



SAP R/3 Enterprise on IBM TotalStorage DS4000 Storage Server

Automated FlashCopy Backup and Layout Considerations



Ralf Schröder (schroedr@de.ibm.com) – Wilhelm Gardt (willigardt@de.ibm.com)

IBM Systems Group – Advanced Technical Support

Sponsor: Eva Lau (elau@us.ibm.com)

Business Development, IBM TotalStorage, STG

Executive Summary

This White Paper describes the potential advantages of using the IBM® TotalStorage® DS4000 storage server in an SAP environment and explains opportunities for SAP customers to help increase the efficiency of their storage infrastructure while reducing costs at the same time. The following subjects are covered:

- **Tuning DS4000 for SAP systems**
- **Choosing the right DS4000 layout for your SAP system**
- **FlashCopy® backup solution for SAP systems on DS4000**
- **How to implement an integrated backup solution**

From the series of performance tests described below, we have concluded that the advanced functions of the IBM TotalStorage DS4000 Storage performed well in the SAP environment. The development efforts described in this White Paper can help clients:

- **Enhance SAP performance by tuning the DS4000**
- **Avoid performance bottlenecks and administrative overhead**
- **Decrease backup window and restore time**
- **Effectively use the storage capacity**
- **Reduce implementation complexity**

The scenarios described below address both the implementation phase of an SAP environment on DS4000 and the continued operation of a productive environment. They are intended to satisfy the need for information about SAP installations on DS4000 and to offer an immediate remedy for many performance problems and difficulty with time-consuming backups.

Table of Contents

EXECUTIVE SUMMARY	1
TABLE OF CONTENTS	2
FIGURES	4
TABLES	4
PREFACE	5
1 INTRODUCTION	6
2 HARDWARE AND SOFTWARE SETUP	9
2.1 LIST OF INSTALLED HARDWARE, SOFTWARE AND FIRMWARE	10
2.2 STORAGE AREA NETWORK ENVIRONMENT.....	11
2.3 BACKUP/RESTORE ENVIRONMENT.....	11
2.3.1 <i>IBM 3584 UltraScalable tape library</i>	11
2.3.2 <i>Tivoli Storage Manager</i>	12
2.4 SAP ENVIRONMENT	12
2.4.1 <i>DS4000 overview</i>	13
2.4.2 <i>SAP installation</i>	15
2.4.3 <i>IBM pSeries and logical partitioning</i>	15
2.4.4 <i>AIX logical volume manager configuration</i>	16
3 DS4000 DISK LAYOUT FOR SAP	19
3.1 TESTING THE PERFORMANCE	19
3.1.1 <i>The application view</i>	19
3.1.2 <i>The server view</i>	22
3.1.3 <i>The DS4000 view</i>	24
3.2 DS4000 PARAMETERS AND PERFORMANCE TUNING.....	26
3.2.1 <i>DS4000 write cache mirroring</i>	26
3.2.2 <i>DS4000 controller cache size and cache block size</i>	27
3.2.3 <i>DS4000 RAID level</i>	27
3.2.4 <i>DS4000 drive loop utilization</i>	28
3.2.5 <i>DS4000 array creation</i>	29
3.2.6 <i>DS4000 segment size</i>	30
3.2.7 <i>DS4000 read-ahead parameter</i>	30
3.2.8 <i>LUN mapping, ownership and multipathing</i>	31
3.3 DS4000 ARRAY LAYOUTS AND RESULTS	31
3.3.1 <i>The basic layout 01</i>	32
3.3.2 <i>The load-balanced basic layout 02</i>	34
3.3.3 <i>Advanced load-balanced layout 03</i>	36
3.3.4 <i>Log separation layout 04</i>	38
3.3.5 <i>Advanced load-balanced layout 04_2 with log separation</i>	39
3.3.6 <i>Load-balanced layout 04_3 with optimized log separation</i>	40
3.3.7 <i>The full-blown layout 05</i>	42
3.4 DS4000 LAYOUT FINDINGS AND RECOMMENDATIONS.....	45

4	AUTOMATED SAP BACKUP SOLUTION WITH DS4000 FLASHCOPY	49
4.1	DEVELOPING A BACKUP SOLUTION FOR SAP ENVIRONMENTS	51
4.1.1	<i>BR*Tools overview</i>	<i>51</i>
4.1.2	<i>Tivoli Storage Manager for ERP</i>	<i>52</i>
4.1.3	<i>DS4000 FlashCopy integration</i>	<i>53</i>
4.2	IMPLEMENTATION COOKBOOK – SINGLE IMPLEMENTATION STEPS WITH TROUBLESHOOTING HINTS....	54
4.2.1	<i>Install and configure the TSM client on the SAP server and backup server.....</i>	<i>55</i>
4.2.2	<i>Install and configure TSM for ERP on the SAP server</i>	<i>56</i>
4.2.3	<i>Create OS users on the backup server</i>	<i>57</i>
4.2.4	<i>Create network filesystems.....</i>	<i>58</i>
4.2.5	<i>Install and configure an Oracle client on the backup server.....</i>	<i>58</i>
4.2.6	<i>Copy Oracle client libraries.....</i>	<i>59</i>
4.2.7	<i>Set up password file authentication</i>	<i>60</i>
4.2.8	<i>Install and configure TSM for ERP on the backup server.....</i>	<i>61</i>
4.2.9	<i>Install BR*Tools.....</i>	<i>61</i>
4.2.10	<i>Create shell scripts to handle DS4000 FlashCopy with Storage Manager CLI.....</i>	<i>63</i>
4.2.11	<i>Create shell scripts to handle volume groups and filesystems</i>	<i>69</i>
4.2.12	<i>Install and configure sudo</i>	<i>70</i>
4.2.13	<i>Development of the SPLITINT script</i>	<i>71</i>
5	CONCLUSIONS AND OUTLOOK.....	76
6	REFERENCES.....	78
APPENDIX.....		80
A1	PERFORMANCE COLLECTION	80
A2	SPLITINT IMPLEMENTATION	81
NOTICES AND DISCLAIMERS		91

Figures

FIGURE 1: HARDWARE SETUP	9
FIGURE 2: CONNECTING DS4000 EXPANSION UNITS TO A DS4000 STORAGE SERVER	13
FIGURE 3: CUMULATIVE READ PERCENTAGE FOR FOUR CLIENT COPIES	20
FIGURE 4: SAP TRANSACTION SM37 – JOB OVERVIEW FOR EVALUATING CLIENT COPY COMPLETION TIME.....	21
FIGURE 5: INTER-DISK ALLOCATION AND THE ADEQUATE VOLUME LAYOUT ON DS4000 (PICTURE TAKEN FROM REFERENCE [20])	23
FIGURE 6: EXAMPLE OF THE DS4000 PERFORMANCE MONITOR	25
FIGURE 7: DOUBLING BANDWIDTH BY USING ALL AVAILABLE DRIVE MINI HUBS (PICTURE TAKEN FROM REFERENCE [20]).....	28
FIGURE 8: DISK LAYOUT FOR TEST SCENARIO 01.....	32
FIGURE 9: SIMPLIFIED DISK LAYOUT FOR SCENARIO 01	33
FIGURE 10: SIMPLIFIED DISK LAYOUT FOR SCENARIO 02	35
FIGURE 11: SIMPLIFIED DISK LAYOUT FOR SCENARIO 03	36
FIGURE 12: SIMPLIFIED DISK LAYOUT FOR SCENARIO 04	38
FIGURE 13: SIMPLIFIED DISK LAYOUT FOR SCENARIO 04_02.....	39
FIGURE 14: SIMPLIFIED DISK LAYOUT FOR SCENARIO 05	43
FIGURE 15: IO/S AVERAGES AND MAXIMUM FROM THE DS4000 PERFORMANCE MONITOR FOR ALL TEST LAYOUTS	46
FIGURE 16: KB/S AVERAGES AND MAXIMUM FROM THE DS4000 PERFORMANCE MONITOR FOR ALL TEST LAYOUTS	46
FIGURE 17: CUMULATIVE CLIENT COPY JOB RUNTIME FROM SAP TRANSACTION SM37 FOR ALL TEST LAYOUTS...47	
FIGURE 18: DATA FLOW FOR DATABASE BACKUP AND RESTORE	50
FIGURE 19: USING BR*TOOLS TO BACK UP ORACLE DATABASE OBJECTS	52
FIGURE 20: INTERACTION OF BR*TOOLS WITH TDP.....	53
FIGURE 21: BR*TOOLS AND DS4000 FLASHCOPY INTEGRATION	54
FIGURE 22: DS4000 STORAGE MANAGER GUI SHOWING ACTIVE FLASHCOPY PAIRS AND REPOSITORY DRIVES....	65
FIGURE 23: DS4000 STORAGE MANAGER GUI SHOWING LUN MAPPINGS	66
FIGURE 24: DS4000 STORAGE MANAGER GUI SHOWING DISABLED FLASHCOPY PAIRS AND REPOSITORY DRIVES	66
FIGURE 25: CONTROL FLOW AND COMMUNICATION CHANNELS OF THE SPLIT FUNCTION	73
FIGURE 26: CONTROL FLOW AND COMMUNICATION CHANNELS OF THE RESYNC FUNCTION	73

Tables

TABLE 1: PERFORMANCE VALUES FOR DISK LAYOUT 01	34
TABLE 2: PERFORMANCE VALUES FOR DISK LAYOUT 02.....	35
TABLE 3: PERFORMANCE VALUES FOR DISK LAYOUT 03.....	37
TABLE 4: PERFORMANCE VALUES FOR DISK LAYOUT 04.....	38
TABLE 5: PERFORMANCE VALUES FOR DISK LAYOUT 04_02	40
TABLE 6: PERFORMANCE VALUES FOR DISK LAYOUT 04_03	41
TABLE 7: PERFORMANCE VALUES FOR DISK LAYOUT 05	43
TABLE 8: SUMMARY OF RECOMMENDATIONS FOR DS4000 SETTINGS FOR SAP WITH ORACLE ON AIX	45
TABLE 9: SUMMARY OF CONCLUSIONS FOR DS4000 RAID ARRAY LAYOUT FOR SAP WITH ORACLE ON AIX	47

Preface

This White Paper was motivated by the observation that many customers who are using SAP as one of their most critical business applications sooner or later experience SAP system performance or backup window problems. These problems are typically related to constant data growth and a corresponding increase in time required for backups. The paper is intended to help the technical community not only in an SAP planning and installation phase, but also help to improve overall system and backup performance after installation.

Our analysis of customers' performance problems revealed that many customers do not spend enough time planning and administering the disk layout on which the SAP system runs. Given the amount of work involved in planning, installing and maintaining an SAP system, it is understandable that factors like disk layout do not become the highest priority, especially when there are that many different sources, including SAP OSS notes, Oracle instructions and other white papers that are intended to help with performance and installation tips.

By contrast, most companies that run SAP focus closely on backup solutions to protect the data they store in the system. Large enterprise companies, can use a variety of good products to help automate the backup process and reduce the amount of time required; however, there are not as many solutions for mid-sized companies, which have similar needs.

To address these requirements, this White Paper provides tools for developing a customized backup strategy using the advanced IBM TotalStorage DS4000 storage server copy services. By offering working example scripts, the hints and tips sections below can help customers or IBM service people to tailor this concept to similar environments.

In addition to recommending backup strategies, the White Paper summarizes a number of layout tips from external sources and verifies the influence of some parameters that are specific to the IBM TotalStorage DS4000 storage server. This information is intended to help administrators find the golden mean between performance gains and disk layout administration efforts.

Many thanks to

Gero Schmidt (IBM Systems Group – Advanced Technical Support)

for his valuable contribution to the performance section of this White Paper.

1 Introduction

One of the greatest challenges associated with installations of the mySAP Business Suite is the exponential growth of data--and consequently its secure backup to tape. Although tape techniques constantly evolve to higher capacity and throughput capabilities, many nevertheless cannot keep up with the need to backup entire mySAP Business Suite databases in an increasingly shorter window of time.

Backing up business-critical data in order to be able to quickly react either to logical database errors or physical errors like hardware failures or disasters is key for faster restoration times. Accordingly, SAP recommends that customers perform backups at least daily online and once a week offline.

The problem is that common online backup techniques require more time and can impede overall system performance of the mySAP Business Suite. And offline backups are seldom possible in a company that operates 24/7. It is clear that the industry requires a new backup strategy.

Periods of low mySAP Business Suite system activity are few and far between because companies tend to run long batch jobs on weekends or during night hours. Since costs can increase dramatically during system downtime, companies need quicker system restoration times in the event of a mySAP Business Suite system failure.

With these factors in mind, IBM equipped its TotalStorage subsystem family with powerful hardware copy services that can be used to accelerate mySAP Business Suite system backup tasks using, for example the TotalStorage Enterprise Storage Server (ESS) FlashCopy option.

For a complete list of available software products, please refer to the Tivoli URL at

www-306.ibm.com/software/tivoli/.

In today's environment, several well-proven mySAP Business Suite backup concepts use IBM TotalStorage subsystem copy services, which provide administrators the technical to implement these concepts beyond the ESS. This White Paper offers guidance on how to exploit the IBM TotalStorage DS4000 storage server functions in conjunction with mySAP Business Suite installations.

Furthermore, it stresses the awareness regarding the importance of planning the disk layout for SAP R/3 installations on DS4000. Although most companies carefully plan their SAP

software rollout strategy and its impacts for the business together with SAP, many neglect the importance of the underlying IT infrastructure. The infrastructure also affects the business in terms of employee frustration or system availability.

Running SAP applications on IBM TotalStorage DS4000 storage servers and optimizing the layout can help companies maximize the high availability and performance capabilities of the DS4000.

This White Paper concentrates on SAP system performance and how it can be optimized by tuning DS4000 parameters, changing DS4000 disk layouts and adhering to the logical layout recommendations that come from SAP. High availability is one of the key capabilities of DS4000.

SAP system performance is a complex subject, even if only considered from a storage subsystem's point of view. One must consider performance from several different points of view: software, server, storage, etc. At the same time, TCO (Total Cost of Ownership) and maintenance efforts have to be taken into account, and these considerations sometimes conflict with the goal of applying all available performance tuning efforts.

This White Paper presents a summary of several storage subsystem performance tips. Some tips are explained in more detail than others. For example, we discuss the effective usage of disk capacity and the efforts to keep the expected performance at the preferred level. Rather than present detailed performance analysis, we have chosen to include practical hints and tips on how to improve and plan performance for SAP applications.

Inside, we discuss two topics:

- 1.) IBM TotalStorage DS4000 storage server volume layout for mySAP ERP
- 2.) Implementation of a mySAP ERP FlashCopy backup solution with DS4000

The goals of Topic 1 is to provide recommendations on how to implement mySAP ERP on an IBM TotalStorage DS4000 storage server in terms of volume layout, performance and DS4000 parameter settings. In order to find a reasonable ratio of layout complexity and SAP system performance versus administration efforts, several possible layouts have been tested. Observations in terms of performance, high availability and response/transaction execution times have been documented.

Topic 2 provides instructions about the manual implementation of a DS4000FlashCopy for DS4000 backup solution using shell scripts, as well as the required theory on several backup concepts. Although the scripts have been developed in a stand-alone SAP R/3 environment, the concept may also be adapted to more complex mySAP Business Suite scenarios.

Chapter 2 introduces the hardware and software setup that has been used during the project. It also provides some basic knowledge about the components and how they interoperate with each other. This lays the groundwork for Chapter 3, which completely covers Topic 1 by discussing DS4000 volume layout and performance considerations for SAP instances.

Next, Chapter 4 provides the theory and single cookbook-like installations steps for developing an automated FlashCopy backup strategy for SAP R/3 that follows all SAP rules and can therefore be certified by SAP. Finally, Chapter 5 summarizes the project and presents some ideas on further development and improvement tasks. The White Paper concludes with references and an appendix with developed scripts.

2 Hardware and software setup

This chapter describes the components of the setup and provides references for further reading. This White Paper focuses on DS4000 performance recommendations and the presentation of a backup solution. It does not cover hardware and software in detail.

The following illustration shows the hardware setup on which all the technical work for this test was done.

The complete setup was used for the development of the backup solution. Performance measurements for different DS4000 logical volume layouts were run on the SAP server and DS4500 storage server only.

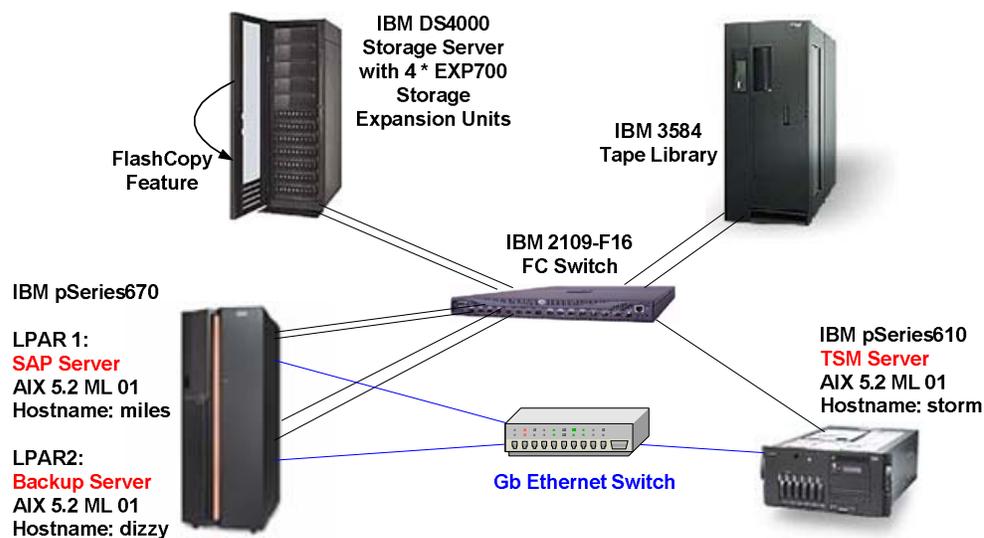


Figure 1: Hardware setup

Remark:

Throughout this White Paper, many configuration and command examples are found. The hostnames used in these examples are:

- *miles*
- *dizzy*
- *storm*

The next chapter contains a list of hardware products, firmware releases and software versions, as well as descriptions for each component and their usage for this project.

2.1 List of installed hardware, software and firmware

Disk storage subsystem

IBM TotalStorage DS4000DS4500

4 * IBM TotalStorage Storage Enclosure EXP 700

Firmware version 05.30.12.00

NVSRAM version N1742F900R830V04

SAN environment

IBM TotalStorage SAN Switch F16

Kernel Version 5.4

Fabric OS Version 3.1.1

SAP server

IBM eServer™ pSeries® 670 LPAR with 4GB RAM and two to four 1.1-GHz POWER4 CPUs

AIX® 5.2 ML 01 64-bit

DS4000 Storage Manager Runtime 8.40

DS4000 Storage Manager Client 8.41

SAP R/3 Enterprise 4.70 IDES

Oracle9i Enterprise Edition Release 9.2.0.4.0 64-bit

Tivoli Storage Manager (TSM) Backup/Archive Client 5.2.2.0

TSM for ERP Systems 3.3.2 for AIX LF 64-bit (also called Data Protection for mySAP.com)

Backup server

IBM pSeries 670 LPAR with 4GB RAM and two 1.1-GHz POWER4 CPUs

AIX 5.2 ML 01 64-bit

DS4000 Storage Manager Runtime 8.40

DS4000 Storage Manager Client 8.41

Oracle Client Release 9.2.0.1

SAP BR*Tools Version 6.20 patch level 124 (Patch level 110 or higher is required.)

Tivoli Storage Manager (TSM) Backup/Archive Client 5.2.2.0

TSM for ERP Systems 3.3.2 for AIX LF 64-bit (also called Data Protection for mySAP.com)
sudo (a free software tool from Courtesan Consulting that allows a UNIX® system administrator to give certain users the ability to run commands as root)

TSM server

IBM pSeries 610 with 1GB RAM and two 375-MHz POWER3 CPUs

AIX 5.2 ML 01 64-bit

IBM AIX Enhanced Tape and Medium Changer Device Driver 8.4.1.0

Tivoli Storage Manager Server 5.2.2.0

Tape library:

IBM 3584 UltraScalable Tape Library with two IBM Ultrium® Tape Drives (LTO®-1)

2.2 Storage area network environment

Servers and storage systems were connected in a Storage Area Network (SAN) environment, which allowed simultaneous storage access. The core of the SAN was an IBM Fibre Channel Switch 2109 with 16 ports. All connections were running at a 2-Gbps speed. See References [1] and [2].

2.3 Backup/restore environment

IBM Tivoli Storage Manager (TSM) Server was installed on an IBM pSeries server running AIX operating system to create data backups and restores. An IBM 3584 Tape Library provided backup storage. Backup/restore clients could access the TSM Server over Gigabit Ethernet. Each of the single components is described below.

2.3.1 IBM 3584 UltraScalable tape library

The IBM 3584 UltraScalable Tape Library is a stand-alone device designed to provide reliable, automated tape handling and storage for unattended mid-range systems and network servers. The basic library is a single storage unit known as the base frame. The library's scalability allows increasing capacity by adding storage units, called expansion frames.

The sharing of a 3584 Tape Library is possible by partitioning the library's storage slots and tape drives into logical libraries. Servers can then run separate applications for each logical library.

The library used for this backup environment contained IBM Ultrium Tape Drives and Linear Tape Open™ (LTO-1 and LTO-2) Ultrium media.

IBM provides device drivers for tape drives and robotics to connect servers to the 3584 Tape Library. Supported operating systems are AIX, HP-UX, Solaris, Linux™, Windows® and more.

Installation steps overview:

- Create a logical library partition.
- Install tape and robotics device drivers on the attached server.
- Connect tape drives to the host system.

2.3.2 Tivoli Storage Manager

Tivoli Storage Manager (TSM) is designed to provide a solution for distributed data and storage management in an enterprise network environment (LAN, WAN and SAN infrastructures). Together with complementary products, TSM is designed to deliver many data management functions, such as data protection for file and application data, record retention, space management and disaster recovery. TSM can be integrated with numerous storage devices and supports more than 30 operating system platforms. See References [9] and [10].

In our TSM environment, we used TSM's basic backup and restore features along with "Tivoli Storage Manager for Enterprise Resource Planning", which acts as an interface to SAP's backup tools (BR*Tools). For an introduction to BR*Tools see section 4.1.1.

Installation steps overview:

- Install Tivoli Storage Manager server on the host.
- Set up TSM configuration:
 - Tape library and associated objects (storage pool, policy set, management class and so on).
 - Hosts to be backed up (TSM clients): SAP server, backup server.

2.4 SAP environment

SAP and Oracle Software were installed on a logical partition of an IBM pSeries 670 running the AIX operating system. The Oracle database was located on logical drives that were created on an IBM DS4000 storage server.

2.4.1 DS4000 overview

IBM TotalStorage DS4000 is a redundant array of independent disks (RAID) storage subsystem that contains the Fibre Channel (FC) interfaces to connect both the host systems and the disk drive enclosures. (Starting in 2004, Serial-ATA disks can also be attached to DS4000 controllers.) The DS4000 Storage Server is designed to provide high system availability using many hot-swappable and redundant components: dual hot swap RAID controllers, power supplies and more.

DS4000 storage is provided by adding up to 16 IBM DS4000 EXP700 expansion units. Each expansion enclosure supports up to fourteen 2Gbps Fibre channel disk drives. DS4000 connection to the disks is established through up to four mini-hubs (two per controller) and fibre channel cabling. A maximum of four mini-hubs is also supported to connect hosts to the DS4000 (direct attach or SAN fabric). See References [3] and [4].

The next figure shows a storage subsystem containing one storage server and two redundant drive loops with four expansion units each.

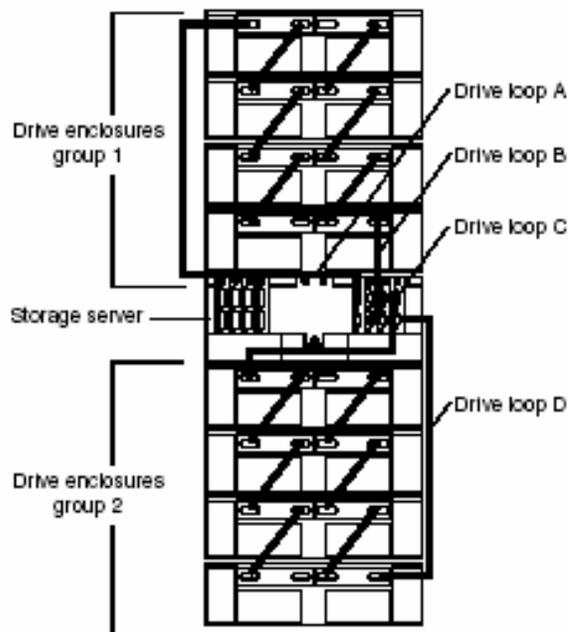


Figure 2: Connecting DS4000 expansion units to a DS4000 storage server

DS4000 multi-path driver software

For DS4000 attachment, most operating systems require a specific multi-path driver: Redundant Dual Active Controller (RDAC), which is delivered on CD-ROM together with the DS4000 subsystem. The CD-ROM also contains a utility called *hot_add*, which allows dynamic configuration of new logical drives to the operating system.

DS4000 Storage Manager and copy functions

The IBM DS4000 Storage Manager software is used to configure, manage and troubleshoot the DS4000 Storage Server. Its front-end component is a graphical user interface (GUI) called Storage Manager Client. It is a Java™-based management tool that is available for various operating systems.

Three copy functions are available with DS4000:

- FlashCopy is designed to create a point-in-time image of a logical drive.
- VolumeCopy is designed to create a complete physical replication of a logical drive to another within the same storage subsystem.
- Remote Volume Mirroring is designed to replicate data between storage subsystems over a remote distance.

The two logical drives that form a copy pair are called “source volume” (or “base logical drive”) and “target volume.”

Administration and copy tasks for the DS4000 can be initiated with Storage Manager Client or a command-line interface called Storage Manager CLI (SMcli). Naturally, access restrictions can be activated on a DS4000. If a password has been set on the DS4000, the password must also be supplied in the CLI command line.

Within this project, no access control has been implemented. Details on DS4000 CLI syntax can be found in Reference [13].

Example:

To display the logical drives that are defined on a DS4000 storage server, simply enter:

```
SMcli 9.155.49.96 -c "show logicalDrives;"
```

where “9.155.49.96” is the IP address of one of two DS4000 controllers.

Installation steps overview:

- Set up DS4000 (firmware/software installation or upgrade, IP configuration).
- Connect DS4000 Expansion Units.
- Install DS4000 Storage Manager Software (incl. RDAC) on hosts.
- Create DS4000 RAID arrays and logical drives.
- Connect DS4000 to host systems:
 - An AIX server must be connected with one to four host bus adapters (HBA) to both DS4000 controllers.
 - For each AIX server, we configured two dedicated input/output paths in the SAN environment: the first HBA was connected to DS4000 controller A, the second HBA to DS4000 controller B.

2.4.2 SAP installation

SAP R/3 and Oracle Enterprise software were installed on the internal system disk of the AIX server. The Oracle data files and the log files were created on separate logical drives on the DS4000 storage subsystem.

Please see the SAP and Oracle manuals (References [16] and [17]) if you are interested in installation details.

The terms “SAP R/3” or “mySAP.com” are sometimes used to represent a wide variety of SAP products (which is not quite correct). To be precise at this point, this project has been realized on an IDES SAP R/3 Enterprise Core 4.70, Support Release 1, Extension Set 1.10 application.

IDES (Internet Demonstration and Evaluation System) demo landscapes are preconfigured installations of mySAP Business Suite components, populated with the data and business processes of a fictional enterprise.

2.4.3 IBM pSeries and logical partitioning

Beginning with AIX 5.1, the AIX operating system supports partitioned environments. A logical partition (LPAR) is the division of a computer’s processors, memory and hardware resources into multiple environments so that each environment can be operated independently with its own operating system and applications. See References [6] and [7] for more details.

For our setup, the SAP server and backup server were both installed on logical partitions of an IBM pSeries 670 with 1.1 GHz POWER4™ processors.

With the LPAR feature, we were flexible regarding CPU and memory resources of the AIX servers:

- During the performance tests, the SAP server used four CPUs and 4GB of RAM.
- For the development of the backup solution, the resources of the SAP server were reduced to two CPUs and 4GB of RAM. The backup server was provided identically.

2.4.4 AIX logical volume manager configuration

Almost all DS4000 disk layouts that are described in this White Paper have been designed to have an even number of logical drives for both Oracle data files (sapdata) and online redo logs (saplog). Half of the “sapdata drives” were assigned to DS4000 controller A, the other half to controller B. Thus, input/output load was distributed over the two controllers (“static” load balancing). The logical drives for the online redo logs (saplog) were treated the same way. The single layouts are described in detail in Chapter 3.

On AIX, DS4000 logical drives are recognized as physical volumes (hdisks). The volumes that were distributed over the input/output paths on the DS4000 were combined in one volume group. Then the partitions of a logical volume were allocated alternately on all available physical volumes that were supplied for a specific logical AIX volume (physical partition striping or PP striping). Oracle data files and online redo logs were created on separate physical disks.

With the AIX system management tools *smitty*, PP striping is specified in the Add a Logical Volume menu by choosing the value “maximum” in the field “RANGE of physical volumes.” In an AIX command line, the option `-e x` must be set to choose PP striping.

On AIX, the utility *fget_config*, which is installed with the DS4000 client software, lists the DS4000 volumes and the preferred DS4000 controller for each volume (see Reference [5]). In AIX, the DS4000 objects are represented by the associated device filenames:

DS4000 disk arrays → dar*

Controllers → dac*

Logical drives → hdisk*

The following example shows two DS4000 logical drives that are assigned to different controllers. It also describes the AIX operation system view of these logical drives and the AIX administration commands needed to create a volume group and a logical volume.

DS4000 Logical Drives:

Logical Drive name: SAPDATA-1

Preferred owner: Controller in slot A

Capacity: 5 GB

Logical Drive name: SAPDATA-2

Preferred owner: Controller in slot B

Capacity: 5 GB

Mappings Storage Partitioning:

Logical Drive-To-LUN Mappings

Logical Drive SAPDATA-1, LUN 2, Host dizzy

Logical Drive SAPDATA-2, LUN 3, Host dizzy

AIX Logical Volume Manager Definitions:

list disk devices

> lsdev -Cc disk

hdisk0 Available 1S-08-00-5,0 16 Bit LVD SCSI Disk Drive

hdisk1 Available 1H-08-01 1742 (700) Disk Array Device

hdisk2 Available 1V-08-01 1742 (700) Disk Array Device

list DS4000 disk configuration for this server

> fget_config -l dar0

dac0 ACTIVE dac1 ACTIVE

hdisk1 dac0

hdisk2 dac1

create a volume group

> mkvg -y sapvg -s 32 hdisk1 hdisk2

0516-1254 mkvg: Changing the PVID in the ODM.

0516-1254 mkvg: Changing the PVID in the ODM.

sapvg

create a logical volume with “PP striping”

```
> mklv -y orasapdata1 -t jfs2 -e x sapvg 20
```

orasapdata1

lists physical volume distribution for logical volume

```
> lslv -l orasapdata1
```

orasapdata1:/oracle/PRO/sapdata1

PV	COPIES	IN BAND	DISTRIBUTION
hdisk1	010:000:000	100%	000:010:000:000:000
hdisk2	010:000:000	100%	000:010:000:000:000

Having introduced the setup and usage of each component, as well as the basic configuration rules that applied for this project, we move in the next chapter to the discussion of optimal disk layouts and how to verify the influence of the layout on SAP system performance.

3 DS4000 disk layout for SAP

The easiest advice for optimal DS4000 performance is to use as many disks as possible and to tune the DS4000 storage server daily. But for many obvious reasons, this is not a realistic strategy. Hence, the goal is to find the golden mean between costs and performance.

This chapter introduces the test technique and discusses how the results will be interpreted and how they can be transferred to operating systems other than AIX.

Afterwards, the operating system view and the associated filesystems are briefly discussed.

Section 3.2 is dedicated to the DS4000 parameters and features that influence the performance of an SAP R/3 system. It introduces several recommendations and explains the aspects that have been examined within this White Paper and for what reason.

Finally, the chosen disk layouts are described in section 3.3, followed by the results that have been achieved with each specific layout. The last section concludes and summarizes the performance section of this White Paper and tries to provide recommendations for running and administering SAP on DS4000.

3.1 Testing the performance

The performance of an SAP system that runs on a DS4000 storage subsystem can be monitored from several perspectives: the application view, the server view and the storage view. Each of these perspectives is critical to estimating the overall system performance.

Therefore, whenever an installed system performs slower than expected, a consultant has to review all three areas to determine the bottleneck (e.g., the runtime of an SAP report does not give any hint on a storage bottleneck or vice versa). In order to help avoid potential bottlenecks and avoid the high cost of analyzing a complete SAP system or landscape, companies need to develop an accurate plan for optimizing the performance of each area.

3.1.1 The application view

The typical workload of an SAP R/3 core component as an OLTP (Online Transaction Processing) application is characterized as almost 100 percent random accesses with a read quota of ~60 to 70 percent. Running SAP on Oracle database, the average transfer request size is by default 8KB.

To simulate a comparable workload with the installed SAP R/3 IDES system a suite of four simultaneously issued client copies was used for several reasons:

- a) A high volume of traffic is generated since the SAP IDES represents a typical small company with several clients representing different business locations, including Frankfurt, Barcelona, Porto and Dallas.
- b) Official benchmarks like the SD benchmark from SAP do not generate enough I/O to stress the subsystem but only to stress the server.
- c) Client copy administration is simple.
- d) Four background jobs generate a high volume of I/O to the subsystem without carrying the server to its performance limits.

Figure 3 demonstrates the progression of the read percentage as monitored with the DS4000 performance monitor (for details see section 3.1.3) for a specific test over the period of approximately 5 hours.

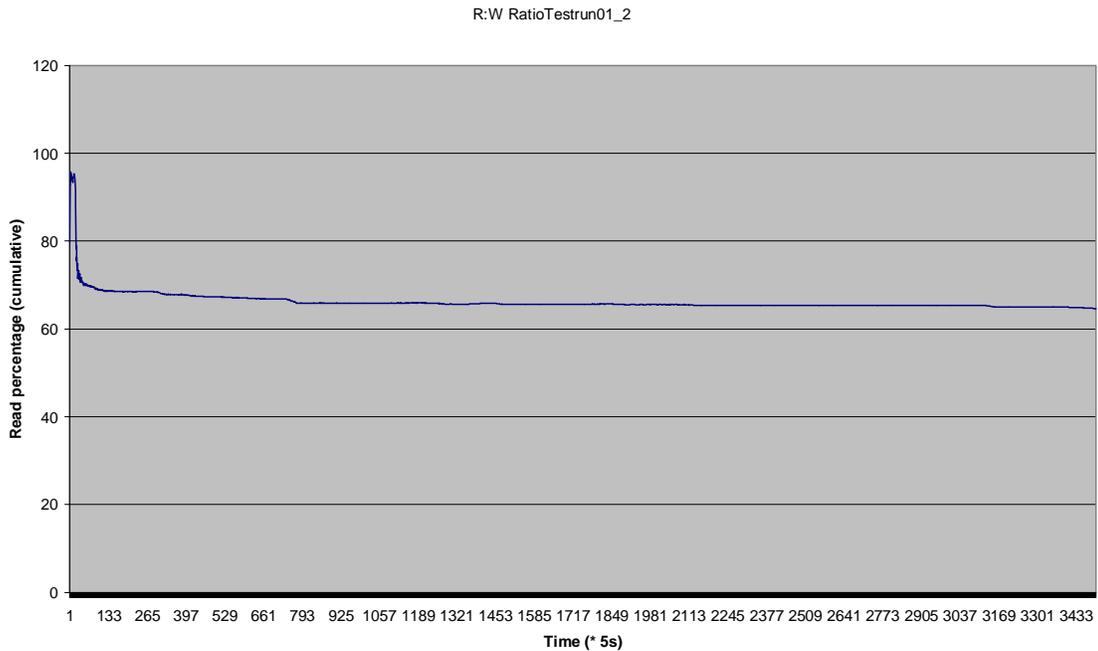


Figure 3: Cumulative read percentage for four client copies

When starting a new task, the initial database state was restored. For this test, one client copy after another was initiated as an immediately started background job with SAP transaction SCCL. Using transactions SCC3 and SM37 the progress and finally the total running time can be monitored for each single job.

Therefore, the reference performance parameter from the application point of view is defined as the sum of the completion time for all four client copy background jobs as displayed in Figure 4.

Job Overview

Job overview from: 24.06.2004 at :: to: 24.06.2004 at ::
 Selected job names: * Selected user names: SAP*
 Scheduled Released Ready Active Finished Canceled
 Event controlled Event ID:
 ABAP program Program name :

Job	Ln	Job CreatedBy	Status	Start date	Start time	Σ Duration(...)	Σ Delay (se...)
CLIENTCOPY 800 -> 900		SAP*	Finished	24.06.2004	09:55:16	21.947	0
CLIENTCOPY 810 -> 910		SAP*	Finished	24.06.2004	09:56:01	16.527	0
CLIENTCOPY 811 -> 911		SAP*	Finished	24.06.2004	09:57:41	20.122	0
CLIENTCOPY 812 -> 912		SAP*	Finished	24.06.2004	09:58:12	19.440	0
SAP_PHL_PERF		SAP*	Released			0	0
SAP_PHL_PERF		SAP*	Finished	24.06.2004	09:44:53	0	2.404
SAP_PHL_PERF		SAP*	Finished	24.06.2004	09:54:53	0	4
SAP_PHL_PERF		SAP*	Finished	24.06.2004	15:36:54	1	19.925
SAP_PHL_PERF		SAP*	Finished	24.06.2004	15:53:53	1	544
SAP_PHL_PERF		SAP*	Finished	24.06.2004	15:57:53	2	184
						78.040	23.061

Figure 4: SAP transaction SM37 – job overview for evaluating client copy completion time

It is important to note that client copies generate no typical I/O load compared with an SAP R/3 production database of 40GB to 60GB, as within only five hours the database grows immensely. Thus, many redo logs will be written and specific SAP filesystems (such as /saparch and /sapdata) grow to more than four times their initial capacity. To avoid test interruptions due to full filesystems, a shell script periodically deleted the archived redo logs (copied to the /saparch directory). Another option would have been to run the Oracle database in NOARCHIVELOG mode, but at that time, the I/O profile would no longer reflect a production R/3 database.

Although client copies are a common way of I/O generation, there are some trade-offs:

- Client copies generate more writes than normal SAP R/3 transactions do
- Simultaneous client copies interfere with each other since they access common cross-client tables. → The runtime of a single client copy does not allow any conclusion regarding the system or subsystem performance.
- The number of operating system processes that are granted to a single client copy varies. → Again, analyzing the runtime of a single client copy does not allow a valid comparison of different test runs.
- Client copies do not represent a normalized benchmark. → I/O and throughput values do not depict the maximum capability of the DS4000

Despite these facts, the total completion time of all client copies together allows comparing the subsystem performance of different test runs.

3.1.2 The server view

From the viewpoint of server performance, there are several ways to tune the operating system AIX, as well as the JFS/JFS2 filesystems, which are not within the scope of this White Paper. For pSeries and AIX performance tuning, consult specific Redbooks such as Reference [18].

However, two topics will be of interest when evaluating a subsystem's performance for the SAP R/3 Enterprise application:

- the LVM (Logical Volume Manager) layout and filesystems
- the monitoring of the performance and the disk usage

3.1.2.1 AIX LVM and filesystem layout

Regarding the filesystems, the layout for the most part is determined by SAP and Oracle (see References [16] and [17]).

According to the SAP installation guide on Oracle (Reference [16]), distributing archive files, redo log files and sapdata files to separate disks is highly recommended. This applies especially to the redo logs, which are written synchronously and cause more I/O activity than any other database files.

Recommendation: Distribute sapdata and redo log files to separate disks for performance and FlashCopy backup reasons. Do not FlashCopy SAP data and logs at the same time!

Certainly, the prerequisites to distributing data and logs must also be made at the DS4000 and LVM level. While DS4000 will be discussed in the next chapter, for AIX LVM it is important to state that because DS4000 systems have their own RAID arrays and logical volumes, we do not work with real physical disks in the system. Functions such as intra-disk allocation, write scheduling and write-verify policies, do not help much, and it is hard to determine the performance benefits when using them. They should be used only after testing. Misusing these functions can lead to degraded performance.

At the same time, we should not forget about the important inter-disk allocation policy. The inter-disk allocation policy is used to specify the number of disks that contain the physical partitions of a logical volume. The physical partitions for a given logical volume can reside on one or several disks in the same volume group, depending on the setting of the range option. By setting the inter-physical volume allocation policy to maximum, we also provide that the reads and writes are shared among PVs (physical volumes) and, in systems like DS4000, also

among controllers and communication paths. If systems are using only one big volume, it is owned by one controller, and all the traffic goes through only one path. This happens because DS4000 uses static load balancing.

Figure 5 illustrates the required setup on DS4000 for utilizing the inter-disk policy:

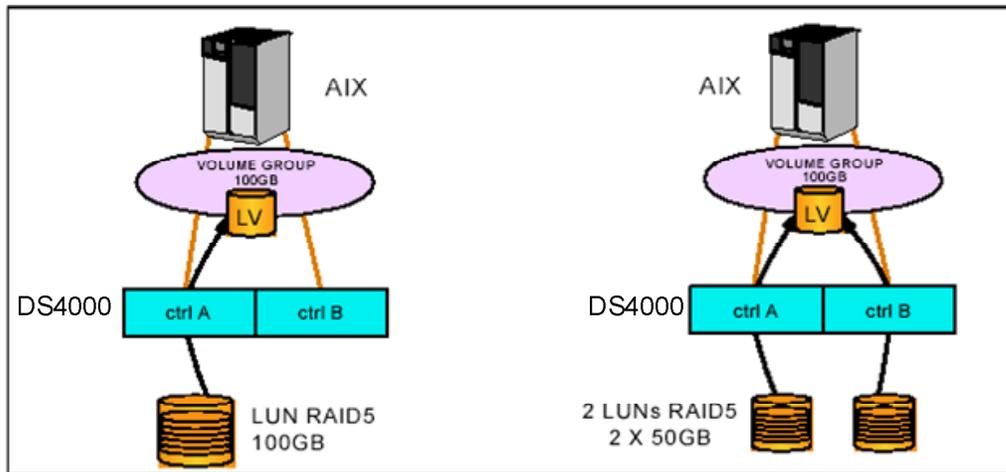


Figure 5: Inter-disk allocation and the adequate volume layout on DS4000 (picture taken from Reference [20])

Recommendation: When using AIX LVM, use an inter-disk allocation policy of maximum in order to spread the physical partitions (PP) of the logical volume across as many physical disks, controllers and communication paths as possible. (This is also referred to as PP striping.)

3.1.2.2 Performance monitoring

Monitoring the performance from the operating system point of view helps allow the quick determination of a bottleneck by using the standard UNIX *iostat* command.

For each test run, the *iostat* measured data for the SAP identified volumes for every interval of 5 seconds with the command

```
iostat hdiskx hdisky... 5.
```

The output has been saved to disk for later analysis. The important parameters to watch out for are:

- *%idle*: shows the percentage of time that the CPU or CPUs were idle and the system did not have an outstanding disk I/O request
- *%iowait*: shows the percentage of time that the CPU or CPUs were idle during which the system had an outstanding disk I/O request
- *%tm_act*: indicates the percentage of time the physical disk was active (bandwidth utilization for the drive)

```
tty:      tin          tout   avg-cpu:  % user   % sys    % idle   %iowait
          1.2          0.0           26.9    28.9     8.7     35.6
```

```
Disks:    % tm_act    Kbps    tps    Kb_read  Kb_wrtn
hdisk4    90.2       17501.6  1248.2  5916     81548
hdisk5    53.2       18697.3  1337.1  25024    68416
```

This example of one *iostat* output of the database server indicates that the system is quite busy executing disk I/O because the overall CPU idle time (for all four CPUs) is very low and the *iowait* percentage is quite high.

Looking at the disk activity time reveals that both hdisks are very busy. (In this example, the sapdata files resided on hdisk4, the log files on hdisk5.) A *tm_act* value of more than 80 percent points to a potential disk bottleneck and explains the system's *iowait* state.

From the *iostat* output, the I/O per second can be monitored, but for this purpose the DS4000 performance monitoring is used, which will be explained in the next chapter.

3.1.3 The DS4000 view

DS4000 performance can be monitored in two ways: either online with the DS4000 performance monitor or by automatically collecting the performance information using a DS4000 SMcli script (see Appendix A1).

Using the performance monitor, it is reasonable to select only the DS4000 disks whose performance values are of interest. Moreover, the subsystem totals and the single-controller throughput may be used to estimate the overall throughput and the load distribution.

Those selections can be made by clicking the Settings button in the Performance Monitor view, as Figure 6 illustrates.

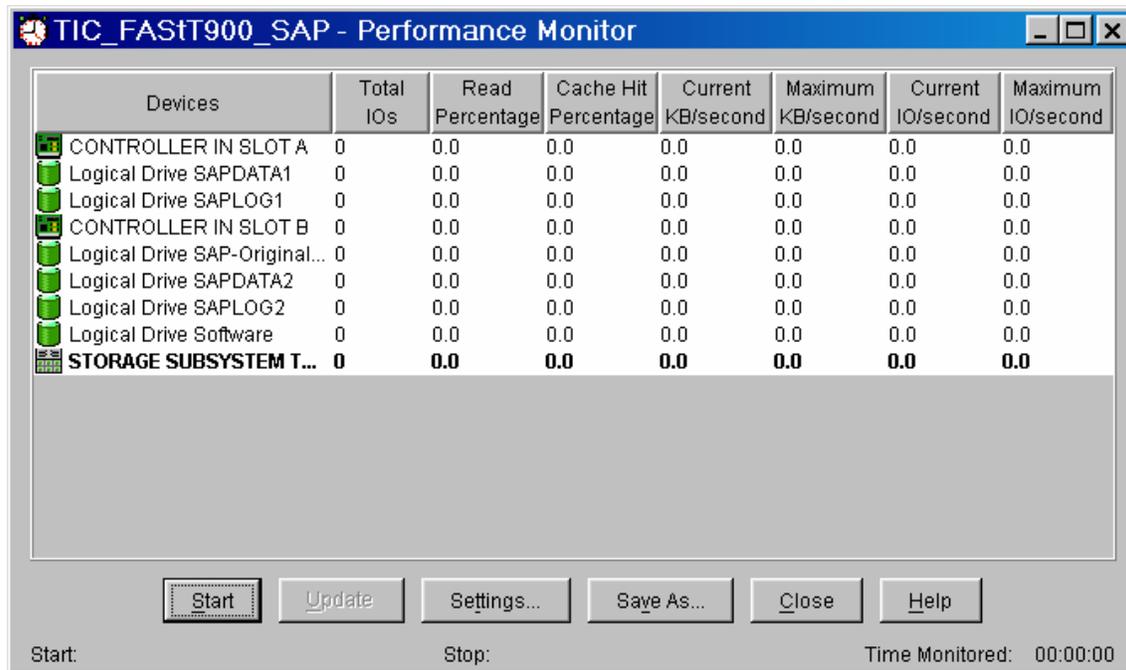


Figure 6: Example of the DS4000 performance monitor

In the Settings menu, it's also possible to change the polling interval, which means that each time the polling interval elapses, the performance monitor re-queries the storage subsystem and updates the statistics in the table.

Using the performance monitor to collect performance data can affect the normal storage subsystem performance, depending on the polling interval that is set. If the monitored storage subsystem begins in or transitions to an unresponsive state, an informational dialog box opens, stating that the performance monitor cannot poll the storage subsystem for performance data.

Since this testing was not of maximum DS4000 performance measurements, the polling interval was set to 5 seconds. Although this value may slightly decrease the maximum throughput of the DS4000, it allows acceptable comparisons of the iostat data from the server point of view and the DS4000 performance data as well as detailed monitoring of the application's I/O behavior.

As stated before, each test consisted of four client copies resulting in a total runtime of four to five hours, which made online monitoring of the DS4000 performance useless. Therefore, SMcli was used to collect the performance data for 3,500 iterations, or almost five hours.

The scripts that were used to collect the required performance data are listed in Appendix A1.

From the perspective of the DS4000 storage server, the key performance parameters that have been mainly analyzed were the storage subsystem totals (here mainly the progression of

current I/O and current KB/s), indicating how one DS4000 disk layout performs compared with another. For the throughput to a single disk, the performance was monitored from the server point of view to verify whether any specific disk tended to be a bottleneck. Load distribution to both DS4000 controllers was done manually with the diversity of tested DS4000 disk layouts.

It should be said again that the objective was not to test DS4000 or SAP performance but rather to compare the influence of several DS4000 disk layouts and parameters on overall SAP performance. Therefore, the next chapter introduces the relevant DS4000 parameters and discusses how those may influence performance.

3.2 DS4000 parameters and performance tuning

Besides the influence of specific disk layouts on performance, which will be discussed in the next chapter, a variety of DS4000 cache parameters and settings affect either the overall DS4000 performance or the performance of a single disk.

Please note that the parameters were neither changed nor tuned during the test, as their influence on performance is already very well documented (see References [19] and [20]), and the tuning of those parameters always depends on a specific environment and setup.

General DS4000 settings affect not only an SAP system running on DS4000 but also all other applications; single disk or cache settings depend on the specific I/O behavior of the applications.

The same is true for remote volume mirroring (RVM), which will not be discussed within this White Paper as the impact of RVM, although its importance certainly depends on size of the database and the amount of data that will be transferred to a remote site.

Again, SAP application tuning as well as operating system tuning is far beyond the scope of this White Paper. For SAP system tuning, usually external consultants are required who are able to analyze ABAP reports or expensive SQL statements. For operating system tuning, see Reference [18].

The next sub sections will summarize the key parameters and their recommended settings for an SAP database.

3.2.1 DS4000 write cache mirroring

Enabling the DS4000 write cache mirroring is highly recommended, as it provides cached data integrity in the event of a controller failure. For SAP installations in production environments, the data integrity is always much more important than performance. Disabling the write cache should be considered only for nonproduction systems.

Since write cache mirroring uses drive-side FC loops, it competes with normal I/O processing on the drive-side loops; therefore, the degradation in performance is extreme for sequential writes.

However, the performance degradation that accompanies the enabling of the write cache mirror decreases as the percentage of writes decreases.

Recommendation: When using write caching, always use DS4000 write cache mirroring for SAP production environments.

3.2.2 DS4000 controller cache size and cache block size

The size of the adapter cache is often driven by marketing and the need to obtain high benchmark scores in publications. Yet, the best performance does not necessarily come with bigger controller cache sizes, since searching the cache as well as a low cache hit rate degrades performance.

The size of the DS4000 controller cache mainly influences the performance of sequential I/O and should therefore be designed to meet those needs. According to Reference [19], the cache size should be just big enough to do read-ahead; for that reason, no general advice is available, since today the cache size of the DS4000 model is fixed with a size of 2GB.

Within the controller's cache, the minimum cache memory allocation unit is a cache block, which can be of two sizes: 4KB and 16KB.

For random workloads with transfers smaller than 8KB, use the 4KB allocation unit; for sequential workloads, it's better to use the 16KB unit.

Recommendation: Use a DS4000 cache block size of 4KB for Oracle/AIX environments that still use the default block size of 8KB or less.

3.2.3 DS4000 RAID level

Typical SAP R/3 enterprise core installations have two different I/O profiles. One is for the SAP data files and is characterized as 100 percent random I/O, with usually around 70 percent reads and 30 percent writes.

In contrast to the data files, the log files and archived log files are accessed sequentially and the write percentage is periodically very high, which is also true for the read percentage—whenever an online log is archived, for example.

According to SAP's recommendations, and aside from the fact that logs and data files should reside on separate physical disks, you should consider choosing different RAID levels for data and log files.

For these tests, RAID 5 arrays will be created for SAP data files, whereas for log files different scenarios explore the influence of RAID 5 versus RAID 1 (for details, please see the layout section [3.3]).

As documented in Reference [21], using RAID 10 might be considered for OLTP applications such as SAP. Random writes on RAID 10 are significantly faster than on RAID 5, but this is recommended only for the highest performance demands, in which case it is worthwhile to consider using ESS as the storage subsystem.

Recommendation: Use at least RAID 5 for SAP data files; consider using RAID 10 if the write percentage is 40 percent or greater.

3.2.4 DS4000 drive loop utilization

To increase bandwidth, move the expansion enclosures to a new or unused mini hub pair (this doubles the drive-side bandwidth). It is possible to perform this reconfiguration with minimal or no disruption to data availability or interruption of I/O. Assuming that the initial configuration is the one depicted on the left in Figure 7, it is reasonable to move the EXP or drive module DM 2 to the unused mini hub pair on the DS4000900, resulting in the cabling configuration shown on the right.

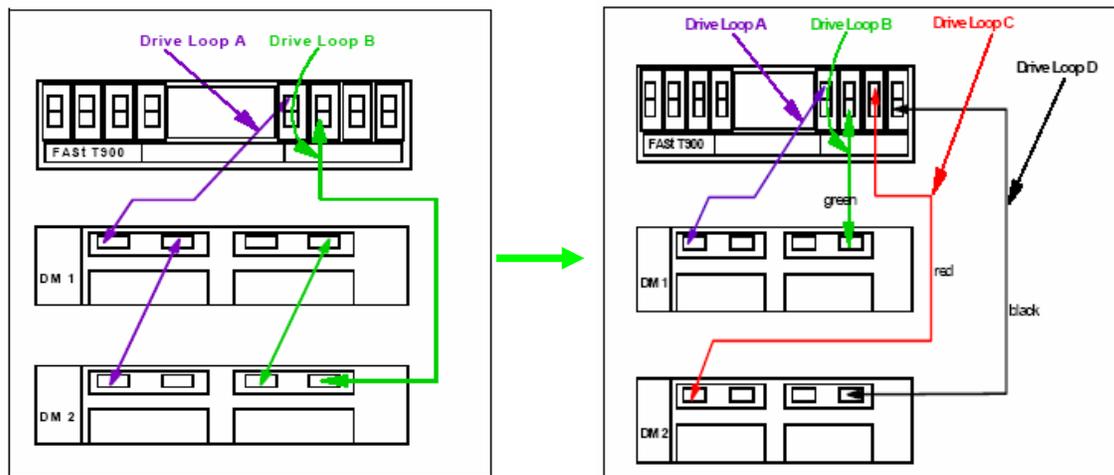


Figure 7: Doubling bandwidth by using all available drive mini hubs (picture taken from Reference [20])

Recommendation: Always use all available drive mini hubs for attaching the expansion units to the DS4000.

3.2.5 DS4000 array creation

This section will summarize some common rules on the creation of DS4000 arrays. Together with the following section, which discusses possible RAID levels, this will help create the arrays needed for the SAP implementation. Since the array creation is the subject of investigation in section 3.3, some basic remarks will be taken into account here.

Channel protection planning

Channel protection is a good way to help make the DS4000 system more resilient against hardware failures. Channel protection means that the arrays are spread across multiple enclosures rather than in just one, to help prevent the failure of a single enclosure from taking a whole array offline. A further potential benefit is a potential performance increase, because multiple ESM (Enclosure Services Monitor) boards along multiple paths (loop side) process the I/O requests.

Number of drives per RAID array

Usually, the more physical drives that belong to an array, the shorter the access time for read and write I/O operations. The only trade-off that has to be taken into account is that the RAID rebuild time increases with the number of disks in an array, which affects the overall system performance. However, as stated in Reference [21], the relation between rebuild time and the number of drives per array is nonlinear, meaning that even when doubling the number of drives per array, the rebuild is less than twice that of the smaller array.

Doubling the number of the physical drives can help increase throughput performance by as much as 50 percent. As a golden mean between capacity, performance and rebuild time typically between 5 and 16 drives per RAID 5 array are recommended. For a specific application, it is sometimes reasonable from a performance perspective to use more drives for an array than needed capacity and place another low-activity application on the same array.

Use both loops to the drives

As each DS4000 enclosure is attached to the controllers via two loops, it is recommended that you create each RAID array in such a manner that both loops are used simultaneously. The easiest way to achieve that is to let the DS4000 automatically create the RAID array. For manual creation, a rule of thumb is that two contiguous drives be accessed via different loops.

Recommendation: Spread the RAID arrays across all enclosures and all loops. For RAID 5, use 5 to 16 drives per array.

Because the subject of investigation in this project was the creation of the RAID arrays and consequently the DS4000 layout, it is advised to refer also to section 3.3 for further recommendations.

3.2.6 DS4000 segment size

The choice of a segment size for a logical volume can have a major influence on performance in both IO/s and throughput. Large segment sizes can help increase the request rate (IO/s) by allowing multiple disk drives to respond to multiple requests. Small segment sizes can help increase the data transfer rate (Mbps) by allowing multiple disk drives to participate in one I/O request. It is suggested that you use a small segment size relative to the I/O size to help increase sequential performance. If the typical I/O size is larger than the segment size, it's better to increase the segment size in order to minimize the number of drives needed to satisfy an I/O request. This is especially true in a multiuser, database (SAP) or filesystem storage environment. Using a single drive for a single request leaves other drives available to service other requests simultaneously.

Possible segment sizes are 8KB, 16KB, 32KB, 64KB, 128KB and 256KB.

The testing of segment size was not part of the investigation for this White Paper; thus throughout all tests a small segment size of 16KB was chosen for all logical volumes following the recommendations that can be verified in Reference [20].

However, for real random workloads with block sizes of 8KB, recent tests pointed to a better performance when using the DS4000 default segment size of 64KB. Therefore, the recommendation differs from the segment size that has been used for this project.

Recommendation: Start using a segment size of 8KB to 16KB for database applications such as SAP. For real random 8KB workloads, consider using a segment size of 64KB.

3.2.7 DS4000 read-ahead parameter

This parameter affects the reading performance; an incorrect setting can have a negative impact. The read-ahead parameter controls the number of additional sequential data blocks that will be cached after a read request. Obviously, if the workload is random, this value should be 0. Otherwise, each read request will unnecessarily prefetch additional data blocks. Because these data blocks are rarely needed, this has a negative impact on performance.

For merely sequential workloads, a good value is between 1 and 8, depending on the particular environment. When using such a setting, a read request causes the prefetching of several sequential data blocks into the cache; this helps speed up subsequent disk access. This leads to a smaller number of I/O transfers (between disk and cache) required to handle the same amount of data, which in a sequential environment is good for performance. A value

that is too high can cause an overall decrease in performance, because the cache is filled with read-ahead data that is never used.

A good approach is to adjust the read-ahead parameter by monitoring the cache hit ration with the DS4000 performance monitor.

Recommendation: For SAP data files, use a read-ahead multiplier of 0; for SAP log files, use 4.

3.2.8 LUN mapping, ownership and multipathing

It is understood that using several paths to the DS4000 can help improve data availability and potentially increase performance.

LUN mapping, also known as storage partitioning, adds a high level of flexibility to the DS4000 Storage Server. It is designed to enable the connection of multiple and heterogeneous host systems, either in stand-alone or clustered mode to the same storage server.

A LUN (Logical Unit Number, which represents the number a host uses to access the logical drive) is the same as a logical drive or logical volume and is defined over a set of drives called an array with a specific RAID level.

However, when assigning logical volumes to the systems, it's very important to remember that the DS4000 Storage Server uses a preferred controller ownership approach for communicating with LUNs. This means that every LUN is owned by only one controller. Therefore at the system level it is important to make sure that traffic is balanced among controllers correctly. This is a fundamental principle for achieving a correct setting of the storage system. Balancing traffic is unfortunately not always a trivial task, which is the motivation of the next chapter's investigation of different array and ownership layouts and their influence on performance.

3.3 DS4000 array layouts and results

This section explains the single layouts that have been tested. Accompanying each layout will be a description of the implementation, the settings and the motivation for picking the layout. Each subsection concludes with a discussion of the performance that has been achieved by choosing the specific layout as well as the work efforts that come with maintaining the structure and the performance in production environments.

So that one disk layout may be easily compared with another, only minor changes occurred from each layout to the next one.

For the first layouts (numbers 01 to 04), a fixed amount of disk drives and expansion units has been used.

Starting with layout 05, more than one expansion unit and, finally, more than 14 disk drives have been used to demonstrate the effects of spreading arrays as much as possible across the DS4000.

For each layout, if the settings are not mentioned in the description, they can be assumed to be as described in the previous section.

During the tests, it was observed that the total job execution speed increased even though the I/O and throughput values remained the same. This can be explained by noting that the four client copies were executed on the operating system level with a varying amount of OS processes. Therefore, an interpretation of the results will also take into account the number of processes.

A summary of all performance and layout recommendations as well as some graphical views of the measured data will be provided at the end of section 3.4.

3.3.1 The basic layout 01

The first layout assumes little tuning of the storage subsystem. However, as a minimal configuration, there was a logical separation of SAP data files from the log files.

The array was built with one EXP and 14 drives (see chapter 2 for the hardware availability).

The picture below illustrates disk layout 01.

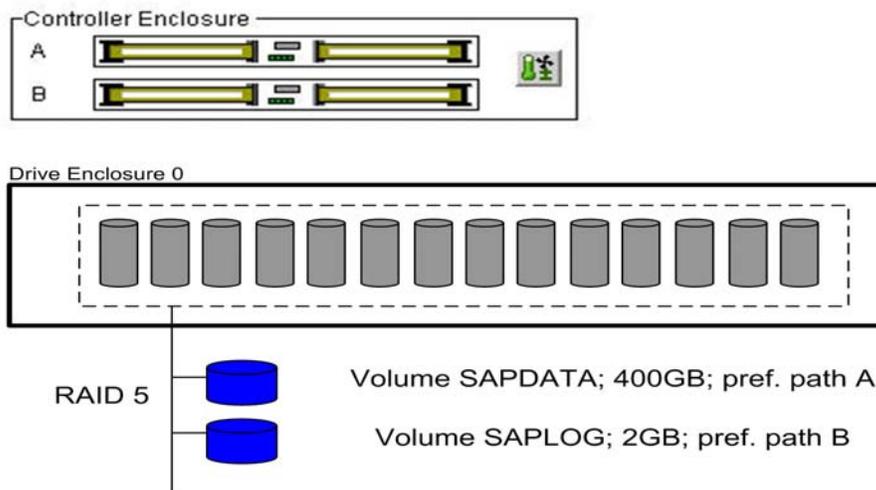


Figure 8: Disk layout for test scenario 01

Figure 8 illustrates one RAID 5 array that was built out of all disks of expansion unit 0 with two logical drives, one for the SAP data files (SAPDATA) and one for the log files (SAPLOG).

For this layout, as well as for all subsequent layouts, the DS4000 parameters were set as follows:

- segment size = 16KB
- cache read-ahead multiplier for SAPLOG = 4
- cache read-ahead multiplier for SAPDATA = 0
- on AIX, an inter-disk allocation policy set to maximum (PP striping), if applicable

All other DS4000 settings have been applied as recommended in the previous sections.

The inter-disk allocation policy applied only to layouts where several SAPDATA or SAPLOG volumes were created.

All subsequent layouts will be illustrated as displayed in Figure 9 for readability reasons. Because the RAID creation was performed as explained in section 3.2.5 by either letting the DS4000 choose the disks for each array or by performing it manually in the same way, the figures omit the disk distribution for each array.

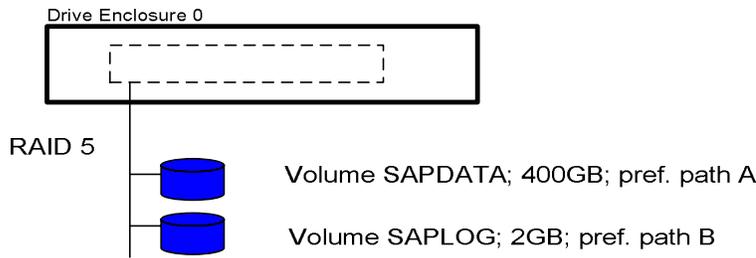


Figure 9: Simplified disk layout for scenario 01

Table 1 displays the values that resulted from this layout. The KB/s and the IO/s have been extracted from the DS4000 performance monitoring.

Total job runtime [s]	79373
Number of processes at OS level	17
Average system idle time [%]	27,05628
Average system iowait time [%]	33,95537
Maximum disk activity tm_act [%]	100 for SAPDATA disk
Maximum KB/s	61772,5
Average KB/s	13987,62
Maximum IO/s	5807

Average IO/s	1115,138
--------------	----------

Table 1: Performance values for disk layout 01

As already stated, the displayed values do *not* represent absolute DS4000 performance data but rather the relative test data, which is of interest only when compared with the results of the other tests.

It is important to mention that the SAP server shows high values of iowait time together with the disk activity percentage. It becomes clear that at the least, the SAPDATA disk is a potential bottleneck. Although Table 1 displays only the maximum *tm_act* value, the complete iostat series of 3,500 iterations was analyzed.

It is obvious that running SAP data I/O and log I/O profiles against the same physical disks comes with a high risk of I/O contention.

From the point of view of cost and administrative effort, this layout is may be cost effective, as one RAID 5 array needs only one disk of parity, which cannot be used for storing data. Administering this layout in case of data growth can easily be done just by adding additional disks to the array.

Conclusion: From a performance perspective, layout 01, although it is economical and easy, is not recommended.

The next layout tries to improve layout 01 by optimizing the load distribution from a DS4000 controller point of view.

3.3.2 The load-balanced basic layout 02

Instead of creating only two logical volumes, this layout differs from layout 01 by now creating four logical volumes and distributing them to each of the DS4000's controllers.

No other changes have been performed, but the database has been restored to its initial state.

Figure 10 shows layout 02:

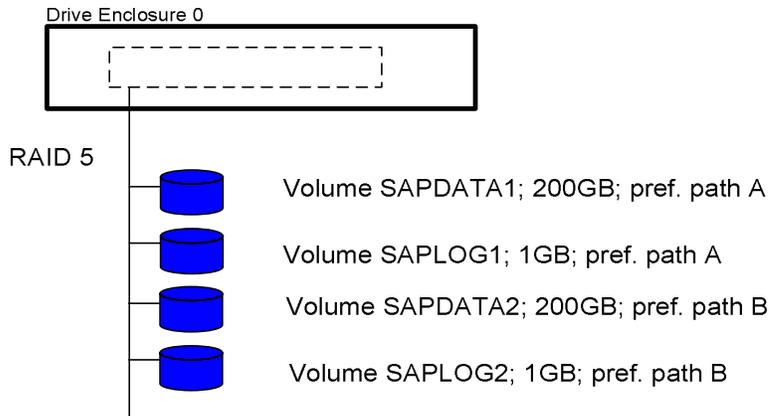


Figure 10: Simplified disk layout for scenario 02

The characteristics of this layout are summarized in Table 2:

Total job runtime [s]	84021
Number of processes at OS level	17
Average system idle time [%]	25,41164
Average system iowait time [%]	36,99629
Maximum disk activity tm_act [%]	97,8 (SAPDATA)
Maximum KB/s	63163
Average KB/s	13942,18
Maximum IO/s	4420,2
Average IO/s	1071,884

Table 2: Performance values for disk layout 02

Although the averages of IO/s and KB/s were slightly lower than those of disk layout 01, the total job execution time was longer. A review of the SAP client copy error log yields the fact that during the tests with disk layout 02, one client copy job terminated with errors. Whenever an error occurs during the client copy, SAP starts to retry the failed operation, which obviously leads to a longer execution time. Client copies that terminated with errors have not been repeated since they were used only for producing I/O to the subsystems.

There can be no apparent broad conclusion out of the two runs beyond the conclusion that the load was balanced more evenly. This may be reflected in a slightly better balanced disk activity (at least below 100 percent). Assuming the SAPDATA disks (and consequently the

total amount of disks that are available for SAP's data files) are bottlenecks, no real I/O or throughput improvement can be expected.

Administration of this layout is slightly more complex because increasing the disk space for the SAP data files requires the simultaneous increase of two logical volumes. This is also true with the AIX level.

Conclusion: Compared to Layout 01, Layout 02 distributed logical volumes for load-balancing reasons at the same cost but with slightly greater administration effort.

The next layout still tries to use the manual load distribution, but it also supports the assumption that the exclusive use of one array by only one controller may help improve performance.

3.3.3 Advanced load-balanced layout 03

Performance interpretations examined in other tests lead to the theory that the exclusive use of one RAID array by one DS4000 controller may help improve the performance due to DS4000 internal I/O management.

Therefore, instead of one array, disk layout 03 consists of two RAID 5 arrays, each with two logical disks that are accessed by only one DS4000 controller. Load distribution occurs by creating the same logical volumes on each array and joining them on the AIX level via PP striping.

Figure 11 illustrates scenario number 03.

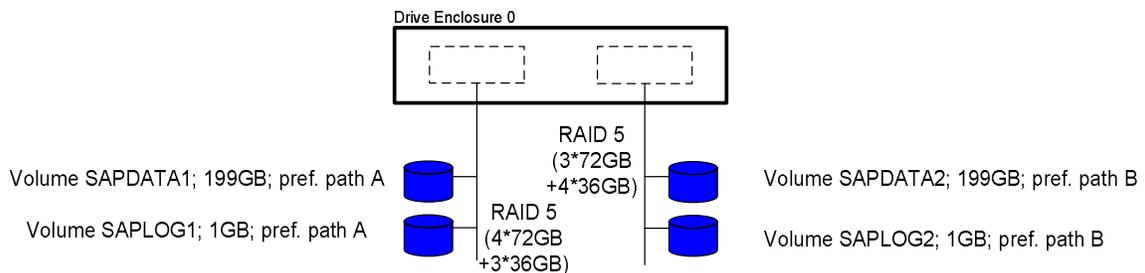


Figure 11: Simplified disk layout for scenario 03

Interpreting the values displayed in the table below allows two conclusions. First, the IO/s and throughput values are very similar to (and even slightly better than) those for disk layout 02, although this configuration actually costs one more data disk because two RAID 5 parity

disks are needed rather than one. Secondly, the total job runtime is very similar to that of test run 02.

Again, the same job terminated with the same errors as in the previous tests, making both test runs comparable.

Total job runtime [s]	84793
Number of processes at OS level	18
Average system idle time [%]	27,67259
Average system iowait time [%]	30,62279
Maximum disk activity tm_act [%]	99,6 (SAPDATA)
Maximum KB/s	68816,8
Average KB/s	14110,48
Maximum IO/s	4520
Average IO/s	1095,393

Table 3: Performance values for disk layout 03

No significant change in runtime, IO/s or throughput was experienced in this test run.

However, a look at the complete IOSTAT data – not just the maximum value – the disk activity percentage is generally lower than those of the previous layouts, strengthening the assumption that the exclusive usage of one array per controller helps better balance the load. This is also reflected in lower system iowait averages.

For layouts 01 to 03, a general conclusion is that although the SAPDATA volumes tend to be a bottleneck for all three layouts, load balancing can help improve by balancing the volumes and the RAID arrays across both DS4000 controllers. This potential advantage's trade-off is higher cost as one additional disk's capacity is lost for parity, as well as greater administrative effort in case of database growth – because here, two RAID arrays have to be expanded in addition to the simultaneous expansion of two data volumes.

Conclusion: Compared to previous layouts, load balancing had potential benefits, but costs may be higher and more administrative effort may be required. Further investigation is needed and will be performed in the next layouts.

The next layouts will take into account SAP's recommendation of separating the log from the data volumes.

3.3.4 Log separation layout 04

For this layout, two RAID arrays have again been created, one for the SAP log files and one for the data files. The findings from layout 02 (the basic load balancing of the logical DS4000 volumes across controllers) have been applied, but the arrays are no longer exclusively used by one controller.

Figure 12 demonstrates layout 04.

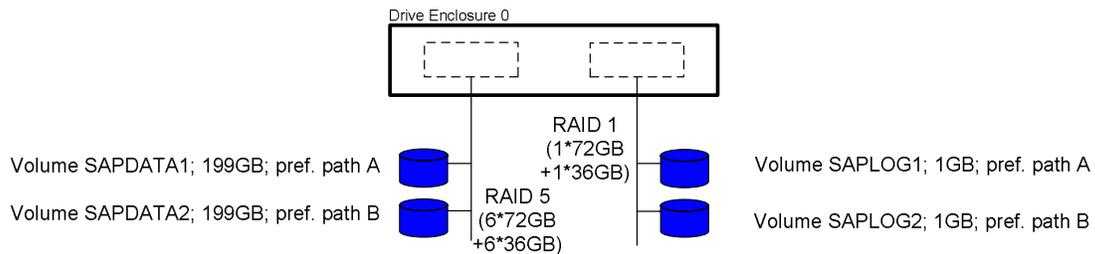


Figure 12: Simplified disk layout for scenario 04

After joining the logical DS4000 volumes for SAPDATA and SAPLOG on the AIX side by way of PP striping, the values displayed in Table 4 were achieved.

Total job runtime [s]	85828
Number of processes at OS level	17
Average system idle time [%]	31,27932
Average system iowait time [%]	30,62654
Maximum disk activity tm_act [%]	99,6 (SAPDATA)
Maximum KB/s	69565,6
Average KB/s	13657
Maximum IO/s	4746,8
Average IO/s	1074,75

Table 4: Performance values for disk layout 04

For this layout, the total job runtime (again, one job terminated with errors) increased and the throughput and I/O values decreased.

It seems that the separation of SAP data files and log files to different physical disks has no advantages. This is deceptive, though, because as long as the amount of disks for the SAP

data files obviously represents a bottleneck (high tm_act values!), the separation of the log files will not improve the performance.

Conclusion: The separation of the logs had no influence on overall performance as long as there were disk bottlenecks; more investigation is needed and will take place in the next scenarios.

For a more detailed investigation of layout 03 and 04, see the next layout, which composes array load balancing and log separation.

3.3.5 Advanced load-balanced layout 04_2 with log separation

As depicted in the figure below, for each of the four RAID arrays only one logical volume was created, guaranteeing the exclusive access of one controller to each array.

On the AIX LVM level, the logical DS4000 volumes have been joined with PP striping means.

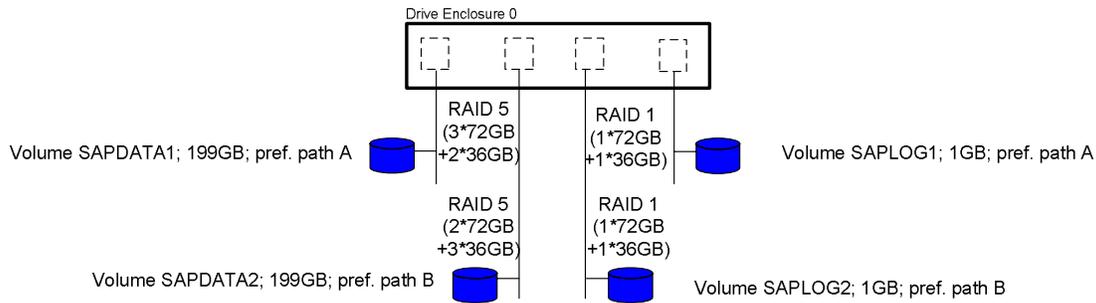


Figure 13: Simplified disk layout for scenario 04_02

Please note that in this layout the usable capacity of one EXP is reduced considerably due to the usage of two RAID 5 and two RAID 1 arrays. Compared with layout 01, the capacity reduction is up to ~23 percent.

The surprising results for this scenario are below:

Total job runtime [s]	66823
Number of processes at OS level	19
Average system idle time [%]	31,70343
Average system iowait time [%]	29,36591

Maximum disk activity tm_act [%]	98,4 (SAPDATA)
Maximum KB/s	77527,8
Average KB/s	15813,54
Maximum IO/s	4890,2
Average IO/s	1233,599

Table 5: Performance values for disk layout 04_02

Although fewer disks were actually used for the SAP data volumes (which still were the performance bottleneck), the throughput and IO/s values increased significantly while the total runtime decreased.

This means that even for a system that is running at its performance limit, the load balancing of arrays to an exclusive controller together with the log separation can help drastically improve overall performance.

Comparing the values with those of layout 04, the throughput increased by about 16 percent, IO/s by about 15 percent and runtime by almost a 30 percent boost in speed.

For the potentially increased performance on a production system, these results justify the reduced capacity and additional administrative efforts that come with the maintenance of this layout.

Conclusion: In layout 04_2, the exclusive usage of each array per controller together with separating the log disks from the data disks is recommended.

Since this layout represents the optimum so far, the next one tries to obtain even better results by intelligently joining Oracle logs according to SAP's recommendations.

3.3.6 Load-balanced layout 04_3 with optimized log separation

In all the layouts discussed so far, tuning has been performed mainly on the DS4000 level. For SAP installations that run on an Oracle database, some SAP recommendations advise how to separate the logs for minimal configurations (that is, no individual disks for each log filesystem).

As documented in Reference [16], Oracle writes its log information into four filesystems:

- origlogA
- mirrlogA

- origlogB
- mirrlogB

Only two filesystems get written to at a time, since Oracle mirrors logs A and B.

In order to optimize disk access, SAP recommends putting the filesystems origlogA and mirrlogB together on one RAID 1 array, and origlogB with mirrlogA on another. This recommendation comes from the fact that during write access to log A (that is, origlogA and mirrlogA are written to), log B gets archived by reading from origlogB, resulting in I/O for both RAID arrays. The same is true when log A is full and a switch to log B takes place.

Therefore, this DS4000 layout looks exactly like the previous layout, except that on the AIX level the two logical log volumes have not been joined via PP striping but rather separated with the position of the filesystems as described above.

Table 6 once again displays the obtained results.

Total job runtime [s]	64469
Number of processes at OS level	17
Average system idle time [%]	32,74
Average system iowait time [%]	22,00
Maximum disk activity tm_act [%]	97,8 (SAPDATA)
Maximum KB/s	66247,4
Average KB/s	15544,22543
Maximum IO/s	4549,6
Average IO/s	1232,090543

Table 6: Performance values for disk layout 04_03

Although in this setup nearly the same throughput and I/O per second values were achieved, the average system iowait percentage was reduced by more than 7 percent. This points to a more balanced utilization of the disks, as expected. This is also reflected in a shorter cumulative job execution runtime, which was decreased by 3.6 percent.

Considering cost and administration efforts, there is no difference in the previous scenario. Following SAP's recommendation, it is advised to distribute the logs as demonstrated here.

Conclusion: In layout 04_3, following SAP's recommendation on Oracle's log position for minimal disk configurations is recommended.

The latest scenario, layout 04_3, reached an unexpected level of performance optimization, the major finding being that the exclusive usage of one RAID array per controller may have a marked influence on the performance.

As described in section 3.2, the use of more disks and more expansion units can help improve the overall performance of the SAP system.

Finally, the next layout examines the influence of more disks and EXPs compared with the restricted usage of one array per controller.

3.3.7 The full-blown layout 05

The layout illustrated in Figure 14 involves the usage of all available disks as well as the distribution over all available expansion units. For each array (the small number in the array box depicts the RAID level), the disks have been picked manually as suggested in section 2.2.5, always consisting of 36GB and 72GB disks. Though it wastes capacity, this was done for reasons of consistency with respect to all previous layouts. RAID 1 arrays use two disks, while each RAID 5 array consists of five.

Some arrays are used solely by one controller. (This is true for all log volumes.) For the data and Oracle archive volumes, a manual load-balanced distribution was assumed in which both logical volumes were joined at OS level by way of PP striping.

Using the findings from layout 04_2, this layout could have been tuned even more; for example, doubling the amount of exclusively owned RAID levels used for the data and archive files, or alternatively by distributing the two logical volumes to separate RAID 5 arrays. However, this has not been done, for three reasons:

1. The performance influence of this distributed layout should be tested compared with the performance of one of the exclusive array usage layouts.
2. The creation and maintenance of such a volume arrangement is hardly reasonable.
3. Using a whole DS4000 for a single SAP system is rather uncommon.

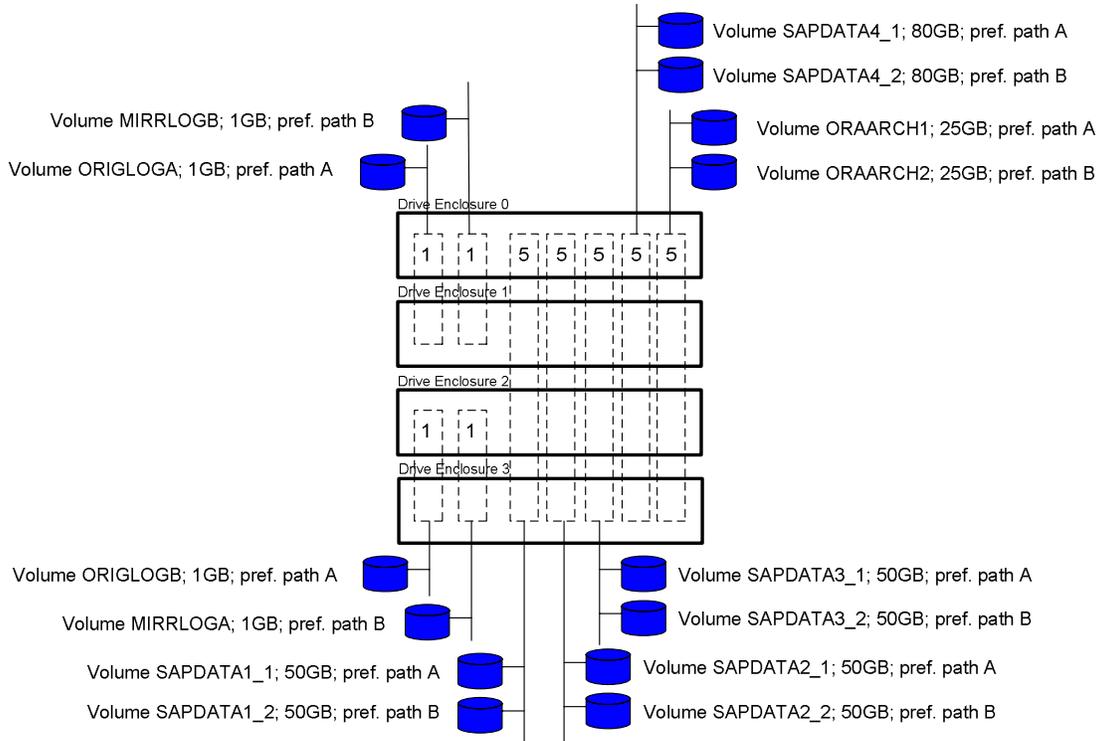


Figure 14: Simplified disk layout for scenario 05

It was expected that this layout would improve the performance significantly; yet surprisingly, the values in Table 7 below show an overall performance reduction compared with the tests 4_2 and 4_3.

Total job runtime [s]	70154
Number of processes at OS level	17
Average system idle time [%]	30,72996
Average system iowait time [%]	34,1
Maximum disk activity tm_act [%]	97,6 (SAPDATA4)
Maximum KB/s	69724
Average KB/s	14495,98
Maximum IO/s	5264,8
Average IO/s	1163,988

Table 7: Performance values for disk layout 05

Cumulative job runtime, as well as the throughput and IO/s values, was worse than those of both preceding layouts, even though the average disk activity was greatly reduced. (The table displays only the maximum, which is a bit misleading.) Consequently, it is assumed that the controllers, suffering from I/O contention, limited the throughput and IO/s behavior.

However, to prove this assumption, more exhaustive testing would be necessary, which is beyond the scope of this White Paper.

Comparing these achievements with those of the earlier tests, it is important to summarize that controller ownership for an array apparently has a much higher than expected influence on performance; it may even exceed the benefits of using more disks and DS4000 expansion units. Nonetheless, using more disks and EXPs led to more I/O.

Yet these findings are true for our workload profile only. Whenever the amount of disks seems to be the bottleneck, it is highly recommended that you use more. Using more expansion units is also still recommended for availability reasons, as well as in those environments where performance is limited by the bandwidth of the mini-hubs, for example.

This scenario concludes the DS4000 layout section. The next section summarizes all recommendations and findings for readability and lookup reasons and illustrates some values that have been displayed only in the tables so far.

3.4 DS4000 layout findings and recommendations

This section puts together all the recommendations and includes some graphical views that allow an immediate performance comparison between different DS4000 layouts.

According to the earlier structure, Table 8 will give general recommendations first:

Recommendation	Distribute sapdata and redo log files to separate disks for performance and FlashCopy backup reasons. Do not FlashCopy SAP data and logs at the same time!
Recommendation	When using AIX LVM, use an inter-disk allocation policy of maximum in order to spread the physical partitions (PP) of the logical volume across as many physical disks, controllers and communication paths as possible (this is also referred to as PP striping).
Recommendation	When using write caching, always use DS4000 write cache mirroring for SAP productive environments.
Recommendation	Use a DS4000 cache block size of 4KB for Oracle/AIX environments that still use the default block size of 8KB or less; adjust according to specific requirements.
Recommendation	Use at least RAID 5 for SAP data files, consider using RAID 10 if write percentage is greater than or equal to 40 percent; adjust according to specific requirements.
Recommendation	Use all available drive mini hubs for attaching the expansion units to the DS4000.
Recommendation	Spread the RAID arrays across all enclosures and all loops. For RAID 5, use an amount of 5-16 drives per array.
Recommendation	Start using a segment size of 8KB to 16KB for database applications such as SAP. For real random 8KB workloads, consider using a segment size of 64KB; adjust according to specific requirements.
Recommendation	For SAP data files, use a read-ahead multiplier of 0, and for SAP log files, use 4; adjust according to specific workload profile.

Table 8: Summary of recommendations for DS4000 settings for SAP with Oracle on AIX

The following figures allow the comparison of single layouts with each other and help to explain the conclusions that Table 9 summarizes.

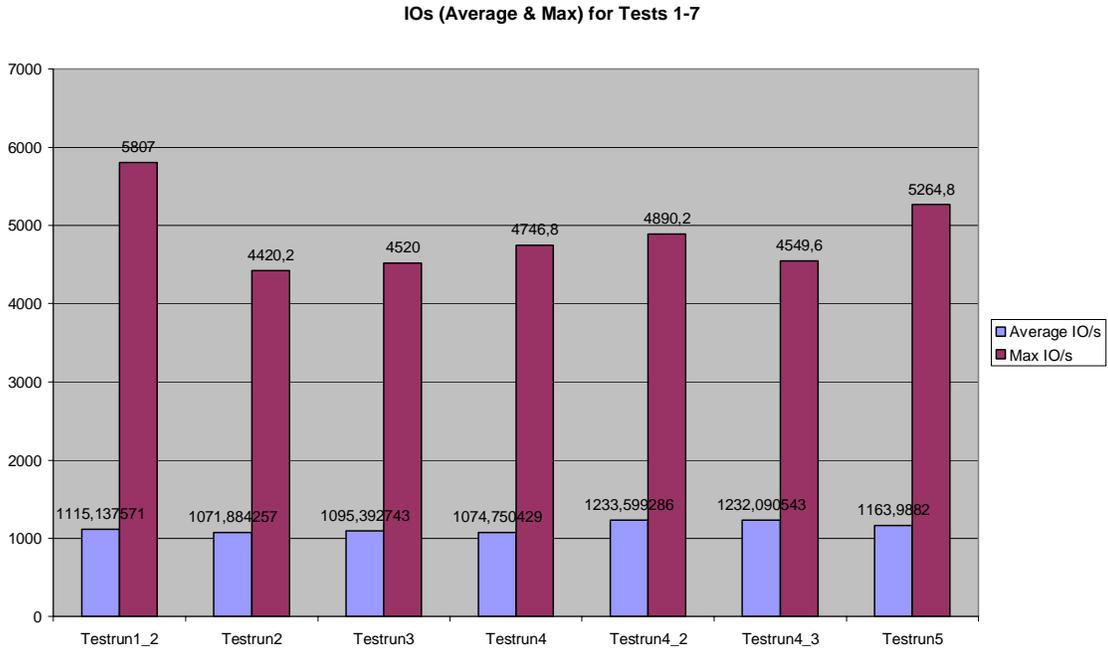


Figure 15: IO/s averages and maximum from the DS4000 performance monitor for all test layouts

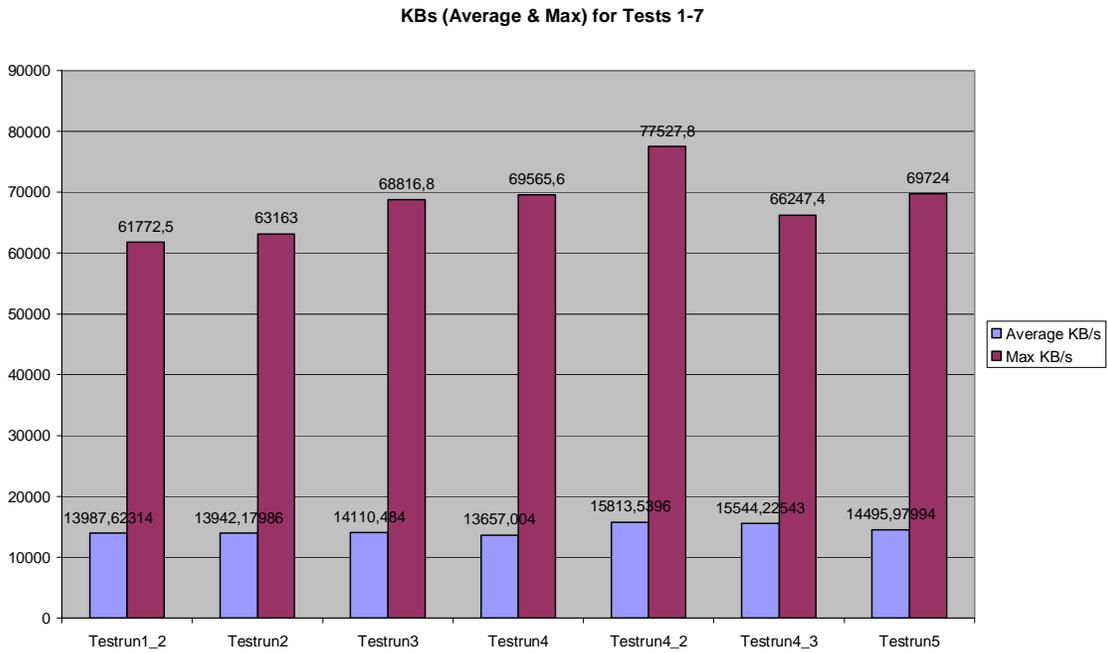


Figure 16: KB/s averages and maximum from the DS4000 performance monitor for all test layouts

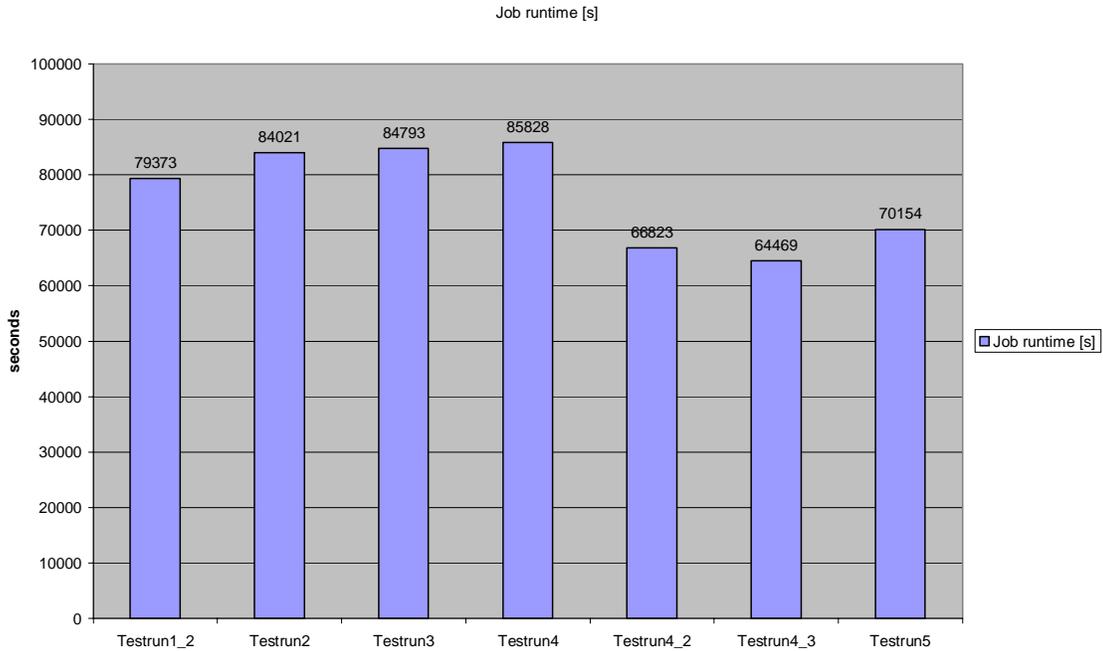


Figure 17: Cumulative client copy job runtime from SAP transaction SM37 for all test layouts

<i>Conclusion</i>	<i>Layout 01 – Although economical and easy, not recommended from a performance perspective.</i>
<i>Conclusion</i>	<i>Layout 02 – Distribute logical volumes for load-balancing reasons at the same cost but with slightly higher administration efforts.</i>
<i>Conclusion</i>	<i>Layout 04 – Exclusive usage of each array per controller together with separating the log disks from the data disks is recommended.</i>
<i>Conclusion</i>	<i>Layout 04_3 – Following SAP’s recommendation on Oracle’s log position for minimal disk configurations is recommended.</i>

Table 9: Summary of conclusions for DS4000 RAID array layout for SAP with Oracle on AIX

As stated earlier, the recommendations and conclusions in this section are not considered to be the optimized performance tuning for all SAP systems running on DS4000. They may serve as a starting point to help deliver some good ratio between performance, capacity and administration costs and, most of all, to demonstrate the variety of tunable DS4000 parameters that affect performance and are worth considering in the planning phase. Because our tests were performed with four simultaneous client copies, they do not completely represent a typical SAP production system. Thus, our findings provide ideas and some orientation rather than represent directives.

For other operating systems or databases, following these suggestions may provide a good approach for performance, since the layout information is independent of the OS and the settings recommendations may be easily adjusted.

4 Automated SAP backup solution with DS4000 FlashCopy

The introduction of this White Paper explains today's business requirements for fast and intelligent backup strategies.

This chapter describes the theoretical background of the SAP backup solution that was developed for our tests:

- the backup tools and interfaces SAP provides
- IBM/Tivoli's interface between BR*Tools and Tivoli Storage Manager
- the integration of DS4000 FlashCopy
- additional tools needed to get the SAP backup solution running

SAP provides concepts and interface definitions to its Oracle database backup tools (BR*Tools), which can be used to develop backup solutions that offer this quality.

1. Split Mirror Disk Backup for Oracle, Reference [15]
2. SPLITINT Interface for Oracle Databases, Reference [14]

For a Split Mirror Disk Backup, the Oracle database must be located on a mirrored disk environment. Using any kind of RAID mirroring is helpful in terms of data availability, but it is not sufficient for backup as a logical database mirror is required. SAP's backup tool (*brbackup*) calls two scripts or programs to split the mirror before the backup and resynchronizes it after the backup has been completed.

If a SPLITINT interface is implemented, *brbackup* calls a script called *splitint* to perform the split or resynchronization.

At first view, both concepts look very similar. We chose a SPLITINT implementation, because in our opinion the SPLITINT interface is designed to offer three benefits:

- a formal definition of the call interface with a well-defined syntax
- possibilities to pass control flow from SAP's backup tools to the *splitint* script and back
- the possibility of SAP certification, as all SAP rules and APIs are obeyed

Integrating a SPLITINT interface with the advanced copy functions of DS4000 allows two major advantages: The window for backing up SAP installations of any size is minimized, and SAP system performance is optimized.

The fact that the DS4000 FlashCopy function is used instead of mirroring, mirror splitting and resynchronization makes this solution a valuable one. No data needs to be copied on the server side. The DS4000 is designed to copy only those data blocks that are changed by an application on the source or target volumes. Mirror resynchronization, which is a task that can be highly time-consuming for the processor, is replaced by some simple actions: The FlashCopy target volumes are deleted or disabled on the DS4000.

Archival of the database logs and database restore will remain on the SAP server. In the event of a database crash, it may be far too much trouble to move the data twice. The following figure illustrates the data flow for a database backup and restore.

The SAP server runs the SAP production environment while the split mirror disks will be mounted on the so-called backup server. The backup data flows from the backup server by way of the TSM server to tape.

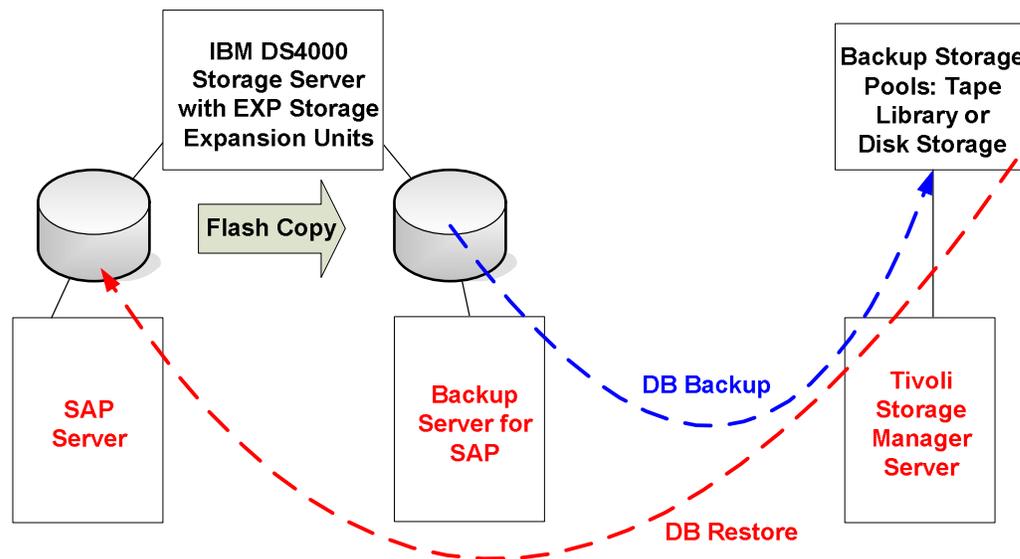


Figure 18: Data flow for database backup and restore

This chapter offers such a scripting solution, which has been developed to automate the backup process. The integration of the solution into SAP's backup tools helps allow backup administration and monitoring from within SAP.

The steps necessary to implement scripts that map to the customer's environment are discussed in a cookbook-like manner. The Appendix includes some example scripts. In addition to helping reduce costs, this approach helps reduce the complexity of implementation and avoid several potential traps, since each step has its own hints and tips section.

Moreover, this implementation is almost independent of the operating system on which the SAP database instance is running, since the scripts are implemented as shell scripts. The

scripts' modularity helps allow adapting this implementation to other IBM storage subsystems, such as Enterprise Storage Server (ESS) or TotalStorage SAN Volume Controller (SVC).

For SAP installations on ESS, IBM/Tivoli offers a product combination designed to clone and backup/restore an SAP environment: TSM for ERP Systems and TDP for Hardware. The two products together offer a highly automated and comfortable solution with support not only from IBM but also from SAP, as these products are SAP certified.

We are focusing on the backup aspect of SAP installations on DS4000 only. Without automated product installation procedures, some manual installation is necessary to implement the solution described here.

4.1 Developing a backup solution for SAP environments

This section describes the theoretical background of an SAP backup solution. It includes a short introduction to SAP's backup tools (BR*Tools). After this, the integration of Tivoli Storage Manager with TSM for ERP Systems is explained. The last step is separating SAP application and backup tasks onto different servers.

The concept and products are not explained in detail. Please see the Appendix, which provides references to product manuals, White Papers and other relevant documentation.

4.1.1 BR*Tools overview

Oracle database backup must take into consideration several file types: control files, data files, online redo logs and offline redo logs (also called an archive).

SAP provides the so-called BR*Tools for Oracle database backup and recovery:

- *brbackup* to back up database files, control files and online redo log files
- *brarchive* to back up the offline redo logs (archive)
- *brrestore* to restore all the objects mentioned above
- *brrecover* to restore and recover the database after severe errors
- *sapdba* is a menu-controlled tool for database administration (startup, shutdown, the starting of other BR*Tools and more)

Figure 19 illustrates the usage of the different BR*Tools for Oracle database backups.

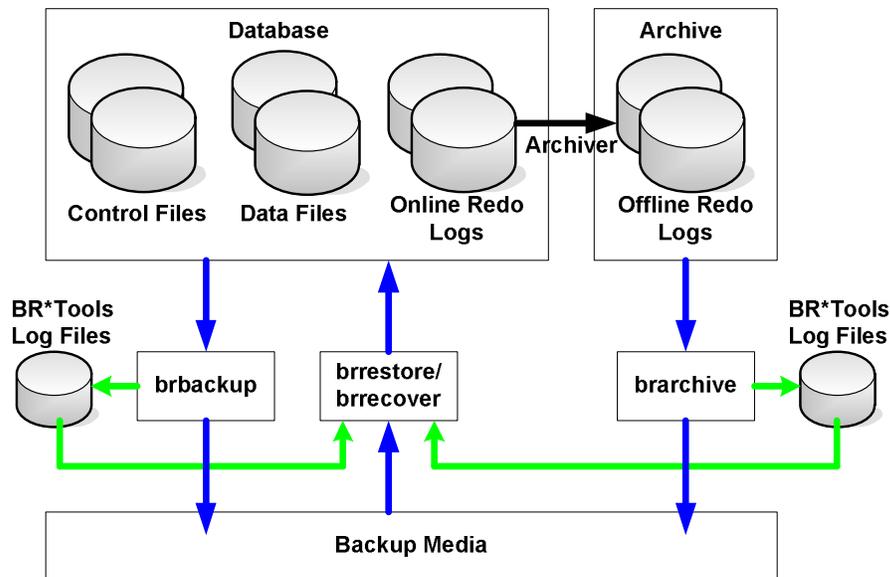


Figure 19: Using BR*Tools to back up Oracle database objects

The backup tools are writing log information to files. This data is necessary for future restore or recovery and can be displayed in SAP's Computing Center Management System (CCMS).

Each of the BR*Tools maintains its own set of log files.

BR*Tools as delivered by SAP enable backup and recovery based on operating system-specific utilities like *cpio* and *dd* or on Oracle's Recovery Manager (RMAN).

Hint:

Currently, *brrecover* offers the full recovery and repair features exclusively. Starting with BR*Tools release 6.20 patch level 100, *sapdba* can no longer be used for database recovery.

Reference: SAP note 602497

Backing up to tape media always involves the use of backup software that controls backups. The interaction of the SAP database backup tools and the specific backup software requires an additional piece of software, which is described next.

4.1.2 Tivoli Storage Manager for ERP

SAP provides an interface description that lets vendors of backup/restore software connect their software products to the SAP BR*Tools. IBM/Tivoli offers TSM for ERP Systems (Reference [22]) to interact with Tivoli Storage Manager. In many documents, the terms Tivoli Data Protection (TDP) for mySAP.com or simply *backint* are used as synonyms. TDP for mySAP.com was the old product name, whereas *backint* is the name of the software executable of TSM for ERP.

The following figure shows the interaction between BR*Tools, TSM for ERP Systems and Tivoli Storage Manager.

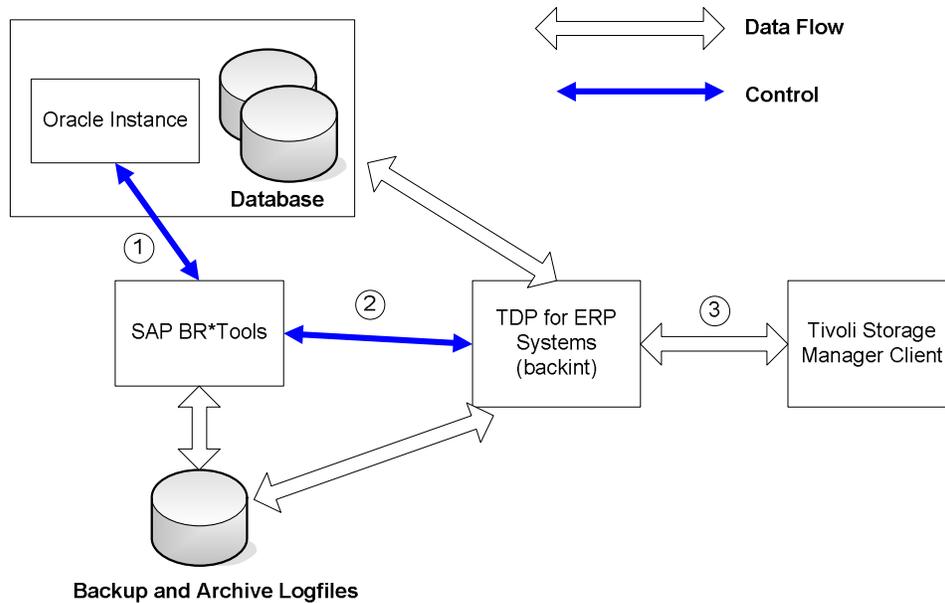


Figure 20: Interaction of BR*Tools with TDP

- *brbackup* causes Oracle to shut down the database (offline backup) or set tablespaces to backup mode (online backup) (1).
- *brbackup* calls *backint* with a list of files to be backed up (2).
- *backint* backs up all the requested file using the TSM client, then returns control to *brbackup* (3).
- In addition, *brarchive* updates and controls the backup of the database redo log files.

4.1.3 DS4000 FlashCopy integration

The DS4000 FlashCopy function creates a point-in-time (t0) copy of those DS4000 logical volumes on which the SAP database files are installed. By attaching these FlashCopy target volumes to another server, it is possible to separate the database backup from the SAP production server, and therefore to help eliminate the performance impact caused by the backup.

Naturally, control flow and data flow are getting more complex in our figures:

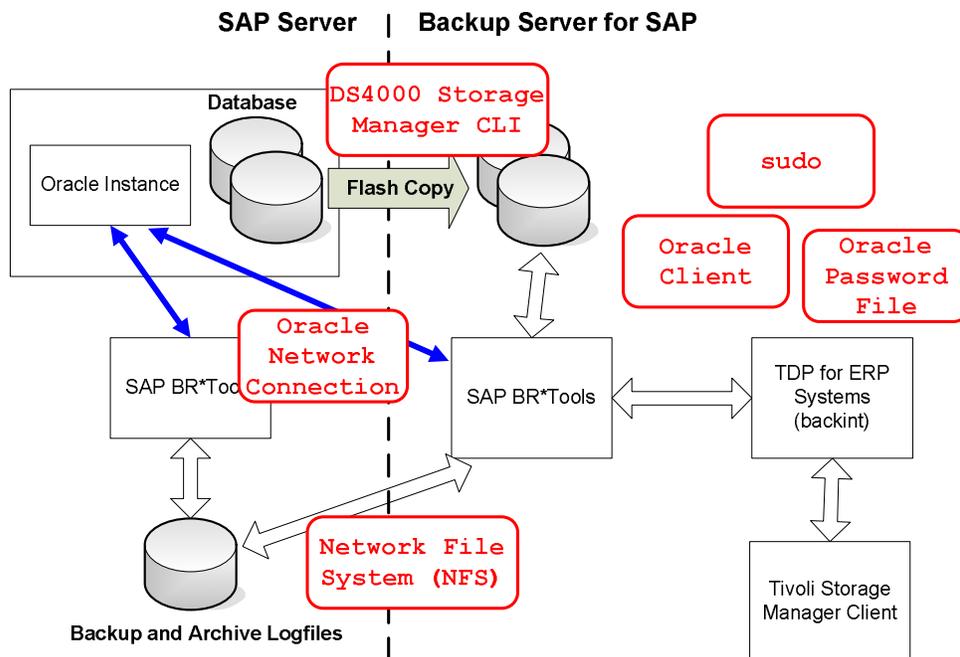


Figure 21: BR*Tools and DS4000 FlashCopy integration

This approach requires additional functionality to establish data and control flow:

- DS4000 Storage Manager CLI to manage DS4000 FlashCopy copies.
- An Oracle client on the backup server and an Oracle Net connection between SAP and the backup server to enable remote communication between BR*Tools (which will be started on the backup server) and Oracle (on the SAP server).
- An Oracle security mechanism that allows remote database administration as access to a database administrator's account is tightly controlled. Further authentication is required to administer a database over a remote connection.
- An operating system feature (here: *sudo*) that enables nonprivileged users to execute system administration commands. *Sudo* will be described in section 4.2.12. BR*Tools are started on the operating system level by the OS user who was created for database administration tasks. Besides starting FlashCopy tasks, this user will also have to handle volume groups and filesystems.

Now that you've had an introduction to the modules needed to back up an SAP system using FlashCopy capabilities, the next section describes systematically how you can create and implement such a solution by keeping in mind the SAP rules to follow.

4.2 Implementation cookbook – single implementation steps with troubleshooting hints

This section describes all the steps that have to be performed to set up a SPLITINT backup for an SAP environment on IBM DS4000 Storage Server and AIX servers.

The reader should be familiar with the administration tasks of mySAP.com (or SAP R/3), Oracle Enterprise Server, Tivoli Storage Manager and UNIX. It would exceed the scope of a White Paper to explain the administration of all these software products. Instead, references to more detailed documentation about these topics are included.

There are also references to SAP notes, which are represented only by their number in the SAP OSS (Online Support System).

4.2.1 Install and configure the TSM client on the SAP server and backup server

On AIX, the TSM backup/archive client is typically installed in the directory */usr/tivoli/tsm/client/ba/bin* and the TSM client API is installed in */usr/tivoli/tsm/client/api/bin64*. The backup/archive client and the API maintain their own sets of configuration files. The client server options file *dsm.sys* keeps a list of TSM servers to which the client can connect. By default, the client node contacts the first server identified in the client system options file. You can specify a different server to contact by entering the *servername* option in your own client user options file *dsm.opt*.

TSM client will create and edit *dsm.opt* when called for the first time. It is not unusual to configure the TSM client first and then copy the configuration file to the API directory. See Reference [8] for more details about TSM client configuration.

Check the TSM configuration files in order to set up TSM client. The files should contain at least these entries (examples):

1. *dsm.sys*

```
SERVERNAME    STORM
TCPSEVERADDRESS  10.0.0.5
NODENAME      miles
```

2. *dsm.opt*

```
SERVERNAME    STORM
```

Three environment variables point to the TSM client executables and the configuration files. Set these variables in the OS user profiles of ora<sid> and <sid>adm:

```
DSM_DIR=<installation directory TSM user interface>
```

```
DSM_CONFIG=<path to the option file including filename> # for backup/restore client
```

```
DSMI_CONFIG=<path to the option file including filename> # for API
```

Example:

```
DSM_CONFIG=/usr/tivoli/tsm/client/ba/bin/dsm.opt
DSMI_CONFIG=/usr/tivoli/tsm/client/ba/bin/dsm.opt
DSM_DIR=/usr/tivoli/tsm/client/ba/bin/
```

The setting of these variables helps prevent the TSM client from using configuration files that you do not maintain. For instance, TSM for ERP System (which will be installed next) brings its own set of configuration files. Using the wrong configuration files may result in TSM error messages when TSM for ERP systems calls TSM.

For example:

```
BKI5014E: Tivoli Storage Manager Error:
ANS0220E (RC400) An invalid option was found during option parsing.
```

4.2.2 Install and configure TSM for ERP on the SAP server

The Tivoli product *Data Protection for mySAP.com* is usually on CD-ROM. You get a license and a usable installation only with the CD-ROM.

You will be guided through a setup procedure, which will ask for information regarding your environment. Your input will be used to edit the configuration file */oracle/<SID>/920_64/dbs/init<SID>.utl* during the setup.

Example:

```
OracleSID = PRO
Connection and login information to TSM Server
  UserID/Password: MILES/password
  Server: storm 10.0.0.5
  Node name: miles
no Oracle RMAN
TSM management classes: STANDARD
```

On AIX, *Tivoli Data Protection for mySAP.com* is installed in the directory */usr/tivoli/tsm/tdp_r3/ora64*. There you can find the executable *backint*, which acts as an interface between BR*Tools and the TSM client.

Edit the following entries in the SAP configuration file

/oracle/<SID>/920_64/dbs/init<SID>.sap

to integrate BR*Tools and *backint*:

```
# backup_type
```

```
backup_type = online
```

```
or
```

```
backup_type = offline
```

```
# backup device type
```

```
# default: tape
```

```
backup_dev_type = util_file
```

Log in to the SAP server as OS user *ora<sid>*, and start the SAP database administration tool (*sapdba*). Back up an Oracle tablespace or archive offline redo logs to check if the products you installed work together.

The installation instructions (Reference [22]) for TSM for ERP have more details.

4.2.3 Create OS users on the backup server

Create the UNIX users and groups that are needed to run SAP (groups *sapsys*, *dba* and *oper* and users *<sid>adm* and *ora<sid>*) on the backup server. Look up the users and their parameters on the SAP server *miles*.

Copy the user profile from the SAP server. SAP- and Oracle-specific profiles must be renamed because the server's hostname is part of the filename.

```
i.e. mv .dbenv_miles.csh .dbenv_dizzy.csh
     mv .sapsrc_miles.csh .sapsrc_dizzy.csh
```

In this environment the servers are called *miles* and *dizzy*.

Hint:

To avoid access permission problems, use the same UIDs and GIDs the SAP server does.

4.2.4 Create network filesystems

Export the directories `/oracle/<SID>/920_64/dbs` and `/oracle/<SID>/sapbackup` on the SAP server and mount these two directories on the backup server.

The first directory stores the configuration files of SAP, Oracle, BR*Tools and TDP. Some must be available on the backup server, and you should avoid redundant maintenance.

The other directory contains the log files of BR*Tools. In our solution, database backups will be performed on the backup server while restore processes are still performed on the SAP server. The two servers must share information about successful and unsuccessful backups.

The path names on the backup server must be the same as those on the SAP server. Also, set directory owners and permissions according to the SAP server.

4.2.5 Install and configure an Oracle client on the backup server

You can use the Oracle software CDs from the SAP CD selection to install the Oracle client, or download the software from Oracle's web page: <http://otn.oracle.com>.

Start the Oracle Installer and select Oracle Client on the Available Products screen.

After the installation is completed, the Net Configuration Assistant will start automatically. You must configure a new Oracle service name in order to access the remote database on the SAP server. Instead of creating a new configuration file, you can also copy the file `/oracle/<SID>/920_64/network/admin/tnsnames.ora` from your SAP server to the backup server. The file `tnsnames.ora` should look similar to the following:

```
PRO.WORLD=
  (DESCRIPTION =
    (SDU = 32768)
    (ADDRESS_LIST =
      (ADDRESS =
        (COMMUNITY = SAP.WORLD)
        (PROTOCOL = TCP)
        (HOST = miles)
        (PORT = 1527)
      )
    )
  )
(CONNECT_DATA =
  (SID = PRO)
  (GLOBAL_NAME = PRO.WORLD)
)
```

Take the time to test the remote database connection by logging in to the database with *sqlplus*.

Example:

```
sqlplus system/manager@PRO
```

```
SQL*Plus: Release 9.2.0.1.0 - Production on Mon Aug 9 17:06:21 2004
```

```
Copyright (c) 1982, 2002, Oracle Corporation. All rights reserved.
```

```
Connected to:
```

```
Oracle9i Enterprise Edition Release 9.2.0.4.0 - 64bit Production
```

```
With the Partitioning option
```

```
JServer Release 9.2.0.4.0 - Production
```

```
SQL> select DBID, NAME from v$database;
```

```
      DBID NAME
-----
686430656 PRO
```

```
SQL> exit
```

```
Disconnected from Oracle9i Enterprise Edition Release 9.2.0.4.0 - 64bit Production
```

```
With the Partitioning option
```

```
JServer Release 9.2.0.4.0 - Production
```

4.2.6 Copy Oracle client libraries

Up to BR*Tools version 4.6D, the programs (*sapdba*, *brbackup* etc) were statically linked to the Oracle client software. Starting with version 6.10, BR*Tool's libraries are linked dynamically and expected to be found in */oracle/client/81x_<bits>/lib* or */oracle/client/92x_<bits>/lib* (<bits> = 32 | 64). Copy the libraries from the appropriate directory on the SAP server to the backup server. If you omit this step, the BR*Tools will deliver an error message like the following:

```
exec(): 0509-036 Cannot load program sapdba because of the following errors:
```

```
0509-150 Dependent module libclntsh.a(shr.o) could not be loaded.
```

0509-022 Cannot load module libclntsh.a(shr.o).

0509-026 System error: A file or directory in the path name does not exist.

Reference: SAP Note 593582

4.2.7 Set up password file authentication

The configuration of an Oracle client enables remote access to the database. If you want to start offline backups of your database, you must be able to shut down and start up the database remotely from the backup server. Oracle provides a special security mechanism for these remote administration tasks, because only a database administrator should be able to perform these tasks.

Perform these steps to enable the authentication of an administrative user by way of a password file (see Reference [11]).

- Log on to the backup server as user ora<sid>.
- Create the password file:

```
ORAPWD FILE=filename PASSWORD=password ENTRIES=max_users  
orapwd file=$ORACLE_HOME/dbs/orapwPRO password=change_on_install entries=100
```
- Log on to the SAP server as user ora<sid>.
- Shut down the Oracle database.
- Append the following line to Oracle configuration file *init<SID>.dba*:

```
remote_login_passwordfile = exclusive
```
- Start the Oracle database.
- Grant an additional role to Oracle user “system”:

```
sqlplus /nolog  
SQL> connect / as sysdba;  
Output: Connected.  
SQL> grant sysoper to system;  
Output: Grant succeeded.  
SQL> exit
```
- Restart the Oracle database.

Now *brbackup* should be able to connect to the database with the username SYSTEM.

Reference: SAP note 131610

4.2.8 Install and configure TSM for ERP on the backup server

At this point, these prerequisites to run *brbackup* on the backup server should be available:

- configuration and log files by way of NFS
- a remote connection to an Oracle database

The procedure for installing and configuring TSM for ERP systems on the backup server is identical to what we described above. See “Install and Configure TSM for ERP on the SAP Server.”

4.2.9 Install BR*Tools

Create two directories on the backup server:

- */sapmnt/<SID>/exe*
- */usr/sap/<SID>/SYS/exe*

Set the directory owner to *<sid>adm* and the group to *sapsys*.

Create directories and links in */usr/sap/<SID>/SYS/exe*:

```
mkdir opt; ln -s /sapmnt/<SID>/exe dbg; ln -s /usr/sap/<SID>/SYS/exe/dbg run
```

These commands will create links and directories similar to this:

```
lrwxrwxrwx 1 proadm sapsys      15 May 21 10:58 dbg -> /sapmnt/PRO/exe
drwxr-xr-x 2 proadm sapsys      512 Apr 14 11:09 opt
lrwxrwxrwx 1 proadm sapsys      24 May 21 10:58 run -> /usr/sap/PRO/SYS/exe/dbg
```

This structure is necessary only because the user profiles refer to these path names. It would also be possible to edit the user profiles. We consider the first alternative to be more fault-tolerant.

Create soft links to *backint* and *backfm* executables in */sapmnt/<SID>/exe*:

```
ln -s /usr/tivoli/tsm/tdp_r3/ora64/backint backint
```

```
ln -s /usr/tivoli/tsm/tdp_r3/ora64/backfm backfm
```

backint is the interface between BR*Tools and the Tivoli Storage Manager client; *backfm* (backint Filemanager) is an expert utility that bypasses the SAP tools to inquire, delete and restore backup sets or individual files. This SAP backup solution does not use it. We created the link only to complete the Tivoli backup environment for SAP.

Download the most current BR*Tools package from SAP Service Marketplace (<http://www.sap.com/services>). At minimum, version 6.20 patch level 110 is required.

See Reference [14] for further details.

Install this package on the backup server and on the SAP server in the directory */sapmnt/<SID>/exe*.

Edit the following entries in the SAP configuration file

```
/oracle/<SID>/920_64/dbs/init<SID>.sap
```

to integrate BR*Tools and *backint* and to activate the SPLITINT process.

Example (with *<SID>* = PRO):

```
# backup_type
```

```
backup_type = online_mirror
```

```
or
```

```
backup_type = offline_mirror
```

```
# backup device type
```

```
# default: tape
```

```
backup_dev_type = util_file
```

```
# decide if resync after the backup is required (e.g. disable flash copy)
```

```
split_resync = yes
```

```
# backup utility parameter file
# default: no parameter file
util_par_file = initPRO.utl

# Oracle instance string to the primary database
# [primary_db = <inst_str> | LOCAL]
# no default
# primary_db = <inst_str>
primary_db = PRO.WORLD
```

You should now be able to open a remote connection to the SAP server with BR*Tools on the backup server.

The next steps will create the prerequisites on the DS4000 and operating system level needed to automate the volume handling.

4.2.10 Create shell scripts to handle DS4000 FlashCopy with Storage Manager CLI

DS4000 Storage Manager is designed to let you create and manage FlashCopy logical drives. A FlashCopy logical drive is a point-in-time image of a standard logical drive in your storage subsystem. The logical drive that is copied is called a base logical drive. When you initiate a FlashCopy command, the controller suspends write operations to the base logical drive for a few seconds while it creates a FlashCopy repository logical drive, a physical logical drive where FlashCopy metadata and copy-on-write data are stored. Additionally, a FlashCopy volume (or target) is created, which can be assigned to a host system.

When you create a FlashCopy volume, you can specify parameters including

- the names of the objects that are created: FlashCopy target volume, repository;
- the capacity of the FlashCopy repository's logical drive (specified as a percentage of the base logical drive);
- the threshold value for the repository logical drive capacity when a warning should be displayed; and
- the sequence number of the array on which to place the repository logical drive.

It is possible to help increase the security of a DS4000 storage server by setting a password. If a command will change the DS4000 configuration, the password must be provided in the SMcli command line (option -p) or in the CLI script.

We implemented two strategies to handle DS4000 FlashCopy:

- Create FlashCopy target volumes for each backup run of an SAP instance and delete those volumes after the backup was finished.
- Create FlashCopy target volumes for an SAP instance only once, disable the volumes after the backup finished and re-create them when a new backup is started.

The following considerations may help you decide which strategy to use:

1. Create/delete FlashCopy.

- The disk space that is needed for the FlashCopy target volumes can be used for the backup of several SAP instances.
- Each time the FlashCopy target volumes are created, they must be assigned to the backup server. These LUN (logical unit number) mappings are lost when the volumes are deleted.

2. Re-create/disable FlashCopy.

- LUN mappings are preserved, as the FlashCopy target volumes are not deleted.
- Re-creation of a FlashCopy target volume is about twice as fast as the initial creation.
- Repository and target volumes are reserved and used exclusively.

The following figures show three development stages of a DS4000 FlashCopy.

The source volumes are called SAPLOG and SAPDATA and are located in RAID arrays 4 and 6 respectively. The FlashCopy target volumes were established in the same arrays and are called SAPLOG-1 and SAPDATA-1 (the default names). The FlashCopy repository drives SAPDATA-R1 and SAPLOG-R1 were created in RAID array 2. (See Figure 22.)

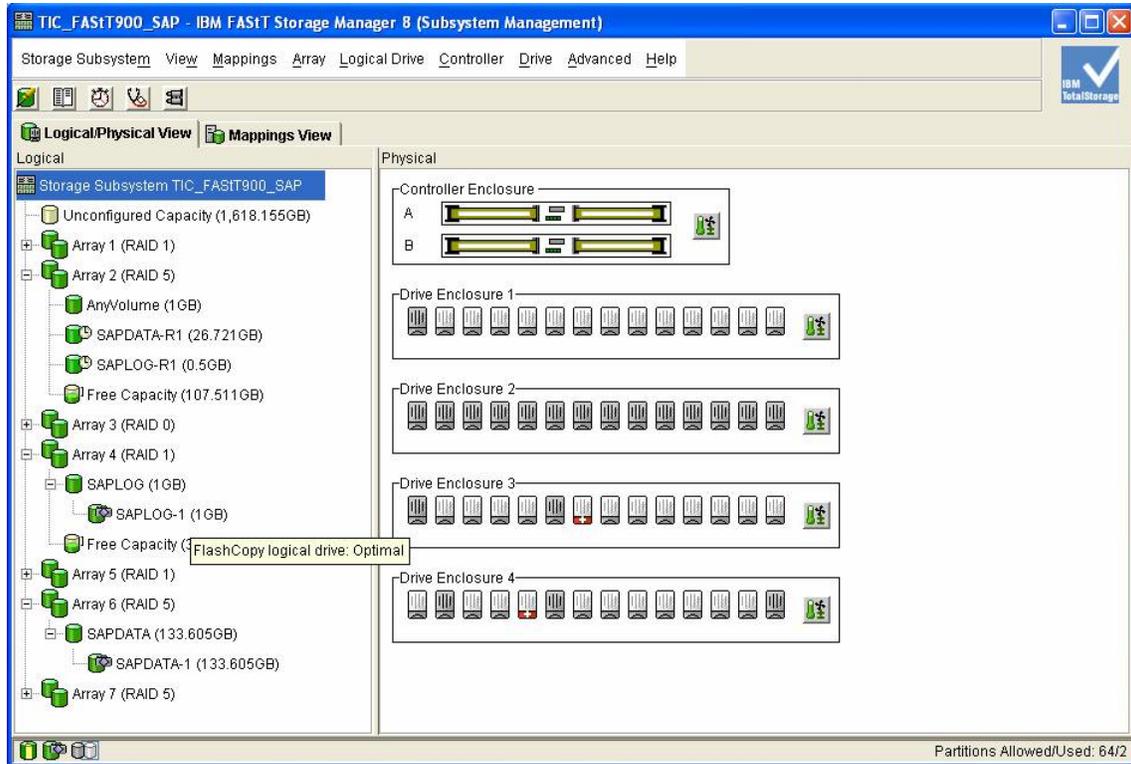


Figure 22: DS4000 Storage Manager GUI showing active FlashCopy pairs and repository drives

The FlashCopy target volumes must be assigned to a host group, host or host port to allow server access. In a SAN fabric, a server's FC adapter is identified by its worldwide name (WWN), which is a 16-digit unique number. This WWN is an attribute to the host port definition on a DS4000.

In the example (see Figure 23), the FlashCopy target volumes SAPDATA-1 and SAPLOG-1 are assigned to the host *dizzy*.

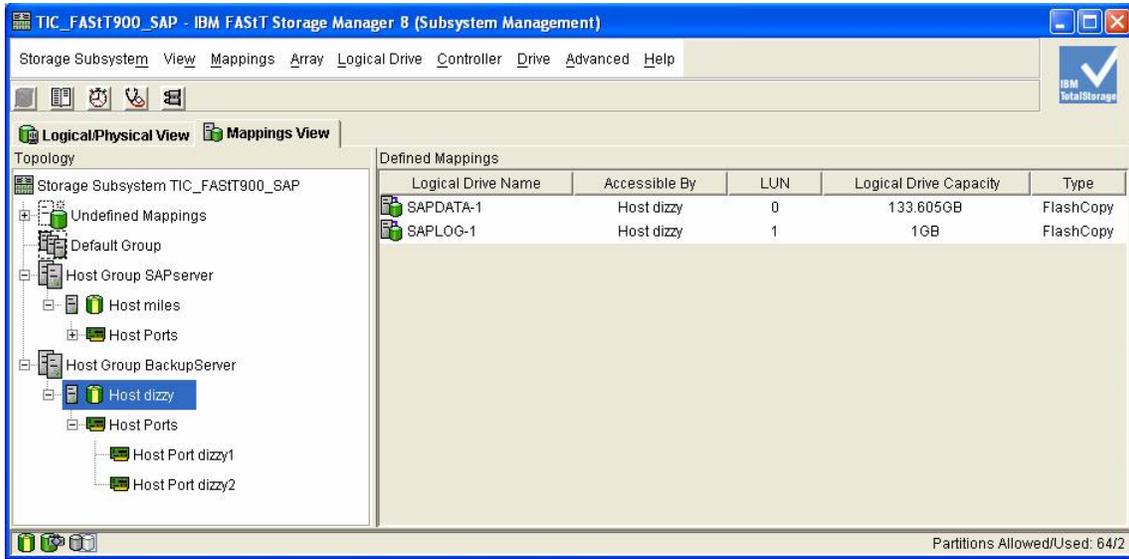


Figure 23: DS4000 Storage Manager GUI showing LUN mappings

You can either delete or disable the FlashCopy target volumes that are no longer required. If you delete the target volumes, the repository and LUN mappings are also destroyed. If you use the disable FlashCopy function, the repository and LUN mappings remain in place. Figure 24 shows the disabled FlashCopy target volumes SAPDATA-1 and SAPLOG-1 and the repository drives SAPDATA-R1 and SAPLOG-R1.

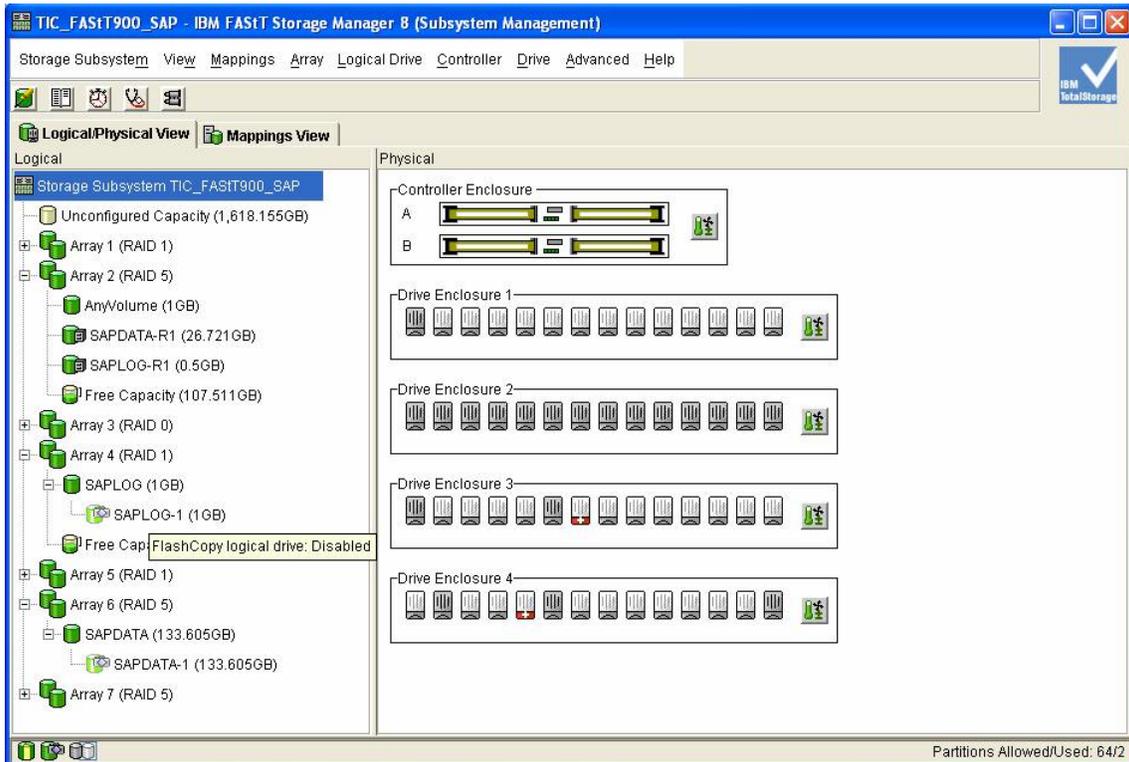


Figure 24: DS4000 Storage Manager GUI showing disabled FlashCopy pairs and repository drives

In the next section, you will find SMcli examples to handle DS4000 FlashCopy. The pictures illustrate how the DS4000 Storage Manager GUI displays the results of the CLI calls. We neither use nor explain all parameters of the copy functions. See the DS4000 documentation in References [12] and [13] for a detailed description.

Create the FlashCopy target volume

SMcli <DS4000 controller IP addresses> -c 'create FlashCopylogicaldrive baselogicaldrive="<Source Name>" repositoryarray=<Array number> repositoryPercentofBase=<Number>';'

For example:

```
SMcli 9.155.49.96 9.155.49.97 -c 'create FlashCopylogicaldrive baselogicaldrive="SAPDATA" repositoryarray=4 repositoryPercentofBase=20;'
```

Remarks:

- The IP addresses of both DS4000 controllers are required to create commands.
- If a name for the target volume is not specified, a default is used: the name of the FlashCopy source volume followed by -1.
- The location, name and size of the repository can be specified when a FlashCopy is initiated. The size is specified as percent of the source volume. The default name for the repository is the name of the FlashCopy source volume followed by -R1.

Assign volume to server

SMcli <DS4000 controller IP addresses> -c 'create mapping logicalDrive="<Target Name>" logicalUnitNumber=<Number> host="<Host Name>";'

For example:

```
SMcli 9.155.49.96 9.155.49.97 -c 'create mapping logicalDrive="SAPDATA-1" logicalUnitNumber=0 host="dizzy";'
```

Remarks:

- The host must be defined on the DS4000 as a server that can be attached to the storage server.
- The logical unit number must be unique for the host specified for that SMcli command.

Disable a FlashCopy pair

```
SMcli <DS4000 controller IP addresses> -c 'disableFlashCopy logicalDrive ["<Target Name>"];'
```

For example:

```
SMcli 9.155.49.96 9.155.49.97 -c 'disableFlashCopy logicalDrive ["SAPDATA-1"];'
```

Re-create a FlashCopy pair

```
SMcli <DS4000 controller IP addresses> -c 'recreateFlashCopy logicalDrive ["<Target Name>"];'
```

For example:

```
SMcli 9.155.49.96 9.155.49.97 -c 'recreateFlashCopy logicalDrive ["SAPDATA-1"];'
```

Delete a FlashCopy pair

```
SMcli <DS4000 controller IP addresses> -c 'delete logicalDrive ["<Target Name>"];
```

For example:

```
SMcli 9.155.49.96 9.155.49.97 -c 'delete logicalDrive ["SAPDATA-1"];
```

Troubleshooting hint:

In some situations, the FlashCopy repository is not automatically deleted with the FlashCopy pair. The displayed error message explains why: “The operation cannot complete because a selected logical drive is the last one belonging to its controller.” In this case, delete the repository using another SMcli call similar to:

```
SMcli 9.155.49.96 9.155.49.97 -c 'delete logicalDrive ["SAPDATA-R1"];
```

As long as other logical volumes have been created on the specific array, you shouldn't run into this error message.

Appendix A2 lists all scripts that have been used for our solution.

4.2.11 Create shell scripts to handle volume groups and filesystems

After the FlashCopy pairs have been created (or re-created), the FlashCopy target volumes must be configured on the backup server and the database components must be made available.

On an AIX server, Logical Volume Manager (LVM) imports the copy of the volume group that was created on the SAP server. Of course, it is possible and sometimes necessary to distribute the database components over or across several volume groups. We used only one volume group to keep our setup simple.

The following LVM actions must occur after the FlashCopy pair is created or re-created:

- scanning the system for new hard disks
- importing and activating the volume group(s)
- mounting the filesystems

Here are examples of AIX commands:

```
cfgmgr
importvg -y fcvg01 hdisk1
mount /oracle/PRO/sapdata1
...
mount /oracle/PRO/origlogA
...
mount /oracle/PRO/oraarch
mount /oracle/PRO/sapreorg
```

These LVM actions must occur before the FlashCopy pair is disabled or deleted:

- unmounting the filesystems
- deactivating and exporting the volume group(s)
- removing the hard disk special files

Here are some examples:

```
mount /oracle/PRO/sapdata1
...
mount /oracle/PRO/origlogA
...
```

```
mount /oracle/PRO/oraarch
mount /oracle/PRO/sapreorg
varyoffvg fcvg01
exportvg fcvg01
rmdev -dl hdisk1
rmdev -dl hdisk2
```

Troubleshooting hints:

If the import of the volume group fails, check the status of the special files that represent the DS4000 storage server and the DS4000 logical volumes on the backup server.

The AIX commands *lsdev -Cc disk* lists the available and configured hard disks; *lspv* shows whether the disks belong to a volume group.

The utility *fget_config*, which is installed with the DS4000 client software, lists the DS4000 controllers, the volumes and the preferred controller for each volume. Each device file *darX* represents one DS4000 storage server; the device files for the controllers are named *dacX*. If only one DS4000 is attached to your server but several device files (*dar0* and *dar1*, for example) were created for it, delete all device files for the DS4000 storage server, controllers and logical volumes, and reconfigure the devices (run *cfgmgr*). If the problem persists, check the Fibre Channel cabling and the SAN fabric zoning. As mentioned earlier, there must be two dedicated input/output paths between the DS4000 storage server and the host.

At this point, we have developed scripts to handle DS4000 FlashCopy, AIX volume groups and filesystems. We named the scripts according to their function: *createflashcopy*, *deleteflashcopy*, *recreateflashcopy*, *disableflashcopy*, *importtargetvgs*, *exporttargetvgs*, *mountfscopies* and *umountfscopies*.

Again, you will find these scripts in Appendix A2.

4.2.12 Install and configure *sudo*

In an SAP environment, backup and restore tasks are run by an unprivileged OS user called *ora<sid>*. This user will have to perform system administration tasks to integrate *brbackup* and the DS4000 FlashCopy function.

On a UNIX operating system, programs generally execute with the user and group access rights of the user who invoked them. Program owners can associate the access rights of the user who invoked them by making the program a set user ID on execution (*setuid*) or set group ID on execution (*setgid*) type. On AIX this is possible with compiled programs, but it is a security issue when using shell scripts.

Instead of developing programs in C or Perl, we installed and configured the freeware utility *sudo* (Super User Do). This useful UNIX program lets a system administrator give certain users the ability to run some (or all) commands as root.

Download *sudo* together with the installation guide from: <http://www.courtesan.com/sudo/>

If a compiled AIX version is not available, you will need a C compiler to create the *sudo* executable. A UNIX *makefile* is provided with the download package.

A readme file describes how to edit the *sudo* configuration file */etc/sudoers*, which contains a list of unprivileged users and the privileged tasks they will be allowed to execute.

Configuration file example:

```
%sapsys ALL=(ALL) NOPASSWD: ALL
%dba ALL=(ALL) NOPASSWD: ALL
```

These two lines allow all members of the user groups *sapsys* and *dba* to run system administration tasks with *sudo*.

Now, if we use *sudo*, the execution of a system administrative task by an unprivileged user is quite simple: `sudo <program or script filename>`

Here are some examples:

```
sudo importvg -y fcvg01 hdisk1
sudo importtargetvgs
```

4.2.13 Development of the SPLITINT script

In a split mirror disk environment, the SAP tool *brbackup* controls task execution:

- database operation (startup/shutdown, begin/end backup mode)
- FlashCopy handling
- LVM operation
- backup of the database and associated files

According to the interface specification by SAP (Reference [14]), *splitint* must implement two functions: mirror split and mirror resynchronization. In our environment, the *splitint* script will comply with this specification but will execute different tasks: create/delete DS4000 FlashCopy or re-create/disable FlashCopy.

Database operation and backup is initiated directly by *brbackup*. The other functions must be provided by a shell script or program called *splitint*. Communication between *brbackup* and *splitint* uses several methods:

- shell script calls
- message files
- semaphores

The tasks of *brbackup* and *splitint* and their communication methods are described in the following two schematic tables.

Split function

brbackup	Communication	splitint
brbackup starts		
brbackup recognizes <code>online_mirror</code> or <code>offline_mirror</code> option and calls <code>splitint</code> with option “-f split”		
	shell script/program call	
		Prepare Split Phase: create semaphore file & wait for deletion of semaphore file
	Send string <code>#BEGIN SPLIT</code> and semaphore filename	
Set tablespaces to backup mode (online backup) or shutdown database (offline backup) delete semaphore file		
	Semaphore file	
		Execute Split Phase: create or re-create DS4000 FlashCopy create another semaphore file & wait for deletion of semaphore file
	Send string <code>#END SPLIT</code> and semaphore filename	
End backup mode or start database		

delete semaphore file		
	Semaphore file	
		Complete Split Phase: configure new disk devices (run AIX command cfmgr) import and activate volume group(s) mount filesystems return control to brbackup
	Send string #SUCCESS	
Backup data files, control files and logs		

Figure 25: Control flow and communication channels of the split function

Resync function

brbackup	Communication	splitint
brbackup calls splitint with option “-f resync”		
	shell script/program call	
		unmount filesystems deactivate and export volume group(s) disable or delete FlashCopy return control to brbackup
	Send string #SUCCESS	
brbackup terminates		

Figure 26: Control flow and communication channels of the resync function

Hints & troubleshooting:

You can start *brbackup* with the type option `-t|type offline_mirror | online_mirror`. As usual, if no command line option is specified, the appropriate entry in the configuration file *init<SID>.sap* is used as the default:

```
backup_type = online_mirror
```

or backup_type = offline_mirror

Consult the *splitint* interface specification (Reference [14]) for the complete syntax description of the list of call parameters.

If the split mirror function fails, *brbackup* resets the tablespace backup mode to normal and terminates with an error.

The *splitint* script can be called manually with the option *-f resync* to clean up the backup server environment after an error.

If a database offline backup fails with Oracle error message ORA-12709 when *brbackup* attempts to start the database using the remote connection, the Oracle Listener is running in an incorrect environment. Check whether the listener was started correctly with *ora<sid>* and whether the correct environment (especially ORA_NLS33, NLS_LANG) was set at the start time. If necessary, restart the listener with the correct user and environment.

Reference: SAP Note 593582

The template for the configuration file *init<SID>.sap* lists the profile parameters *<split_cmd>* and *<resync_cmd>*. These two entries are used for the configuration of Split Mirror Disk Backup for Oracle together with *backup_type = offline_split | online_split*.

Do not mistake Split Mirror Disk Backup for the SPLITINT interface that we used for the backup solution described here.

As mentioned before, you can no longer use *sapdba* for database recovery starting with BR*Tools release 6.20 patch level 100. Instead, you must use *brrecover* to restore, recover or repair a database.

If you call *sapdba* by mistake, the tools will not accept the backup types *online_mirror* and *offline_mirror* that are specified in the *brbackup* logs. The associated backups will not be recognized as valid. The error messages will look similar to the following:

```
SAPDBA: Bad format in detail BRBACKUP log file.
```

```
(2004-06-07 08.23.13)
```

```
Filename: "/oracle/PRO/sapbackup/bdnjtqvy.anf"
```

```
Line : 17
```

```
Additional info: LOC_091:
```

(see /oracle/SID/sapbackup/bdnlfwjc.anf*)

Reference: SAP note 602497

This section concludes the cookbook for developing your own *splitint* interface to automate an SAP system backup using DS4000 FlashCopy. You may find our *splitint* implementation in Appendix A2 together with all other reference scripts for administering DS4000 FlashCopy or AIX LVM objects.

The procedure described within this chapter has been thoroughly tested by performing several backups (online and offline) and restores and has proved to work very reliably.

Please note that the scripts fulfill a minimum level of automation, which certainly can expand according to your individual flavor. For each individual scenario, several aspects such as security, automation and availability have to be taken into account; that prevents a universal shell script approach and exceeds the scope of these tests.

The next chapter discusses the results that have been reached within this testing and offers suggestions for improvement when adapting the proposed findings to a specific customer environment.

5 Conclusions and outlook

The scope of this White Paper is to discuss two areas of interest that have been identified. The first is how to configure an IBM TotalStorage DS4000 storage server in an SAP environment.

Several DS4000 parameters and volume layouts have been evaluated from different points of view, such as capacity management, performance and pure administrative maintenance expenses.

Although every SAP environment is rather complex and differs from all others, the summarized recommendations concerning DS4000 parameter and volume layouts in section 3.4 serve as a good starting point. They also will help you avoid unsatisfactory SAP performance, user frustration and unnecessary reconfiguration projects.

Moreover, it was shown that the SAP performance can improve even for a SAP system whose number of available disks is already at its limit.

This may enable customers to use an existing infrastructure longer instead of forcing them to make new investments for improving performance.

From the point of view of IBM support, the optimization of the DS4000 performance according to this White Paper's suggestions may help to reduce critical performance situations and complaints, therefore improving customer satisfaction.

The second area of interest is more complex, as it concerns the automation of SAP database backups using DS4000 FlashCopy. The potential advantages of this technique are well known to several high-end enterprise companies that use IBM TotalStorage Enterprise Storage Server (ESS) as their primary storage subsystem. ESS and IBM/Tivoli software offer a suite of robust products designed to automate and take advantage of ESS CopyServices functions for producing online backups.

Results from this testing allow customers to now use their IBM DS4000 storage server to help get some of the same advantages that are available with ESS. The cookbook along with the well-tested example scripts let IBM service organizations and customers alike develop their own backup strategy quickly with DS4000 FlashCopy.

The example scripts can be used in an AIX/Oracle environment right away. The scripts can also be adapted for other operating systems, widening the applicability of this approach.

Besides the main potential advantages of easy backups and fast database restoration times, customers may also benefit from efficient DS4000 FlashCopy algorithms that can help reduce the capacity needed for the FlashCopy backup by as much as 80 percent of the original capacity. Together with the possibility of reusing this disk space for different SAP system backups, this may lead to a more efficient usage of the DS4000 investment.

Moreover, customization options permit improvement according to the automation and security level required for each individual environment.

6 References

- [1] IBM Redbook Introduction to Storage Area Networks (SG24-5470-01, April 2003)
- [2] IBM Redbook IBM SAN Survival Guide (SG24-6143-01, August 2003)
- [3] IBM Redbook TotalStorage FAStT 600/900 (now DS4300/DS4500) and Storage Manager 8.4 (November 2003, SG24-7010-00)
- [4] IBM Redbook IBM TotalStorage FAStT 700 (now DS4400) and Copy Services (July 2002, SG24-6808-00)
- [5] IBM TotalStorage FAStT (now DS4000) Storage Manager Version 8.4 Installation and Support Guide for AIX, HP-UX, and Solaris (September 2003, GC26-7574-00)
- [6] AIX 5L Version 5.2 – AIX Installation in a Partitioned Environment (SC23-4382-04, July 2003)
- [7] IBM Redbook The Complete Partitioning Guide for IBM@server pSeries Servers (SG24-7039-01, October 2003)
- [8] Tivoli Storage Manager for UNIX Backup-Archive Clients Installation and User's Guide Version 5.1
- [9] Tivoli Storage Manager for AIX Administrator's Guide Version 5.2
- [10] Tivoli Storage Manager for AIX Administrator's Reference Version 5.2
- [11] Oracle9i Database Administrator's Guide, Release 2 (9.2), Chapter "Database Administrator Authentication"
- [12] IBM TotalStorage FAStT (DS4000) Storage Manager: Copy Services Users Guide (GC26-7561-00)
- [13] Help function in the Enterprise Management window of DS4000 Storage Manager (search for SMcli)
- [14] BC-BRS SPLITINT Interface for Oracle Databases (Interface Specification), June 2003
- [15] Oracle Database Administration - Split Mirror Disk Backup for Oracle (SAP White Paper), November 2003
- [16] Installation Guide: SAP® R/3 Enterprise on UNIX: Oracle - Using SAP R/3 Enterprise Core 4.70 Support Release 1, SAP R/3 Enterprise Extension Set 1.10 Document Version 1.05 - February 28, 2003
- [17] Installation Guide: SAP® Software on UNIX: OS Dependencies - For Components based on SAP® Web Application Server 6.20 Document Version 1.10 - June 13, 2002
- [18] Redbook SG244810: Understanding IBM @server pSeries Performance and Sizing
- [19] xSeries Fibre Channel: Performance and Optimization by Charles Stephan - Performance Engineer, xSeries Development Performance for I/O Subsystem, April 2001

[20] IBM TotalStorage: FAStT (now DS4000) Best Practices Guide, Redpiece 3690 by Bertrand Dufasne, Bernd Baeuml, Carlos De Nobrega, Ales Leskosek, Stephen Manthorpe and Jonathan Wright

[21] IBM FAStT 700 (now DS4400) Storage Server - AIX® Performance Measurements by Siebo Friesenborg and Bruce McNutt, October 2002

[22] IBM Tivoli Storage Manager for Enterprise Resource Planning (5698-APR): Data Protection for mySAP.com® Technology Installation & User's Guide for Oracle Version 3 Release 3. 10

Appendix

This scripts in these appendices are intended to provide examples of how several automation steps can be accomplished. .

A1 Performance collection

The following script was run with SMcli in order to collect the DS4000 performance monitor values for (3,500*5s=) 17,500 seconds.

```
on error stop;
set performanceMonitor interval=5 iterations=3500;
upload storageSubsystem file="/home/scripts/perf_data/perf04_3.txt" content=performanceStats;
```

This script was called by another script that was executed on the SAP backup server just before a series of four client copies was started. This parent script can be verified below:

```
#!/usr/bin/ksh

rsh miles lsvg -p sapvg|grep hdisk|awk '{ print $1 }' >/tmp/hdisklist.tmp
B=""
while read -r xx
do
    S=${S}${B}${xx}
    # echo $xx
done < /tmp/hdisklist.tmp
echo iostat will monitor $S

/usr/SMclient/SMcli 9.155.49.96 9.155.49.97 -f /home/scripts/perf_data/DS4000perf&
rsh miles /home/scripts/cleanarchive&
rsh miles iostat $S 5 >> iostat_miles_run4_3.txt
```

The script initially detects the hdisks that are used for the SAP environment and starts the iostat monitor for those disks only. Afterwards, it calls the previous script to collect the DS4000 performance data. The output data is collected in specific output files.

A2 splitint implementation

In this appendix, you will first find the splitint realization that calls some other functions, which also have been realized as shell scripts and are listed below.

Korn shell script function that checks the return code of a command or program and exits if an error is detected

(This function is used in all scripts that are listed below.)

```
function chk
{
RC=$?
if [ $RC -ne 0 ]; then
    echo $RC; exit $RC
fi
}
```

SPLITINIT script for SAP backup solution using the IBM DS4000 FlashCopy function

```
#!/bin/ksh
```

```
SCRIPTDIR=/home/scripts
```

```
CREFC="/usr/local/bin/sudo $SCRIPTDIR/createflashcopy"
```

```
DELFCE="/usr/local/bin/sudo $SCRIPTDIR/deleteflashcopy"
```

```
DISFC="/usr/local/bin/sudo $SCRIPTDIR/disableflashcopy"
```

```
RECFCE="/usr/local/bin/sudo $SCRIPTDIR/recreateflashcopy"
```

```
EXPVGE="/usr/local/bin/sudo $SCRIPTDIR/exporttargetvgs"
```

```
IMPVGE="/usr/local/bin/sudo $SCRIPTDIR/importtargetvgs"
```

```
MOUFS="/usr/local/bin/sudo $SCRIPTDIR/mountfscopies"
```

```
UMOFS="/usr/local/bin/sudo $SCRIPTDIR/umountfscopies"
```

```
SPLITINTLOG=/oracle/PRO/sapbackup/splitint.log
```

```
function chk
```

```
{
RC=$?
if [ $RC -ne 0 ]; then
    echo $RC >>$SPLITINTLOG 2>&1
    exit $RC
fi
}

echo SPLITINT called >>$SPLITINTLOG 2>&1
date >>$SPLITINTLOG 2>&1

# read and check splitint parameters
# that's what is expected: -u <user_id> -f <function> -i <in_file> -t
<type> -s <split options>
# this script only uses the parameters <function> and <split_options>
# where:
# <function> = split|resync
# <split options> = configuration file for br*tools/DS4000 flashcopy
integration
#
echo $* >>$SPLITINTLOG 2>&1
KEYWORD2=$3
KEYWORD5=$9
#echo $KEYWORD2 $KEYWORD5 >>$SPLITINTLOG 2>&1

case $KEYWORD2 in
-f)
    FUNCTION=$4
    case $KEYWORD5 in
-s)
        #echo split options parameter found >>$SPLITINTLOG 2>&1
        SPLITOPTIONS=${10}
        ;;
*)
        #echo split options parameter not found >>$SPLITINTLOG 2>&1
        SPLITOPTIONS=" "
        ;;

```

```
    esac    # end case KEYWORD5
        ;;
*)
    echo Error: wrong parameter sequence >>$$SPLITINTLOG 2>&1
    exit 98
        ;;
esac    # end case KEYWORD2

echo Function: $FUNCTION Split options: $SPLITOPTIONS >>$$SPLITINTLOG 2>&1

case $FUNCTION in

split)

    SEM1=/tmp/semaphore1
    SEM2=/tmp/semaphore2

    # create semaphore file to communicate with brbackup
    touch $SEM1

    # inform brbackup that split preparation has been completed
    echo "#BEGIN SPLIT" $SEM1

    # wait for completion of brbackup which alters the db tablespaces to
    backup mode or stops the db
    # max. time expected for this operation: 4 seconds
    while ( [ -e $SEM1 ] )
    do
        sleep 1
    done

    # execute the SPLIT here
    RETCODE=0

    # recreate flash copies
    $CREFC >>$$SPLITINTLOG 2>&1
```

```
RC=$?; let RETCODE=$RETCODE+$RC

# create semaphore file to communicate with brbackup
touch $SEM2

# inform brbackup that split has been completed
echo "#END SPLIT" $SEM2

# wait for completion of brbackup which starts the database
# max. time expected for this operation: 4 seconds
while ( [ -e $SEM2 ] )
do
    sleep 1
done

if [ $RETCODE -ne 0 ]; then
    echo Error: split phase did not succeed >>$SPLITINTLOG 2>&1
    exit $RETCODE
fi

# complete split phase here
# import volume groups, mount filesystems etc.
RETCODE2=0

# import VG
$IMPVG >>$SPLITINTLOG 2>&1
RC=$?; let RETCODE2=$RETCODE2+$RC
# mount file systems
$MOUFS >>$SPLITINTLOG 2>&1
RC=$?; let RETCODE2=$RETCODE2+$RC

if [ $RETCODE2 -ne 0 ]; then
    echo MZ Complete split phase did not succeed >>$SPLITINTLOG 2>&1
    exit $RETCODE2
fi
```

```
# return control to brbackup which finally performs the backup
echo "#SUCCESS"

        ;;
# case statement - end of split section

resync)

        # execute the RESYNC here
        RETCODE=0

        # umount file systems
        $UMOFS >>$SPLITINTLOG 2>&1
        RC=$?; let RETCODE=$RETCODE+$RC
        # export VG
        $EXPVG >>$SPLITINTLOG 2>&1
        RC=$?; let RETCODE=$RETCODE+$RC
        # disable FlashCopy
        $DELFC >>$SPLITINTLOG 2>&1
        RC=$?; let RETCODE=$RETCODE+$RC

        if [ $RETCODE -ne 0 ]; then
                echo Error: resync phase did not succeed >>$SPLITINTLOG 2>&1
                exit $RETCODE
        fi
        # return control to brbackup
        echo "#SUCCESS"

        ;;
# case statement - end of resync section

*)
        echo Error: unknown splitint function >>$SPLITINTLOG 2>&1
        exit 99
        ;;
esac
```

```
exit 0
```

```
create FlashCopy for DS4000 logical volumes sapdata and saplog and assign  
the new LUNs to the SAP backup server "dizzy"
```

```
#  
# both IP Addresses are required for create commands  
# providing only one IP address results in error message 1011  
SMcli 9.155.49.96 9.155.49.97 -c 'create FlashCopylogicaldrive  
baselogicaldrive="SAPDATA" repositoryarray=2 repositoryPercentofBase=20;'  
chk  
SMcli 9.155.49.96 9.155.49.97 -c 'create FlashCopylogicaldrive  
baselogicaldrive="SAPLOG" repositoryarray=2 repositoryPercentofBase=50;'  
chk  
SMcli 9.155.49.96 9.155.49.97 -c 'create mapping logicalDrive="SAPDATA-1"  
logicalUnitNumber=0 host="dizzy";' ; chk  
SMcli 9.155.49.96 9.155.49.97 -c 'create mapping logicalDrive="SAPLOG-1"  
logicalUnitNumber=1 host="dizzy";' ; chk  
exit 0
```

```
delete FlashCopy target volumes sapdata-1 and saplog-1
```

```
#!/bin/ksh  
#  
SMcli 9.155.49.96 9.155.49.97 -c 'delete logicalDrive ["SAPDATA-1"];'  
SMcli 9.155.49.96 9.155.49.97 -c 'delete logicalDrive ["SAPLOG-1"];'  
  
# Delete repository manually. Reason:  
# Unable to delete logical drive using the Delete Logical Drive command at  
line 1.  
# Error 189 - The operation cannot complete because a selected  
# logical drive is the last one belonging to its controller  
# owner and could not be successfully deleted.  
SMcli 9.155.49.96 9.155.49.97 -c 'delete logicalDrive ["SAPDATA-R1"];'  
SMcli 9.155.49.96 9.155.49.97 -c 'delete logicalDrive ["SAPLOG-R1"];'
```

exit 0

recreate FlashCopy target volumes sapdata-1 and saplog-1

#

```
SMcli 9.155.49.96 9.155.49.97 -c 'recreateFlashCopy logicalDrive  
["SAPDATA-1"];'; chk
```

```
SMcli 9.155.49.96 9.155.49.97 -c 'recreateFlashCopy logicalDrive  
["SAPLOG-1"];'; chk
```

exit 0

disable FlashCopy target volumes sapdata-1 and saplog-1

#

```
SMcli 9.155.49.96 9.155.49.97 -c 'disableFlashCopy logicalDrive  
["SAPDATA-1"];'; chk
```

```
SMcli 9.155.49.96 9.155.49.97 -c 'disableFlashCopy logicalDrive ["SAPLOG-  
1"];'; chk
```

exit 0

import volume group on DS4000 FlashCopy target volumes

#!/bin/ksh

cfgmgr

fget_config -l dar0

importvg -y fcvg01 hdisk1; chk

exit 0

export volume group on DS4000 FlashCopy target volumes

```
#!/bin/ksh
```

```
varyoffvg fcvg01; chk  
exportvg fcvg01; chk  
rmdev -dl hdisk1; chk  
rmdev -dl hdisk2; chk  
exit 0
```

mount SAP-specific file systems

```
#!/bin/ksh
```

```
mount /oracle/PRO/sapdata1; chk  
mount /oracle/PRO/sapdata2; chk  
mount /oracle/PRO/sapdata3; chk  
mount /oracle/PRO/sapdata4; chk  
mount /oracle/PRO/origlogA; chk  
mount /oracle/PRO/origlogB; chk  
mount /oracle/PRO/mirrlogA; chk  
mount /oracle/PRO/mirrlogB; chk  
mount /oracle/PRO/oraarch; chk  
mount /oracle/PRO/sapreorg; chk  
df -k|grep PRO  
exit 0
```

unmount SAP-specific file systems

```
#!/bin/ksh
```

```
umount /oracle/PRO/sapdata1; chk  
umount /oracle/PRO/sapdata2; chk  
umount /oracle/PRO/sapdata3; chk  
umount /oracle/PRO/sapdata4; chk
```

```
umount /oracle/PRO/origlogA; chk
umount /oracle/PRO/origlogB; chk
umount /oracle/PRO/mirrlogA; chk
umount /oracle/PRO/mirrlogB; chk
umount /oracle/PRO/oraarch; chk
umount /oracle/PRO/sapreorg; chk
df -k|grep PRO
exit 0
```

SAN Fabric Zoning

```
IBM_2109_F16:admin> alishow
Defined configuration:
cfg:   SAPonDS4000
       miles_DS4000900_2; miles_DS4000900_1; dizzy_DS4000900_1;
       dizzy_DS4000900_2
zone:  dizzy_DS4000900_1
       DS4000900_1; p670dizzy1
zone:  dizzy_DS4000900_2
       DS4000900_2; p670dizzy2
zone:  miles_DS4000900_1
       DS4000900_1; p670miles3
zone:  miles_DS4000900_2
       DS4000900_2; p670miles2
alias: DS4000900_1
       20:04:00:a0:b8:0f:82:78
alias: DS4000900_2
       20:05:00:a0:b8:0f:82:78
alias: p670dizzy1
       10:00:00:00:c9:30:95:66
alias: p670dizzy2
       10:00:00:00:c9:2e:27:0e
alias: p670miles1
       10:00:00:00:c9:2c:e5:ab
alias: p670miles2
```

10:00:00:00:c9:2e:26:42

alias: p670miles3

10:00:00:00:c9:2e:26:20

Effective configuration:

cfg: SAPonDS4000

zone: dizzy_DS4000900_1

20:04:00:a0:b8:0f:82:78

10:00:00:00:c9:30:95:66

zone: dizzy_DS4000900_2

20:05:00:a0:b8:0f:82:78

10:00:00:00:c9:2e:27:0e

zone: miles_DS4000900_1

20:04:00:a0:b8:0f:82:78

10:00:00:00:c9:2e:26:20

zone: miles_DS4000900_2

20:05:00:a0:b8:0f:82:78

10:00:00:00:c9:2e:26:42

Notices and Disclaimers

Copyright © 2004 by International Business Machines Corporation. All Rights Reserved.

No part of this document may be reproduced or transmitted in any form without written permission from IBM Corporation.

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This information could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or programs(s) at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead. It is the user's responsibility to evaluate and verify the operation of any non-IBM product, program or service.

Performance data for IBM and non-IBM products and services contained in this document was derived under specific operating and environmental conditions. The actual results obtained by any party implementing such products or services will depend on a large number of factors specific to such party's operating environment and may vary significantly. IBM makes no representation that these results can be expected or contained in any implementation of any such products or services.

THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR INFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. IBM is not responsible for the performance or interoperability of any non-IBM products discussed herein.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Trademarks

IBM, the IBM logo, FlashCopy, AIX, Enterprise Storage Server, Tivoli, pSeries, eServer e-business on demand, Ultrium, Linear Tape Open, and TotalStorage are trademarks of International Business Machines Corporation in the United States, other countries, or both.

Please visit our alliance website: www.ibm-sap.com

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java is a registered trademark of Sun Microsystems in the United States and other countries.

Sudo is distributed under the following ISC-style license:

Copyright (c) 1994-1996,1998-2004 Todd C. Miller
<Todd.Miller@courtesan.com>

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

Other company, product or service names may be trademarks or service marks of others.