of the drive has no significant influence on the number of I/Os per second. The I/O rate primarily depends on how fast the disk drives are able to seek and re-position the read/write heads.

Implementing a database of 70 GB using drives with different storage capacity gives the following table:

# of Drives	Drive Type	RAID	Capacity	I/Os per Sec
2	72.8 GB	1	70.4 GB	130
4	36.4 GB	0+1	70.4 GB	260
8	18.2 GB	0+1	70.4 GB	520
16	9.1 GB	0+1	70.4 GB	1040

**Table 2, Increasing I/O Capabilities of a Volume Depending on Size of Disk Drives**

This comparison shows that the number of I/Os a disk subsystem can provide is highly dependent on the number of disk drives that build a volume or a LUN. If a volume is not able to provide the amount of I/O required, there are only two or three ways in which this can be compensated for. Either reduce the amount of I/O required by moving some of the data to another volume, try to cache more data in the database data buffer, or increase the number of drives for this volume.

## Disk Speed

The performance of the disk subsystem can benefit from having highly sequential accessed data implemented on the fastest rotating drives. At the moment, Compaq offers disk drives that rotate at either 10,000 or 15,000 rpm. The online log, mirror log, archive and reorg are areas that benefit from using 15,000 rpm disk drives.

In an average mySAP.com system, 40-50 percent of access to the database consists of sequential table scans. But in the case of Oracle, a table's data cannot be expected to be stored consecutively on the disk; especially not in a system that has been running for some time without any database reorganization. So, even if data is being read logically sequentially, the physical read of data from disk cannot be expected to be sequential.

It is important to understand that a customer will get a better return on investment by investing in more memory for the database server than by investing in the fastest rotating hard disks.

## Caching Data

In a mySAP.com system, cache or buffers exist on several different levels, where the lowest level is on disk/disk controller level and the highest is the buffers in the application server.

From a mySAP.com system point of view, there are several different caches involved in the system. The most important are:

Cache Type	Comment	Context Awareness	Hit Rate Should Be
Application, information	Cached R/3 tables	Very high	100%
Application, work areas	Roll	Very high	100%
Application, work areas	Page	Very high	Should be high, but not much used in R/3 4.x
Application, work areas	SAP extended memory, OS heap	Very high	As high as possible
Database	Data and SQL statement	High	As high as possible
File system	OS read-ahead and write-back	Very low	Low in OLTP (direct access); high in sequential access
Disk controller	Read-ahead and write-back	very low	Low in OLTP (direct access); high in sequential access

**Table 3, Cache Overview in a mySAP.com System**

The closer the data is to the application server, the more knowledge the system has about the data in the buffer. Also, the closer the data is to the application server, the more expensive it is to fetch it from disk. The closer the data is to the disk, the less understanding of the use of the data is available, therefore no intelligent access methodology can be implemented in the OS file system or disk controller cache except for read-ahead.

Today's disk controllers like the Compaq SMART Array Controller and Compaq StorageWorks HS Controllers are able to detect multiple threads or streams of sequential access, therefore enabling them to read ahead from several areas on the disk. A mySAP.com system will not benefit significantly from large disk controller caches because most of the data reads are random or direct.

These array controllers have algorithms built in that enable the controller to write data to the disk in the most efficient ways in respect to minimizing the re-positioning of the read/write heads, using write-back cache.

An additional benefit of write-back cache is that data can be cached in peak periods where the disks are not able to write data at the rate it is received from the application. The larger the (write-back) cache on the controllers, the longer it will take for the cache to be saturated. When the cache is saturated, the amount of data the controllers are able to receive depends entirely on how fast it is able to write to the disks. In OLTP systems, where the majority of writes are direct or random, the higher the amount of I/O the logical volumes are capable of, the faster the disk subsystem is able to accept data from the application.

In large retail, production planning and other systems where a high amount of data is being uploaded in batches, a large number of disk drives is required. These drives have to be connected via several controllers and host bus adapters to minimize the write time as much as possible.

When a disk subsystem in an OLTP environment like mySAP.com is facing performance problems, it comes down to the number of I/Os the disk subsystem is able to provide and not the amount of megabytes the disk subsystem is able to transfer per second. This results in the need for more I/Os per second. To be able to get sufficient I/Os per second, the number of disk drives required will result in a data capacity that exceeds the actual requirements for storing the customer's data.

Therefore, the number of disk drives is essential to the performance of the disk subsystem. The smaller the capacity of the disk drives installed, the more the number of drives to be installed and eventually the more array controllers and host bus adapters, all resulting in a higher amount of I/O per second.

However, reducing the required number of I/Os for a specific volume is a much more efficient way of solving the problem. Installing a larger cache can do this, but not a larger disk controller cache. Because the database knows much more about the data than the disk controller, increasing the database cache will be much more efficient than enhancing the disk subsystem. Making the database cache larger will result in a larger hit rate, thus reducing the number of reads the disk subsystem has to deal with. Reducing the number of reads will result in more bandwidth to writes. And because it is much faster to read from memory than from a well performing disk subsystem, the average database request time will decrease, thus benefiting all application servers and users.

A positive side effect from increasing the database cache or buffer and thus reducing the I/O load is a reduced load on the database server's CPUs. The reason for this is that when the database engine is not able to find the required information in the database buffer, the engine has to look up in the index eventually reading additional index information from disk about where to find the required data on disk. Each time a disk read is required; the part of the OS taking care of disk access has to be engaged. All this looking up in the index, going through the OS file system, device driver, interrupt handling, etc. puts additional load on the database server's CPU. Therefore, reducing the number of disk reads also reduces the load on the database server's CPUs.

The amount of additional memory that needs to be installed on the database server is difficult to predict. In an Oracle system, the ratio between logical reads and physical reads for a typical and reasonably responsive system is in the range of 30:1. By installing more memory this ratio can be raised to 50:1. Systems with a ratio better than 100:1 have been seen. Also, systems with more than 5 GB assigned to the Oracle SGA have been seen.

## Backup and Restore Considerations

When designing a disk subsystem for mySAP.com, backup and restore considerations also have to be taken into account. Questions like:

- Does the customer require 24/7 uptime?
- Is the system going to be backed up online or offline or online on workdays and offline at the weekend?
- What is the size of an offline backup window?
- How does online backup affect daily production?
- What are the requirements for a system restore in case of disaster?

Answers to these questions have an influence on the backup strategy and on the design of the disk subsystem. It can happen that there is no standard backup strategy that fits for a productive database server based on some of the mentioned requirements. For customers like these, HSG80-based snapshot/cloning functionality (described in *StorageWorks Rapid Restore and Non-Disruptive Backup for mySAP.com*) might be the answer.

Customers who need to implement a disaster tolerant solution between distant computing sites and combine this with a backup solution can benefit from Compaq SANworks Data Replication Manager, the HSG80's controller-based peer-to-peer remote copy capability. This functionality can also be achieved by using the RDBMS capabilities of Oracle or SQL Server and implementing logical log shipping. These techniques are very dependent on specific customer requirements and are beyond the scope of this paper.

To plan a disk configuration for mySAP.com with a standard backup strategy directly to tape or a tape library system in mind, the following considerations have to be taken into account. The type of tape, backup management software and storage subsystem in a FC-based solution has to be certified in the Compaq Enterprise Backup Solution (EBS) matrix. The current compatibility matrix is available at <http://www.compaq.com/products/storageworks/ebs>.

To get the fastest backup, it is essential for the tape device to stream constantly without interruption. This means that the tape device receives data as fast as it is able to write it to tape (feed speed). The higher the compression rate, the larger the amount of data that has to be transferred from the disk subsystem to the tape device. In the case of the current SDLT technology we have seen the need for a feed speed of +20 MB/s per tape device (72 GB/hour). Experience has shown that an HSG80-based subsystem, for example, can deliver +120 GB/h in a mySAP.com environment. It is therefore necessary to ensure that there is sufficient bandwidth on the I/O path to the accessed devices to complete a backup in the available timeframe. A utility supporting the backup sizing process in an EBS environment is available at <http://www.compaq.com/products/storageworks/ebs/ebstoolsbackupsizing.html>

Bear in mind that each time the disk capacity on a database server is increased, the backup capabilities have to be reviewed to ensure that the system is able to back up the data in the available backup window and meet the restore requirements.

## Enhanced Availability and Multiple Data Pathing

There are several ways in which the disk subsystem can be enhanced to increase availability:

- Online spare disks
- Redundant power supplies
- Dual redundant array controllers
- Mirrored cache modules
- Redundant external cache batteries

This functionality is available on all Compaq Fibre Channel based solutions outlined in the sample configurations in this document. A very efficient and common way to have a no single point of failure (NSPOF) subsystem configuration is to implement a redundant data path from the server to the disk subsystem. With at least two host bus adapters (HBA), more than one physical path to the volumes is established. If a component in one data path fails (HBA, cable, FC switch, controller), it is the task of the OS to switch the I/O stream to the remaining path. A further benefit of having more physical paths to the storage is that, under normal conditions, the I/O workload can be spread dynamically between the available paths. This dual rail functionality is implemented either in the operating system (OVMS and Tru64) or by the **SANworks Secure Path** layered software (Windows, SUN, HP, IBM).

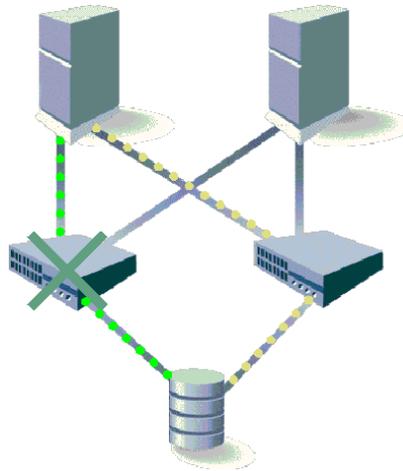


Figure 1, Dual Path Failover

## General Recommendations

This section outlines the disk space requirements for mySAP.com for Oracle and SQL Server.

### Database Layout

#### Oracle

The SAP installation procedure requires six tablespace sets by default. Table 4 lists the areas where a mySAP.com installation based on Oracle requires disk space. The listing does not contain areas for the operating system and paging/swap areas:

Description	R/3 Name	Min. Size
SAP executables		200 MB
Oracle executables		200 MB
Oracle archive file	SapArch	300 MB
Space used for database reorganization	SapReorg	600 MB
Oracle online log files, group 1, 3	OrigLogA	41 MB
Oracle online log files, group 2, 4	OrigLogB	41 MB
Mirrored Oracle online log files, group 1, 3	MirrorLogA	41 MB
Mirrored Oracle online log files, group 2, 4	MirrorLogB	41 MB
Oracle trace files	SapTrace	20 MB
Space used for Oracle backup	SapBackup	10 MB
Space used for SAPDBA	SapCheck	20 MB
Space for statistical information	SapStat	10 MB
Oracle tablespace set #1	SapData1	
Oracle tablespace set #2	SapData2	
Oracle tablespace set #3	SapData3	
Oracle tablespace set #4	SapData4	
Oracle tablespace set #5	SapData5	
Oracle tablespace set #6	SapData6	

**Table 4, Oracle Basic Fileset**

Table 5 shows which tablespaces reside in which tablespace sets. The value *Initial* refers to tablespaces that are read only under normal circumstances during the startup phase or written during the shutdown phase of the system. For large systems, which in this context means systems with very high database activity, the I/O load on PSAPROLL, PSAPROLLBIG and PSAPTEMP can become a bottleneck. Because PSAPROLL and PSAPROLLIG are tables with system-wide influence, special attention has to be taken to ensure that the RAID sets for these tables provide sufficient I/O capabilities.

Tablespace Set	Tablespace	Read Intensive	Write Intensive	Usage
SapData 1	SYSTEM	No	No	Oracle system information
	ROLL	Yes	Yes	Rollback segments
	(ROLLBIG)	(Yes)	(Yes)	Rollback segments
	TEMP	Yes	Yes	Temporary segments
	BTABD	Yes	Yes	Operational data
	DDICD	No	No	Data dictionary data
SapData2	POOLI	Yes	Yes	Indexes
	BTABI	Yes	Yes	Indexes
	USER1I	No	No	Indexes
	CLUD	Yes	Yes	Cluster table data
SapData3	SOURCED	No	No	Program source data
	LOADI	Initial	No	Indexes
	ELxyl	Initial	No	Indexes
	CLUI	Yes	Yes	Indexes
	DOCU	No	No	Indexes
	STABD	Yes	Yes	Master data
SapData4	PROTI	Yes	Yes	Indexes
	PROTD	Yes	Yes	Temporary data
	LOADD	Initial	No	Compiled program data
	SOURCEI	No	No	Indexes
	ESxyl	No	No	Indexes
	USER1D	No	No	User-defined data
SapData5	ELxyD	Initial	No	Compiled program data
	STABI	Yes	Yes	Indexes
	DOCUD	No	No	Documentation data
	POOLD	Yes	Yes	Pool table data
SapData6	DDICI	No	No	Indexes
	ESxyD	No	No	Program source data

**Table 5 Oracle Tablespace Layout**

To design a configuration for high I/O capacity, the SapData1 area can be split into two or three areas: SapData1 and a new SapData7 and eventually a SapData8 in such a way that SapData7 contains PSAPROLL, PSAPROLLBIG and PSAPTEMP, or SapData7 contains PSAPROLL and PSAPROLLBIG and SapData8 contains PSAPTEMP. SapData1 contains the remaining tables.

Analyzing several databases ranging from 270 GB to over 900 GB from manufacturing and retail systems gives the minimum and maximum sizes shown in Table 6. The entry "Static" covers all the tables marked as *No* or *Initial* under Read/Write Intensive in Table 5.

Tablespace	Min. Size MB	Max. Size MB	Min. Reads in K	Max. Reads in K	Min. Write in K	Max. Write in K
Static	5,697	19,196	408	3,050	28	313
PSAPROLL	2,048	60,702	24	2,806	935	10,232
PSAPTEMP	2,995	9,314	19	511	9	316
SYSTEM	307	512	67	485	1	6

**Table 6, Min. and Max. Size of "Static" Tables plus Roll, Temp and System**

Table 6 shows that “Static” tables are almost static in size whether the total database is 270 GB or 900 GB. The amount of I/O is expected to be small for the “Static” tables compared to PSAPROLL. Therefore, it makes sense to place some of the “Static” tablespaces and SYSTEM on the same RAID set as the PSAPROLL, PSAPROLLBIG and PSAPTEMP.

The following table shows a proposed layout using **eight** tablespace sets:

Tablespace Set	Tablespace	Read Intensive	Write Intensive	Usage
SapData1	BTABD	Yes	Yes	Operational data
	DDICD	No	No	Data dictionary data
	POOLI	Yes	Yes	Indexes
SapData2	BTABI	Yes	Yes	Indexes
	USER1I	No	No	Indexes
	CLUD	Yes	Yes	Cluster table data
SapData3	ELxyl	Initial	No	Indexes
	CLUI	Yes	Yes	Indexes
	STABD	Yes	Yes	Master data
	PROTI	Yes	Yes	Indexes
SapData4	PROTD	Yes	Yes	Temporary data
	ESxyl	No	No	Indexes
	USER1D	No	No	User-defined data
	ELxyD	Initial	No	Compiled program data
SapData5	STABI	Yes	Yes	Indexes
	POOLD	Yes	Yes	Pool table data
	DDICI	No	No	Indexes
SapData6	ESxyD	No	No	Program source data
SapData7	ROLL	Yes	Yes	Rollback segments
	(ROLLBIG)	(Yes)	(Yes)	Rollback segments
	LOADI	Initial	No	Indexes
	DOCUI	No	No	Indexes
	SOURCED	No	No	Program source data
SapData8	SYSTEM	No	No	Oracle system information
	TEMP	Yes	Yes	Temporary segments
	LOADD	Initial	No	Compiled program data
	DOCUD	No	No	Documentation data
	SOURCEI	No	No	Indexes

**Table 7, Tablespace Layout for Larger Systems with Eight Tablespace Sets**

Table 7 shows a layout where several of the low read/write-intensive tablespace sets have been moved to SapData7 and SapData8 to make use of the excessive storage capacity installed on these RAID sets to further enhance the I/O capabilities.

## SQL Server

SAP recommends the following minimum configuration:

RAID	R/3 Name	Min. Size	Description
1	MSSQL7	65 MB	MS SQL Server 7 executables
	MSSQL7DB	45 MB	master, msdb
	TEMPDB	280 MB	Area for temporary data
1	<SAPSID>LOG<n>	1.5 GB	R/3 transaction log file
5	<SAPSID>Data<n>	6 GB	Data files

**Table 8, Microsoft SQL Basic Files**

For integrity reasons, the MASTER and MSDB are separated from the normal data files. In larger installations with high I/O activity, the TEMPDB can benefit from being implemented on a separate RAID-0/1 consisting of several drives striped together to increase the amount of I/O available. In such systems, the volume(s) for data files should also be implemented as RAID-0/1 instead of RAID-5.

Microsoft SQL does not split a table into a data part and index part; instead, SQL distributes its information as equally as possible over all its data files. Distributing these data files equally over all available volumes ensures I/O distribution over all volumes (and controllers). Therefore, the optimum disk configuration would be an even number of equal sized volumes distributed over several array controllers.

The data files should be implemented in such a way that the amount of data per RAID set is the same relative to the number of disks going into a RAID set. Optimal file sizes are 4 MB, 8 MB, 16 MB or 32 MB. The following table shows examples:

# of Drives	Drive Type	RAID	Capacity	I/O Sec	Files	I/O per File
6	18.2 GB	0+1	52.8 GB	390	3 x 16 GB	130
6	18.2 GB	0+1	52.8 GB	390	6 x 8 GB	65
6	18.2 GB	0+1	52.8 GB	390	12 x 4 GB	33
8	18.2 GB	0+1	70.6 GB	520	4 x 16 GB	130
8	18.2 GB	0+1	70.6 GB	520	8 x 8 GB	65
8	18.2 GB	0+1	70.6 GB	520	17 x 4 GB	31
10	18.2 GB	0+1	88.2GB	650	5 x 16 GB	130
10	18.2 GB	0+1	88.2GB	650	11 x 8 GB	60
10	18.2 GB	0+1	88.2GB	650	22 x 4 GB	31
12	18.2 GB	0+1	105.8 GB	780	3 x 32 GB	260
12	18.2 GB	0+1	105.8 GB	780	6 x 16 GB	130
12	18.2 GB	0+1	105.8 GB	780	12 x 8 GB	65
12	18.2 GB	0+1	105.8 GB	780	26 x 4 GB	31

**Table 9, I/O per File and Disk Drive**

If, for some reason, it is necessary to have RAID sets consisting of 12 and 10 disks, there should be six 16-GB files on the 12-drive RAID set and five 16-GB files on the 10-drive RAID set. This ensures that the load on each drive is the same.

## Sample Configurations

The sample configurations in this section can be regarded as a starting point for various mySAP.com installations on Compaq StorageWorks platforms for a Tru64 or a Windows environment. It might be necessary to introduce configuration changes for various reasons and customer-specific requirements. The SQL Server sample configurations are based on the MSA 1000 Storage Array because the MSA1000 is currently only supported in a Windows environment. SQL Server configurations can be built on the HSG80-based storage arrays; the Oracle configurations on Windows can also be configured on the MSA1000.

### ***Tru64 Cluster Configurations***

When enhancing a configuration with a Tru64 Cluster solution, bear in mind that at least 4 more LUNs need to be configured for the MA8000. This is because the member-specific boot disks and the cluster common root have to be on the shared storage as well as the Quorum disk and that all have to be in a RAID 1 configuration.

### ***MSCS with Microsoft SQL Server and Oracle Fail Safe Configurations***

For details on implementing a Microsoft cluster solution in a mySAP.com environment, please refer to the <http://emea-apps.eur.compaq.com/icc-sap/technical/hasapmssol.asp> Web page.

## Entry-level Configuration

### **Oracle**

The first example is a minimum mySAP.com configuration based on Oracle using the Compaq StorageWorks MA8000, the HSG80 Enterprise/Modular Storage RAID Array Fibre Channel solution. This configuration can be regarded as an entry-level setup for the MA8000 starting with 11 x 18.2-GB drives with initial 91 GB available for SAPData and 1 Hot Spare disk. The MA8000 is configured with 3 split bus disk shelves that contain 42 drives at maximum.

Disks	RAID	Capacity	Description
1 x 18.2 GB			Online spare
6 x 18.2 GB	5	91.0 GB	Oracle tablespace files SAPData1
2 x 18.2 GB	1	18.2 GB	Archived log
2 x 18.2 GB	1	18.2 GB	Online redo log, Backup, Check, Reorg, Stat, Trace, SAP&Oracle Exe

**Table 10, Oracle Entry-level Disk Configuration for mySAP.com**

Configuring RAID 0+1 on the HSG80 is achieved via striped mirror sets. A 12-member RAID 0+1 set consists of a stripe set of 6 two-member mirror sets. To enhance availability, each member of a mirror set should reside in a different storage enclosure. The default stripe width is fine.

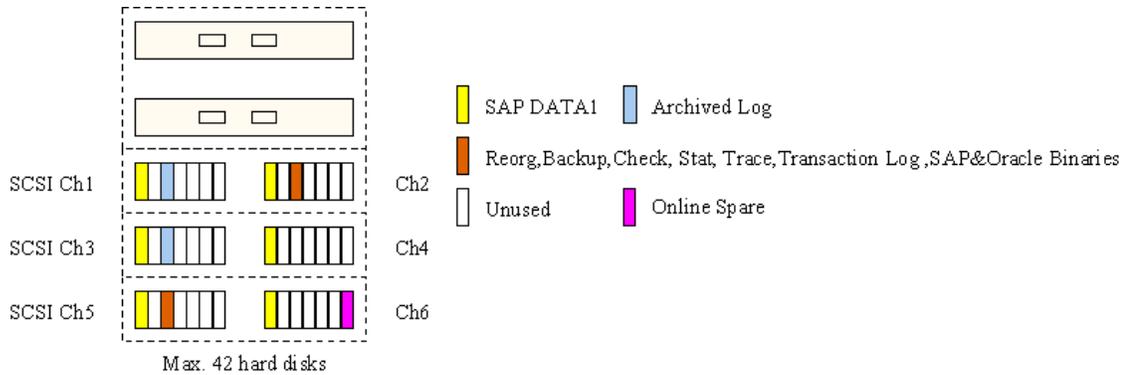


Figure 2, StorageWorks MA8000 - Oracle Entry Level

### Microsoft SQL Server

Using SAP’s recommended minimum configuration for a clustered SQL Server system, we get the following configuration:

RAID	R/3 Name	Min. Size	Description
1	MSSQL7	65 MB	MS SQL Server 7 executables
	MSSQL7DB	45 MB	master, msdb
	TEMPDB	280 MB	Area for temporary data
1	<SAPSID>LOG<n>	1.5 GB	R/3 transaction log file
5	<SAPSID>Data<n>	6 GB	Data files

Table 11, Microsoft SQL Basic Files for mySAP.com

For integrity reasons, the MASTER and MSDB are separated from the normal data files. In larger installations with high I/O activity, the TEMPDB can benefit from being implemented on a separate RAID-0/1 consisting of several drives striped together to increase the amount of I/O available. In such systems, the volume(s) for data files should also be implemented as RAID-0/1 instead of RAID-5.

SQL Server does not split a table into a data part and index part; instead, SQL Server distributes its information as evenly as possible over all its data files. Distributing data files evenly on all available volumes ensures I/O distribution over all volumes.

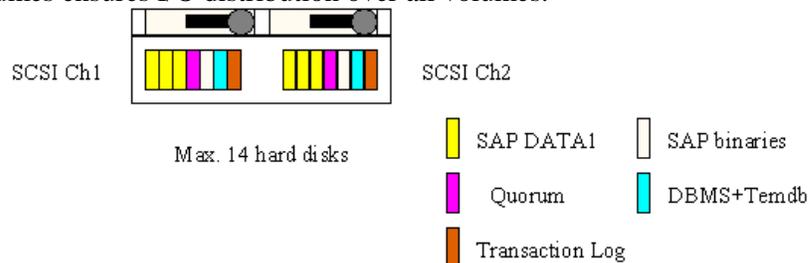


Figure 3, StorageWorks MSA1000 - SQL Server Entry Level

Figure 3 illustrates an MSA1000 configuration providing about 90 GB for SAPData on a 6 x 18.2-GB-member RAID 5 set.

## Medium Configuration

### Oracle

The following configuration can be regarded as a medium-level setup for the EMA12000 starting with 52 x 18-GB disk drives providing 455 GB raw (RAID 0+1) disk space with an initial 327 GB available for SAPData. This configuration can grow by 273 GB (83 percent) for SAPData to 82 drives with 728 GB raw disk space just by adding additional drives. In the medium configuration the EMA12000 is configured with single bus shelves that allow 84 drives at maximum.

In addition, it is possible to upgrade from entry level to medium by reconfiguring from split bus to single bus. This needs application downtime but no backup restore to grow the capacity for SAPData with the same performance capabilities as stated with the entry level configuration.

Disks	RAID	Capacity	Description
2 x 18.2 GB			Online spare
36 x 18.2 GB	0+1	327.6 GB	Oracle tablespace files SAPData
2 x 18.2 GB	1	18.2 GB	Online redo log
4 x 18.2 GB	0+1	36.4 GB	Archived log
4 x 18.2 GB	0+1	36.4 GB	Backup, check, reorg, stat, trace
2 x 18.2 GB	1	18.2 GB	Ora executables
2 x 18.2 GB	1	18.2 GB	SAP executables

Table 12, Oracle Medium Configuration

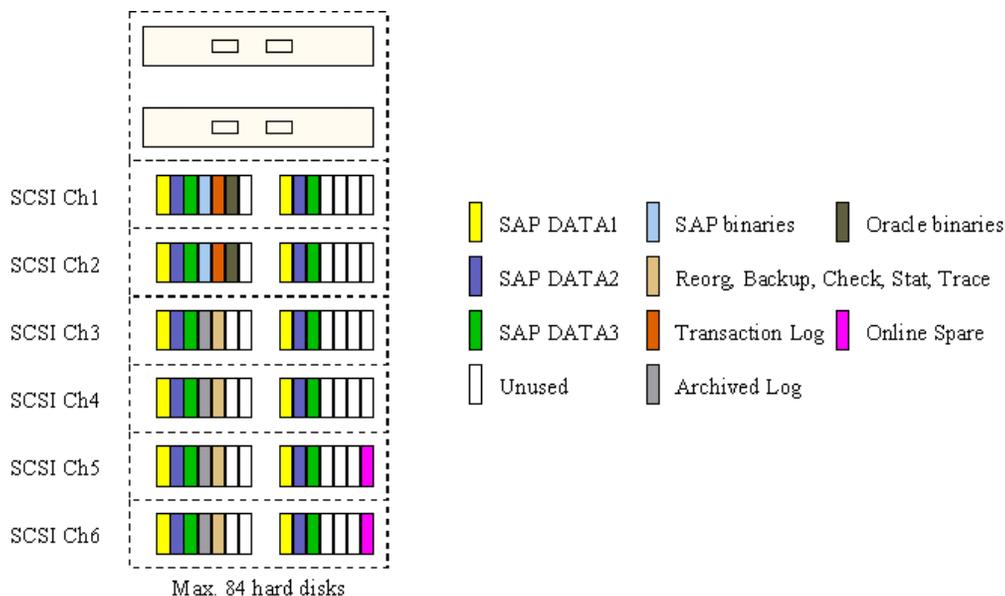


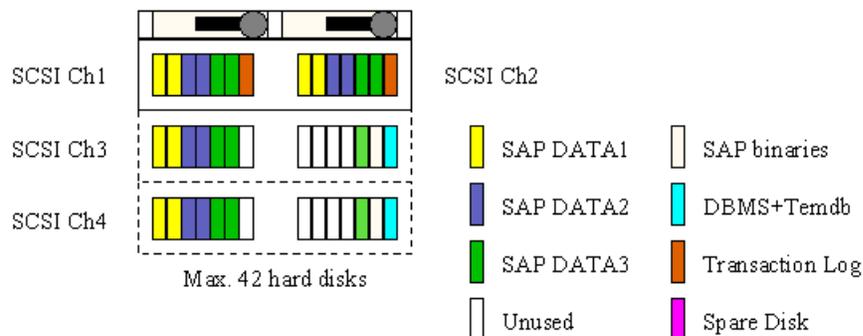
Figure 4, StorageWorks EMA12000, Oracle Medium

### SQL Server

The SQL Server medium configuration is the entry-level configuration plus two additional storage shelves, which activates SCSI channels 3 and 4 on the MSA 1000.

Disks	RAID	Capacity	Description
2 x 18.2 GB			Online spare
2 x 18.2 GB	1	18.2 GB	SAP binaries
2 x 18.2 GB	1	18.2 GB	DBMS+Tempdb
2 x 18.2 GB	1	18.2 GB	Transaction log
8 x 18.2 GB	0+1	72.8 GB	SAP Data1
8 x 18.2 GB	0+1	72.8 GB	SAP Data2
8 x 18.2 GB	0+1	72.8 GB	SAP Data3

**Table 13, SQL Server Medium Configuration**



**Figure 5, StorageWorks MSA1000 - SQL Server Medium**

Figure 5 illustrates an MSA1000 configuration providing about 220 GB for SAPData on three 8 x 18.2-GB-member RAID 0+1 sets.

## High-end Configuration

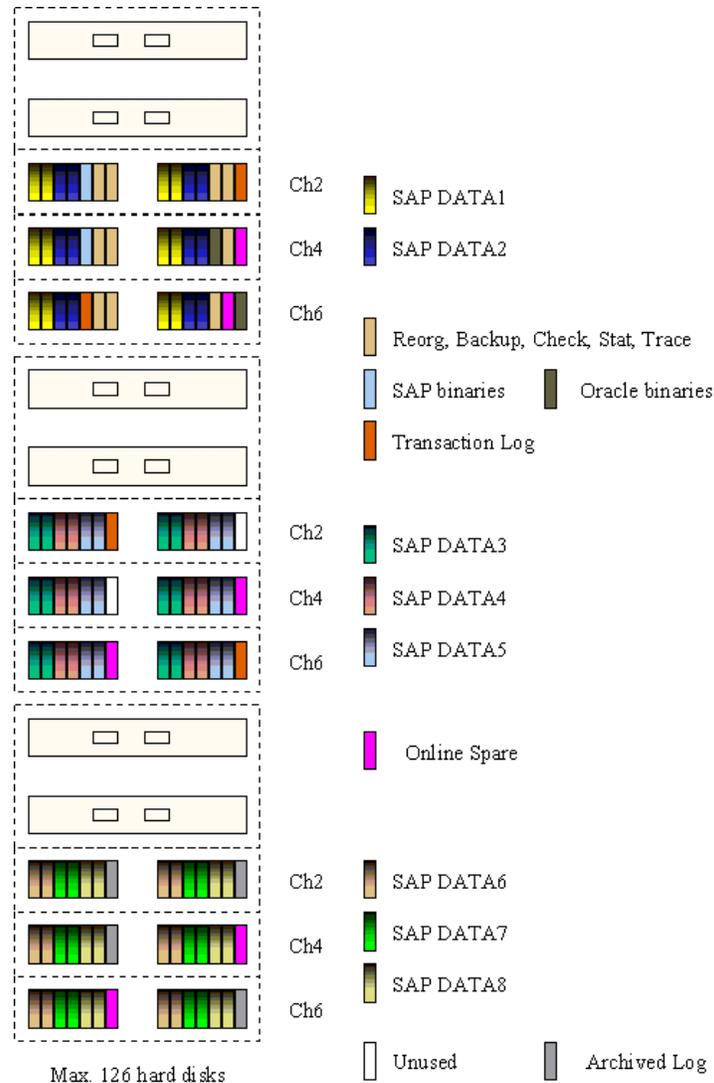
### Oracle

The two high-end configurations in Figure 6, based on three HSG80 controller pairs for performance optimization, and Figure 7, based on two HSG80 pairs for capacity optimization are configured in the EMA12000 product package.

Disks	RAID	Capacity	Description
6 x 18.2 GB			Online spare
96 x 18.2 GB	0+1	873.0 GB	Oracle tablespace files SAPData
4 x 18.2 GB	1	36.4 GB	Online redo log
4 x 18.2 GB	0+1	36.4 GB	Archived log
10 x 18.2 GB	0+1	91.0 GB	Backup, check, reorg, stat, trace
2 x 18.2 GB	1	18.2 GB	Ora executables
2 x 18.2 GB	1	18.2 GB	SAP executables

**Table 14, Oracle High-end Configurations**

Figure 6 illustrates an HSG80-based configuration for a nearly 1-terabyte configuration.



**Figure 6, StorageWorks EMA12000 - Oracle High-end Performance**

The main difference between the high-end performance configuration in Figure 6 and the high-end capacity configuration in Figure 7 is the performance advantage of the three controller pairs in a single 42U rack compared to the 44 free disk slots for a customer environment where the growth rate in terms of demand for capacity is not that fixed during the initial sizing process. This gives the opportunity to further grow the mySAP.com database without the costs of an additional HSG80 controller pair, but needs the floor space of two 42U racks initially.



Figure 7, StorageWorks EMA12000 - Oracle High-end Capacity