



Technical Report

Solutions for SAP Test Data Migration Server with NetApp Storage

NetApp SAP Competence Center, NetApp
May 2010 | TR-3797

EXECUTIVE SUMMARY

This document describes the best practices for deploying NetApp® storage solutions in support of SAP® Test Data Migration Server (TDMS) to improve the efficiency and speed of the data migration process.

TABLE OF CONTENTS

| | | |
|----------|--|----------|
| 1 | INTRODUCTION | 3 |
| 1.1 | SCOPE AND ASSUMPTIONS | 3 |
| 1.2 | BUSINESS CHALLENGES FACING THE SAP CUSTOMER | 3 |
| 1.3 | TECHNOLOGY CHALLENGES OF AN EXPANDING SAP LANDSCAPE | 3 |
| 2 | NETAPP STORAGE IN AN SAP TEST DATA MIGRATION SERVER ENVIRONMENT | 4 |
| 2.1 | SOLUTIONS IN A HOMOGENEOUS NETAPP STORAGE LANDSCAPE | 5 |
| 2.2 | SOLUTIONS IN A HETEROGENEOUS STORAGE LANDSCAPE | 6 |
| 3 | CONCLUSION | 8 |
| 4 | REFERENCES | 9 |

1 INTRODUCTION

This document describes the storage solutions used to support an environment that uses the SAP Test Data Migration Server product. It addresses the architecture and processes to efficiently integrate TDMS into a NetApp storage environment. It discusses scenarios in which the storage landscape is only NetApp and in which third-party storage is installed along with NetApp storage.

1.1 SCOPE AND ASSUMPTIONS

The recommended solutions in this document are generic with respect to most SAP applications as well as to the size and scope of the implementation. This technical report assumes a basic understanding of the technology and operation of NetApp, SAP, and databases.

1.2 BUSINESS CHALLENGES FACING THE SAP CUSTOMER

Corporations deploying SAP software today are under pressure to reduce cost, minimize risk, and control change by accelerating deployments and increasing the availability of their SAP landscapes. Business leaders and project managers must coordinate with corporate IT management to optimize the scheduling and availability of systems to support rapid prototyping and development, frequent parallel testing or troubleshooting, and appropriate levels of end-user training.

The ability to access these systems as project schedules dictate, with current data sets and without affecting production operations, often makes the difference between SAP projects being delivered on time and within budget or not. Frequently, the volume of data in production systems limits the ability to quickly and efficiently create test and development systems. The provisioning, management, and expense of storage resources to support complete copies of production data can be prohibitive. Furthermore, once the system is copied, it must be reconfigured to be able to interface with other SAP systems in the landscape. It may also be necessary to remove sensitive financial or human-resource information to protect the company and its employees. All of these factors contribute to increasing the cost and complexity of managing an SAP environment.

1.3 TECHNOLOGY CHALLENGES OF AN EXPANDING SAP LANDSCAPE

A typical SAP production landscape consists of several different SAP systems. Just as important to the successful operation and management of these production instances are the many nonproduction instances used to support them. The ability to maintain nonproduction SAP instances is particularly important when the business initiates a new SAP project. In many cases, project development begins with the creation of a new SAP instance from a copy of the current production system.

When an SAP instance is created, either from a copy of an existing SAP system or through an SAP installation, additional instances are established based on the needs of the customer and their development and integration teams. SAP has long encouraged its customers to maintain separate development and test instances for each production instance. In practice, SAP's standard three-system (development, test, and production) landscapes often expand to include separate instances for technical sandbox, business development sandbox, development consolidation, and user training.

Multiple test systems are also common. Sometimes referred to as quality assurance or staging instances, they are typically the last instance in which changes or modifications to the SAP system reside before being moved into production. Another characteristic of these systems is that they can be exact replicas of the IT infrastructure used to support the production instance. It is possible that for every production instance of SAP, up to 10 copies of that instance could be used to create and support it. Compound this with the many different SAP applications, such as ERP, CRM, BI, SCM, SRM, and Portal, and the number of systems to support can become very large.

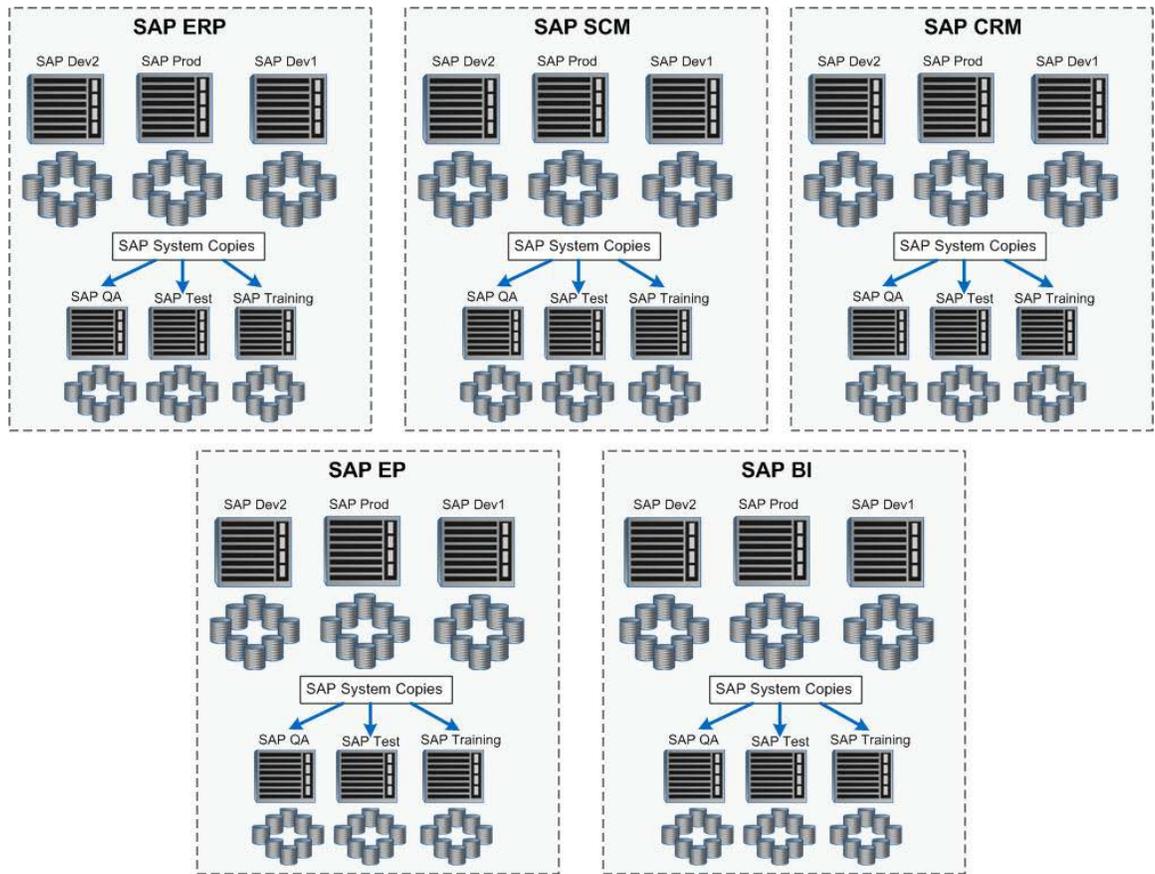


Figure 1) Example of an SAP landscape with many systems.

At various points in the application lifecycle, the instance copy is resynchronized with production to make sure that tests are performed on the current data. The speed and efficiency of this process affect both instances and any projects that rely on those systems. The demand on IT to produce and manage nonproduction instances varies from customer to customer and from project to project.

Two other factors increase the demands on the IT infrastructure:

- More test and training systems are needed to accelerate test cycles by allowing parallel independent operation.
- Test and training systems need to be refreshed often to improve the quality of testing and training.

If the infrastructure that is supporting SAP systems and related applications is inflexible and difficult to operate or manage, or if high cost-of-ownership barriers develop within IT, the ability of business owners to deploy new and improved business processes can be negatively affected.

2 NETAPP STORAGE IN AN SAP TEST DATA MIGRATION SERVER ENVIRONMENT

SAP customers constantly look for ways to simplify and streamline their IT operations. As databases supporting SAP applications grow in size, there is a desire to reduce the amount of data that is replicated in development and test environments. However, for proper testing to occur, development and test systems require a representative group of data. With the highly integrated nature of SAP modules, it is not possible to simply eliminate some tables and keep others. As a result, SAP has developed a solution called SAP Test Data Migration Server, which extracts data from a source

system, based on criteria set by the customer, and imports it into a test or development system. TDMS can significantly reduce the amount of data that is duplicated in SAP development and test systems.

TDMS is a solution for migrating transaction data from production environments to test environments. SAP has found that 10% to 20% of tables in an SAP system contain 80% to 90% of the data. As a result, the amount of data in a test system can be greatly reduced by setting parameters to migrate only a subset of the transaction data in the production system. Customized data and customer tables are migrated entirely.

2.1 SOLUTIONS IN A HOMOGENEOUS NETAPP STORAGE LANDSCAPE

A recommendation of the TDMS solution—and a requirement in some cases—is that no changes should be made to the production source system database while the TDMS data selection processes are reading the data. If the source system is active, data inconsistency at the target system is possible. NetApp storage works in conjunction with TDMS to allow data integrity without downtime at the production server. A clone of the production system can be created on the NetApp storage by using NetApp FlexClone[®] functionality. The cloned system can then be used as the source for the TDMS processing, eliminating the need for downtime on the production system. Because the FlexClone copy is created based on a Snapshot[™] copy of the source production system, very little storage is used. Also, the clone can be deleted after the TDMS process is complete, eliminating the need to maintain a complete copy of the production database. NetApp FlexClone functionality can also be leveraged to take a point-in-time Snapshot copy of a complete SAP landscape consisting of, for example, SAP ERP, BW, and CRM, thus providing a consistent data set for TDMS activities extracting a consistent subset of data across the complete SAP environment.

TDMS allows data selection from a certain date to the current date. For some customers, this could cause a regulatory conflict. Data from the current financial period could be imported into the development and test systems, giving employees access to sensitive information. NetApp technology allows customers to choose a Snapshot copy of the production system prior to the current reporting period to use as a source for the data extraction so that sensitive data is not part of the extraction. If sensitive data is not an issue, FlexClone copies can be created on a regular basis to make sure that the latest data from the production system is migrated to the test system by using the TDMS process.

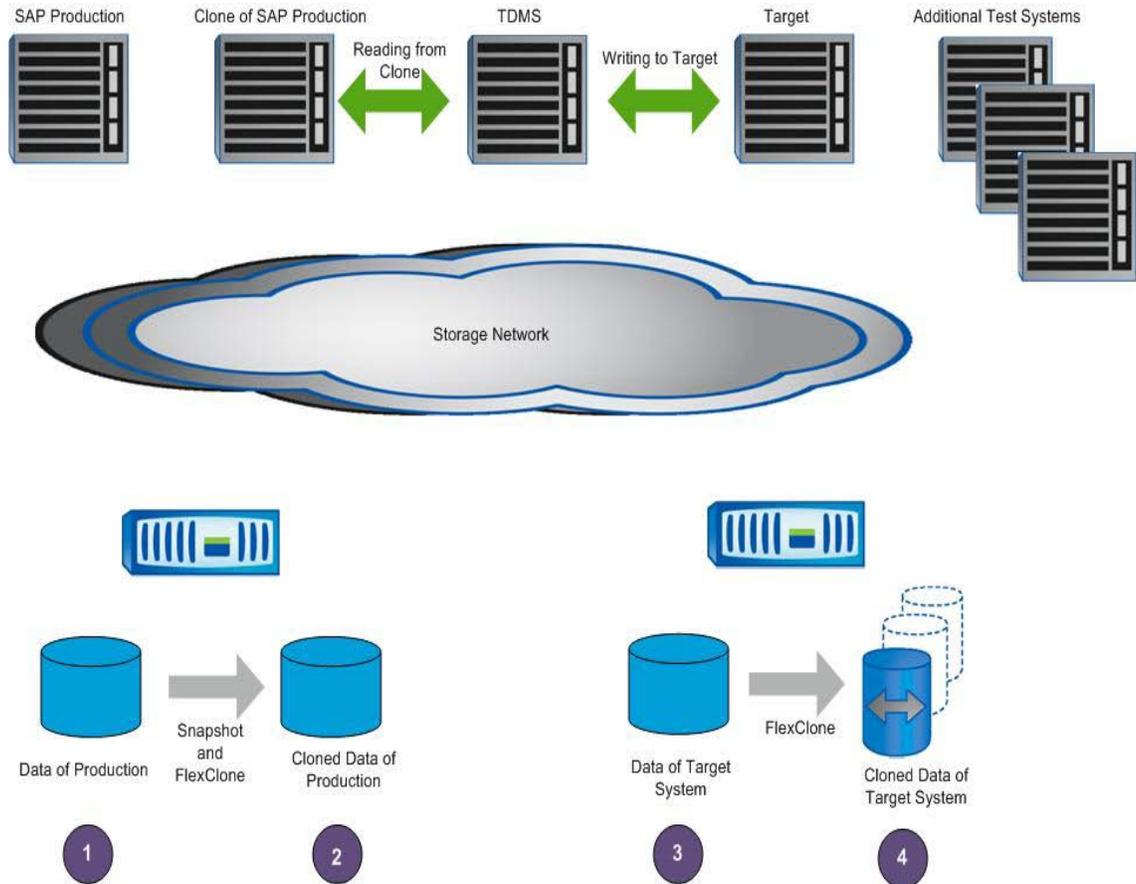


Figure 2) NetApp storage in an SAP TDMS environment.

Figure 2 illustrates the process of using NetApp storage in an SAP TDMS environment. Several steps are performed to migrate the data between the source production system and the target test systems. The steps in the figure correspond to the following activities:

1. A copy of the data in the production system is requested.
2. A clone of the production data is created by creating a FlexClone copy based on a Snapshot copy. There is no impact on the production system.
3. TDMS reads data from the FlexClone copy of the source system and writes to the target system.
4. Additional target systems can be created as FlexClone copies, so that the data in the target system is not changed and can be reused for multiple test cycles.

2.2 SOLUTIONS IN A HETEROGENEOUS STORAGE LANDSCAPE

If the source system does not reside on NetApp storage, then it is not possible to create a FlexClone copy for the extraction. Instead, you have two options for migrating the data onto NetApp storage to take advantage of FlexClone copies for testing.

Figure 3 illustrates one option for migrating data in a TDMS environment with a heterogeneous storage infrastructure. The first step is to use a replication technology to move the data from the source system to NetApp storage. When the system has been migrated, a FlexClone copy can be made and the TDMS extraction process can be run to move the source data into the target system. Once the data is in the target system, it can be cloned multiple times to allow parallel testing to occur or to reduce the need for repeated data loads. This method allows the SAP production system to remain active during the data extraction.

1. Production data is replicated from the non NetApp storage to NetApp storage.
2. A copy of the replicated data is requested for a TDMS extraction.
3. A clone of the replicated data is created by creating a FlexClone copy based on a Snapshot copy. There is no impact on the replication process.
4. TDMS reads data from the FlexClone copy of the source system and writes to the target system.
5. Additional target systems can be created as FlexClone copies, so that the data in the target system is not changed and can be reused for multiple test cycles.

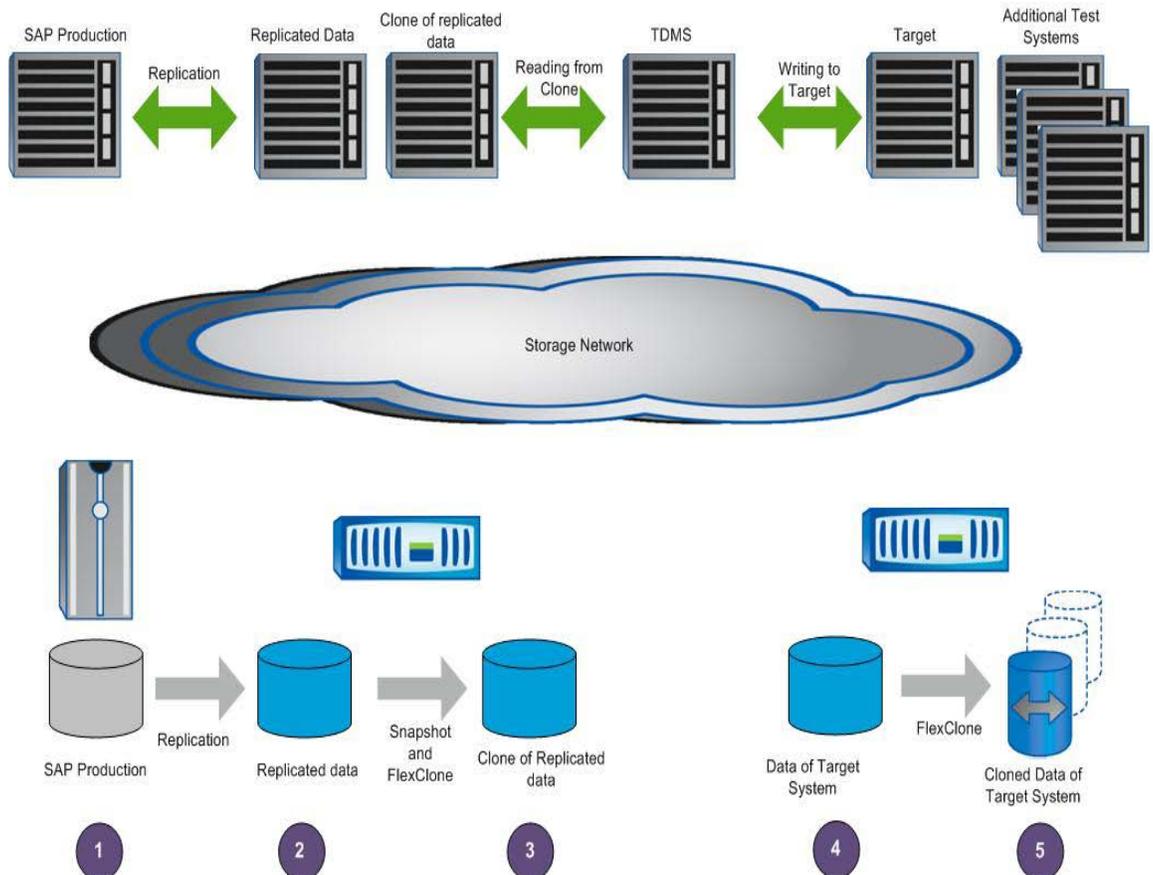


Figure 3) SAP TDMS with heterogeneous storage and replication.

Figure 4 illustrates the other possible scenario for customers with a heterogeneous storage infrastructure. Data is extracted directly from the SAP production system to a target system on NetApp storage. Although this method is not ideal with regard to data inconsistency and production availability, it does offer benefits because the target system is on NetApp storage. When creating multiple test systems from the original target system by using FlexClone copies, it is not necessary to perform the TDMS extracts repeatedly to refresh the test systems. As a result, the production system is not affected as frequently.

1. TDMS processes are run directly on the production system.
2. TDMS reads data directly from the SAP production system and writes to the target system.

- Additional target systems can be created as FlexClone copies, so that the data in the target system is not changed and can be reused for multiple test cycles.

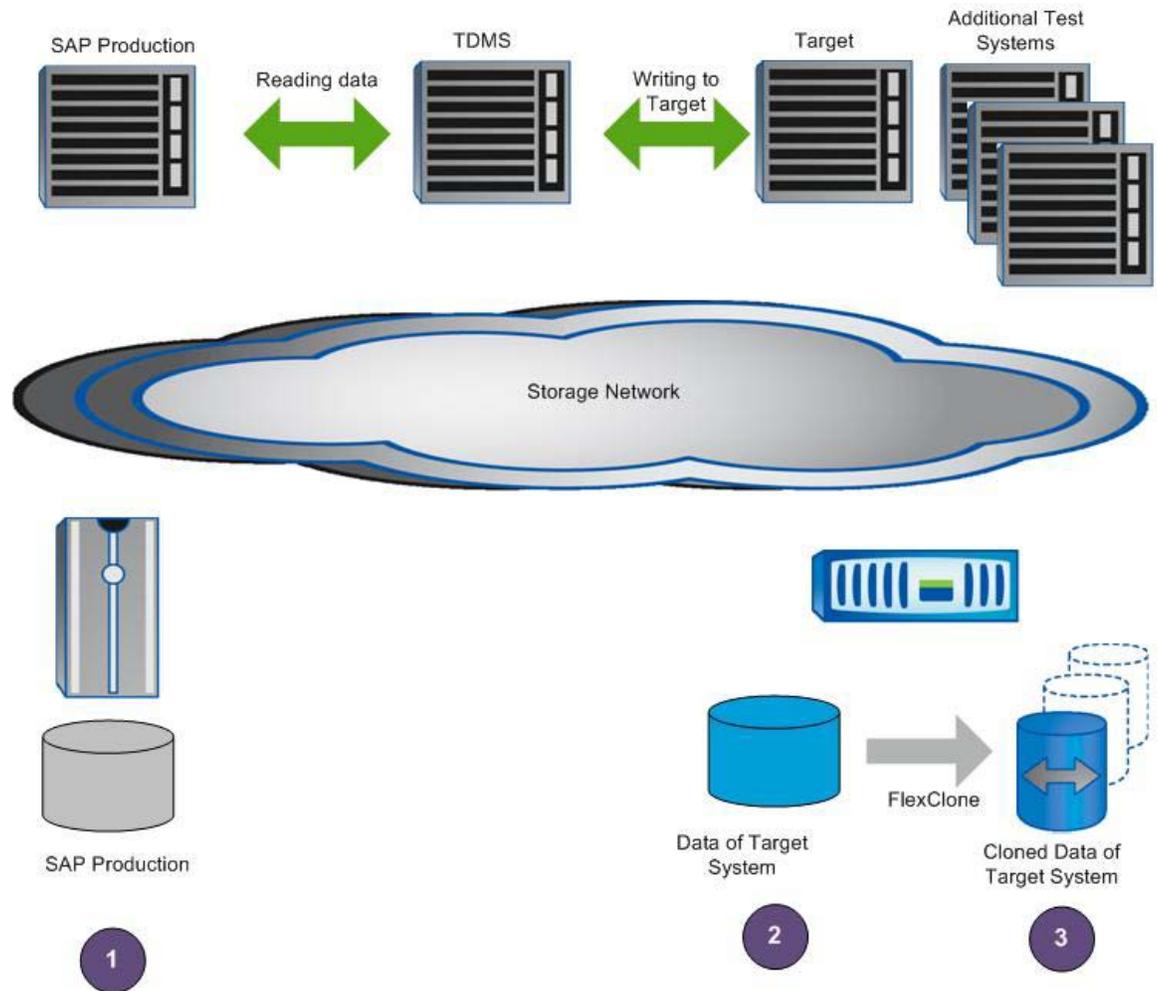


Figure 4) SAP TDMS with heterogeneous storage and direct extract from the source system.

3 CONCLUSION

A key requirement for successfully managing an SAP environment is the ability to have current data for test, development, training, and sandbox systems. However, because of the explosive growth of data in SAP production systems, it is not always feasible to create complete copies for these purposes. SAP developed TDMS to overcome this obstacle. When combined with NetApp Snapshot and FlexClone technologies to allow the fast and efficient creation of SAP systems, TDMS extractions can happen without limiting production availability. Furthermore, NetApp Snapshot copies allow customers more flexibility when choosing the end date for the extraction, thus greatly reducing regulatory risks. Once the data has been extracted, multiple copies can be made to allow parallel testing and development activities without changing the data in the original target. If more tests need to be performed, it is possible to copy the data again without having to run the data extract, saving time and resources.

NetApp storage is the perfect solution for SAP development and test environments with TDMS. NetApp storage enables customers to reduce their infrastructure costs by using less storage. It also simplifies administration and reduces risk by allowing more testing with higher-quality data in a shorter period of time.

4 REFERENCES

SAP ON NETAPP

[Enterprise Application Development and Test](#)

[NetApp and SAP Partnership](#)

[TR-3533—SAP on UNIX/Oracle/FCP](#)

[TR-3540—SAP System Copy with NetApp Storage](#)

[TR-3444—SAP on Windows/Oracle](#)

[TR-3585—SAP on Windows/SQL Server Best Practices](#)

[TR-3682—SnapManager for SAP Best Practices](#)

SAP

[SAP Test Data Migration Server](#)

[Global SAP Homepage](#)

[SAP Service Marketplace](#)

[SAP Developer Network](#)

[SAP Help Portal](#)

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein must be used solely in connection with the NetApp products discussed in this document.



NetApp

www.netapp.com

© Copyright 2010 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, FlexClone, and Snapshot are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. SAP is a registered trademark of SAP AG. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-3797