



---

ZFS STORAGE  
APPLIANCE

An Oracle Technical White Paper  
September 2014

# Implementing Oracle ZFS Storage Appliance OST Optimized Duplication in a Symantec NetBackup Data Protection Environment

## Table of Contents

Introduction .....	2
Solution Fundamentals .....	3
Installation and Configuration Overview.....	4
Remote Site Duplication Implementation .....	5
Implementation Best Practices.....	6
Using a Common Media Server for Duplication .....	6
Setting Up and Configuring the Storage .....	6
Setting Oracle ZFS Storage Appliance Pool Size .....	6
Using Oracle ZFS Storage Appliance Projects and Shares.....	7
Setting Up and Configuring NetBackup .....	8
NetBackup Storage Server Access .....	8
NetBackup Disk Pool and Storage Unit Configuration .....	8
Configuring a NetBackup Storage Lifecycle Policy (SLP).....	9
Tuning the Operational Configuration .....	9
Supporting Mixed Windows and UNIX Environments .....	10
Service Interruption Considerations .....	10
Conclusion .....	11
References.....	12

## Introduction

Symantec NetBackup introduced the OpenStorage Technology (OST) API which enables hardware vendors of intelligent storage systems to perform certain backup and recovery operations that were previously processed at the media server level. These operations include backup image copy creation, data deduplication, data replication, and writing images directly to tape. Typically, these operations have been processed at the NetBackup media server level, utilizing the media server resources as well as possibly inefficiently utilizing network bandwidth. With OST, the media server still coordinates the activity, but the processing is done at the storage appliance level.

As an initial pathway to bridge the NetBackup OST capabilities with the advanced features of the Oracle ZFS Storage Appliance, the Oracle ZFS Storage Appliance Plug-in for Symantec NetBackup OpenStorage v1.0 has been tailored for OST Optimized Duplication. This plugin enables the efficient copying (duplication) of backup image data directly between two Oracle ZFS Storage Appliance products, utilizing their inherent replication technology. Oracle ZFS Storage Appliance shares (either NFS or CIFS/SMB) are used as NetBackup storage units, allowing for manual or automated image copying.

This white paper describes the fundamentals regarding the plug-in's OST Optimized Duplication implementation with the Oracle ZFS Storage Appliance and provides an overview of the installation process. Additionally, it introduces a remote site use case and overall implementation best practices. It is intended for an audience already familiar with Symantec's Veritas NetBackup and the Oracle ZFS Storage Appliance.

For readability purposes for this paper, the plug-in may be referred to as the Oracle ZFS Storage Appliance OST plug-in. However, note that the full name -- Oracle ZFS Storage Appliance Plug-in for Symantec NetBackup OpenStorage -- is always used in formal references as well as the Oracle ZFS Storage Appliance plug-in downloads location on the Oracle Technology Network site. See the References section at the end of this document for this and other helpful technical documentation listings.

## Solution Fundamentals

The Oracle ZFS Storage Appliance OST plug-in is used to quickly copy Symantec NetBackup backup images between two Oracle ZFS Storage Appliance products. It offloads the processing from the media servers to the Oracle ZFS Storage Appliance by utilizing continuous replication technology. The general dataflow is shown in Figure 1.

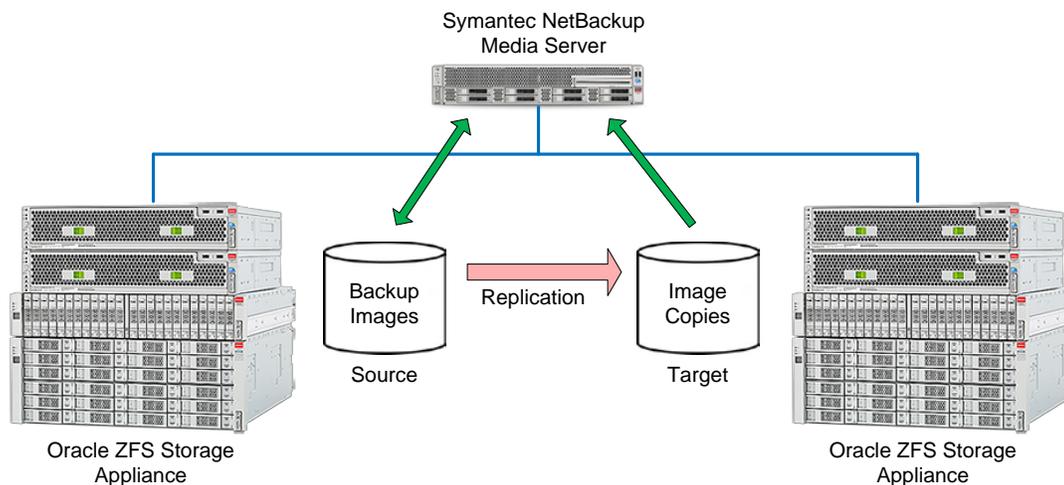


Figure 1. OST Optimized Duplication dataflow with Oracle ZFS Storage Appliance

The Oracle ZFS Storage Appliance products are configured as storage servers in NetBackup, and their shares are then configured as disk pools and storage units. The Source storage unit is used as the initial destination for backups that are expected to be duplicated. The Target storage unit is used exclusively for storing copies of the backup images contained on the Source storage unit.

A key benefit to this solution is that the data backed up to the Source starts the replication process to the Target immediately, before NetBackup even initiates the duplication request. Thus, when the actual NetBackup duplication process is invoked, it only takes seconds to complete, even with large images. In order for NetBackup to recognize the duplicated images in its catalog, either a manual copy operation must be started, or the copy must be invoked through a NetBackup Storage Lifecycle Policy (SLP). Regardless of the NetBackup duplication method, the action must be performed by a media server that has access to both the Source and the Target. Otherwise, the operation will fail.

## Installation and Configuration Overview

The high-level steps for installation and configuration are:

- 1 Download and install the plug-in on the designated NetBackup media servers according to the platform-specific installation instructions. The installation will stop and restart the NetBackup services on the server, so installation should be performed at a time when the desired media server(s) can be modified. **The plug-in must be installed on each media server that will have access to any of the NetBackup storage units.**
- 2 Configure the Source Oracle ZFS Storage Appliance replication service, select or create the appropriate pool, and create and configure a project to be used as the image repository on the Source Oracle ZFS Storage Appliance. Then, replicate and configure the project on the Target Oracle ZFS Storage Appliance, and create and configure the shares.
- 3 In NetBackup, configure each Oracle ZFS Storage Appliance as a storage server, assign the shares to their respective disk pools, and create the associated storage units. The Source storage unit is now available as a destination for backup images, and the Target storage unit can receive duplicated images from the Source storage unit.

Figure 2 provides a logical representation of the bottom-up configuration elements.

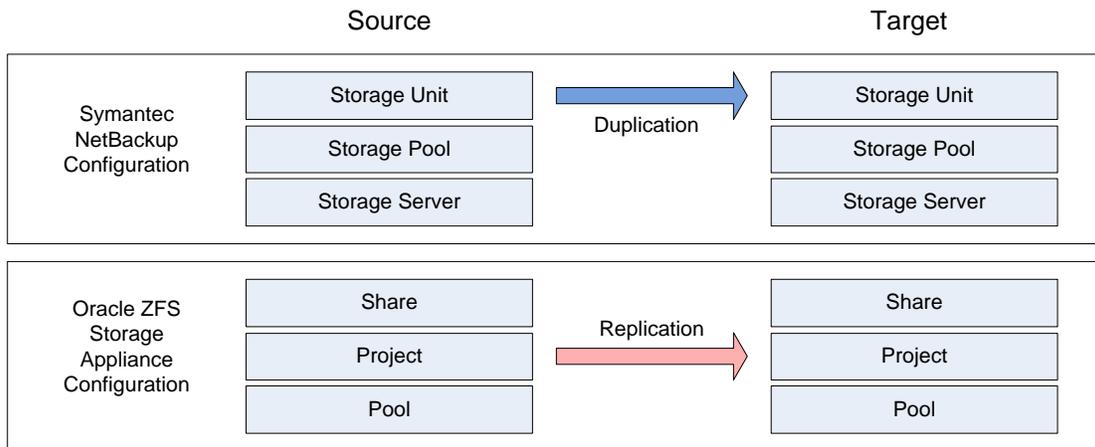


Figure 2. OST Optimized Duplication Solution Configuration Elements

The detailed installation and configuration steps are available in the *Oracle ZFS Storage Appliance Plug-in for Symantec NetBackup OpenStorage Administration Guide*. Both the plug-in and the documentation can be downloaded from the Oracle ZFS Storage Appliance plug-in downloads page listed in References at the end of this document.

## Remote Site Duplication Implementation

In its simplest configuration, Optimized Duplication can be utilized within the same data center. However, the solution can also be applied in disaster recovery scenarios between local buildings, between cities, or even across the country. Figure 3 presents a possible remote site data protection configuration.

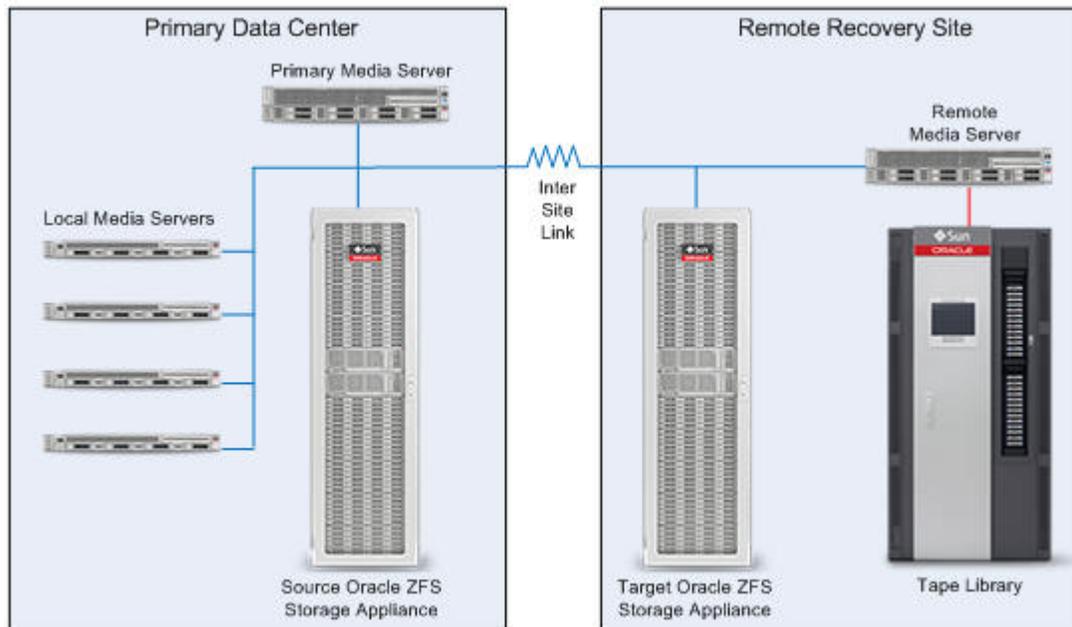


Figure 3. Example OST Optimized Duplication configuration to a remote location

In the Figure 3 example, multiple NetBackup media servers at the local site are able to store image backups to the Source storage unit. At least one media server, however, must have connectivity to both the Source and Target storage units in order to coordinate the Optimized Duplication copy operation. At the remote site, the media servers need only have access to the remote Oracle ZFS Storage Appliance in order to perform recovery operations or other data protection, such as backing up to tape.

## Implementation Best Practices

The following sections provide guidance for setting up the OST Optimized Duplication environment.

### Using a Common Media Server for Duplication

By default, NetBackup selects media servers based on preferences and availability. If a preferred server is not busy, it will tend to use (but not always) that media server. If the preferred server is busy, NetBackup will select another available server. That selection may cause a failure if the selected server does not have access to both the Source and Target storage units.

In order for the OST Optimized Duplication operation to succeed, the selected media server must have access to both storage units, which is set with the master server's `COMMON_SERVER_FOR_DUP` parameter. By default, this parameter is set to "default," which can result in the action just described that can produce a failure.

Change the `COMMON_SERVER_FOR_DUP` parameter setting to "required," which forces NetBackup to make sure that the selected media server has direct access to both storage units. If the media servers are busy, the duplication action will wait until the selected media server is free rather than attempt to locate another server.

The command to change the setting is:

```
nbbemcmd -changesetting -machinename <master server name>  
-COMMON_SERVER_FOR_DUP required
```

### Setting Up and Configuring the Storage

The architectural design and configuration of the storage in both the Oracle ZFS Storage Appliance and within NetBackup is critical to proper operation.

#### Setting Oracle ZFS Storage Appliance Pool Size

Since this solution uses Oracle ZFS Storage Appliance replication, the available space in the Target pool must be as large (if not larger) in size as the Source pool. Replication failures will occur if the Target pool runs out of space. Space utilization must be closely monitored to avoid this situation.

To ensure the Target pool has the space it requires, dedicate a fixed-size storage pool on both the Source and Target appliances. The correct size of these pools depends on the amount of backup data to be stored and its storage duration. Additional disks can be added to these pools in the future if more storage is needed. Dedicated pools, however, are not required.

### Using Oracle ZFS Storage Appliance Projects and Shares

Creating a single Oracle ZFS Storage Appliance project to hold all the shares that will be used for OST Optimized Duplication is recommended. The replication is configured at the project level. The setup in Figure 4 depicts a configuration using a single share, while Figure 5 shows a configuration using multiple shares.

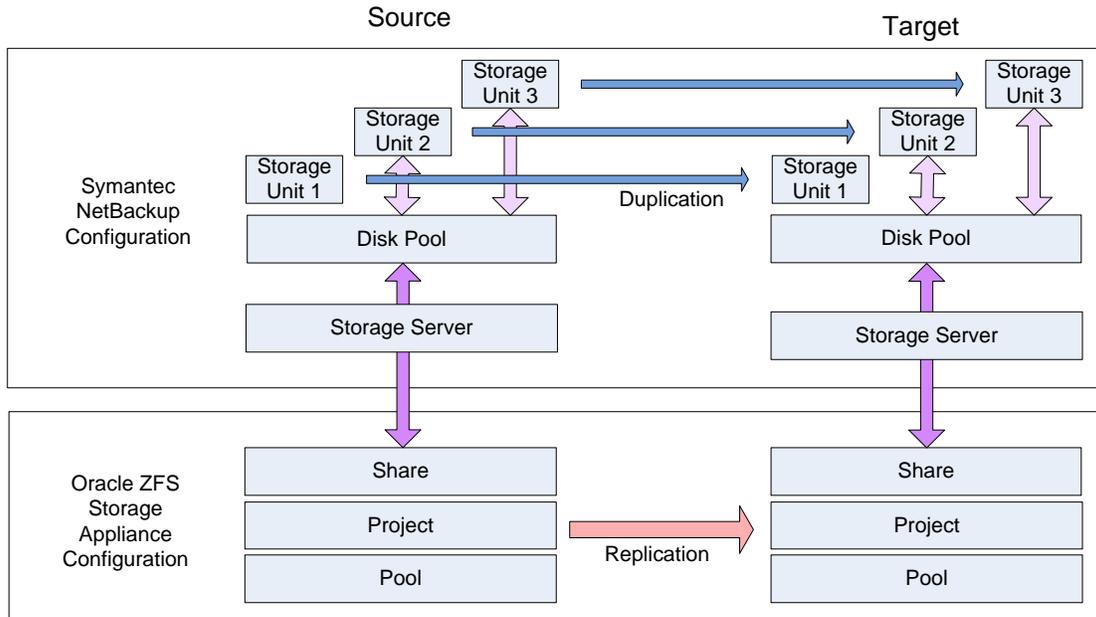


Figure 4. OST Optimized Duplication configuration using a single share

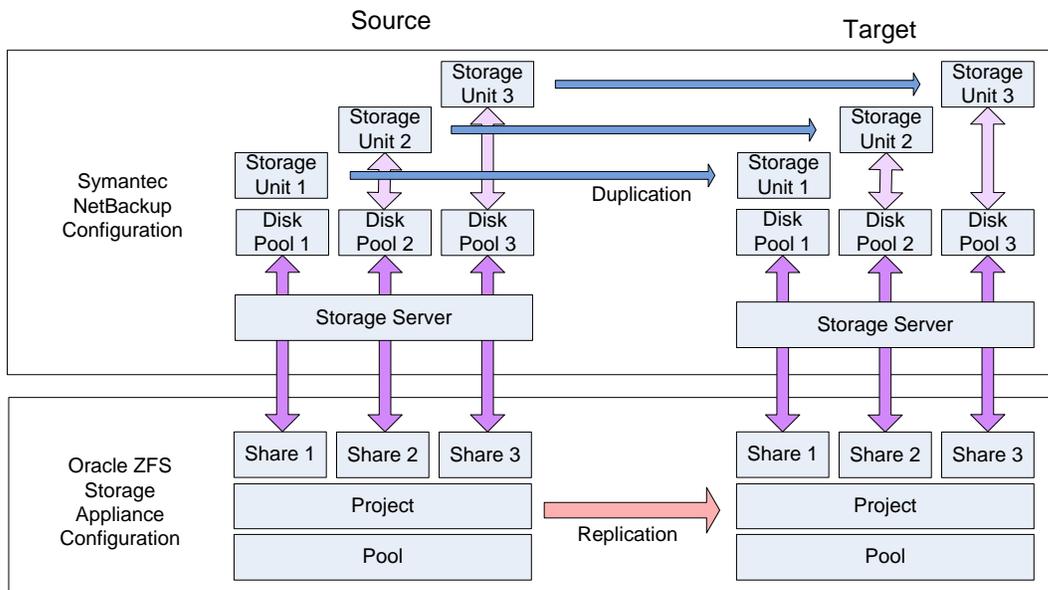


Figure 5. OST Optimized Duplication configuration using multiple shares

Each share in the configuration becomes its own disk pool in NetBackup. For most environments, the single share approach shown in Figure 4 would be adequate. The multishare environment in Figure 5 would only be necessary if a more compartmentalized design was required.

### Setting Up and Configuring NetBackup

Consider the following in order to properly set up the OST Optimized Duplication solution within NetBackup.

#### **NetBackup Storage Server Access**

Not all media servers need to access both storage servers. In the Figure 3 example, only the local media servers and the primary media server need access to the Source storage server (Oracle ZFS Storage Appliance). Only the primary media server and the remote media server need access to the Target storage server (appliance). In the presented solution, only the primary media server will have the ability to perform the OST Optimized Duplication between the storage units.

All of the local media servers could have access to both storage servers, and therefore be able to perform the duplication. The remote media server, however, should not have access to the Source storage server. Otherwise, the remote media server could be selected by NetBackup to perform a local backup. The local backup would back up to the media server, which would write the data back to the Source storage server, which would then get replicated to the Target storage server. Thus, the inter-site link would have the same data crossover three times. Depending on the bandwidth between the local and remote sites, this could greatly impact inter-site operations.

Each media server that will perform the OST Optimized Duplication action must have all of the shares configured on that host. For example, assume an environment where two media servers can be chosen by NetBackup for OST Optimized Duplication, and that two sets of disk pools will be used: d1\_source and d1\_target, d2\_source and d2\_target. All four of the underlying shares must either be mounted to both servers (UNIX media servers), configured to use SMB (Windows media servers), or both (mixed environments). An improperly configured setup will cause failures in duplication operations and some configuration tasks.

#### **NetBackup Disk Pool and Storage Unit Configuration**

Each NetBackup disk pool can only be configured with one Oracle Storage Appliance share, as is demonstrated in Figures 4 and 5. This configuration restriction is due to the one-to-one relationship required with the underlying replication. The NetBackup disk pool is then used to configure the storage unit(s).

At a functional level, there are no benefits to creating multiple storage unit source-target pairs. The more storage unit source-target pairs that are created, the more cluttered the setup, and administration becomes more complex.

NetBackup can only back up its images to the Source storage unit. Only a NetBackup duplication operation can copy an image from the Source storage unit to the Target storage unit. Any NetBackup policy that attempts to write a backup image directly to the Target storage unit will fail.

### Configuring a NetBackup Storage Lifecycle Policy (SLP)

SLPs can be used for the routine duplication of backup images between the two storage units. Within the SLP definition, two steps are needed to automate the duplication:

- Operation 1 – This can either be the initial backup or a subsequent duplication with a destination of the Source storage unit. An alternate read server does not need to be specified. Set the desired retention period.
- Operation 2 – This needs to be duplication with a destination of the Target storage unit. The alternate read server should be selected as a server that has access to both the Source and Target storage units. Set the desired retention period to the same as in Operation 1.

Due to the replication technology, the space that is used by the backup image will only be made available again when both the Source and the Target images have expired. Thus, there is no benefit to having a shorter retention period on the Source storage unit's data. It would actually be less beneficial to expire it sooner, especially when in a remote site configuration (Figure 3.) With a shorter retention at the Source storage unit, a restore would have to come across the inter-site link.

In the case of the configuration in Figure 3 where a backup is performed at a remote site, an additional SLP step would be:

- Operation 3 – This would be a duplication, as well, with a destination of the desired tape library. The alternate read server should be selected as the remote media server since it would have access to both the Target storage unit and the tape library. Set the desired tape retention period.

### Tuning the Operational Configuration

After a configuration is operational, it will take some amount of monitoring to make sure that the system is running smoothly. One potential area of concern is the size of the stored backup images and the amount of replication bandwidth between the Source and Target Oracle ZFS Storage systems. It is possible that a backup to the Source appliance can complete and a duplication operation started before the data replication to the Target appliance has completed. This can occur if the size of the backup image is considerably

large and the inter-site network bandwidth is relatively slow. The duplication operation will fail under these conditions.

To work around this situation, disable the SLP until a time that it is known that the replication would have completed. For example, a weekly full backup may take eight hours to complete. However, the replication may take up to 12 hours to complete. A script could be used in this case to deactivate the SLP just prior to the start of the backup job and activate it again 24 hours later to complete duplication.

Another solution might be to reduce the amount of data backed up through individual backup policies or, similarly, reduce the maximum fragment size of the Source storage unit. The goal in using these options is to create smaller “chunks” of data to replicate and duplicate. However, this is much less convenient.

### Supporting Mixed Windows and UNIX Environments

The Oracle ZFS Storage Appliance Plug-in for OST supports mixed platform environments to the storage units. For the interface to function properly, though, it is critical that the resource naming be identical between Windows and UNIX environments.

For example, the duplication shares may be presented to the Windows media server as:

```
\\zfssa1.domain.com\ost_duplicate
```

```
\\zfssa2.domain.com\osttarget_duplicate
```

On the UNIX media server, the shares must be mounted as:

```
/oraclezfs_ost/zfssa1.domain.com/ost_duplicate
```

```
/oraclezfs_ost/zfssa2.domain.com/osttarget_duplicate
```

The plug-in will fail with media access errors if the shares are not correctly set.

### Service Interruption Considerations

Understanding the operational behavior in the following service interruption scenarios should prompt some foresight in order to effectively manage them.

- **Temporary Replication Interruption** – This can occur if the network link between the two Oracle ZFS Storage Appliance products is interrupted. If the outage is short term, the replication should restart as soon as the network connectivity is restored. Any duplication operations that may start during that timeframe can fail since the replicated data will not be available on the Target Oracle ZFS Storage Appliance. If it will be an extended outage, then the associated SLP should be deactivated until the connectivity is restored and replication has completed.
- **Primary Site Failure** – In this event, the Target/remote Oracle ZFS Storage Appliance will continue operating under the parameters of the Target storage unit. Its contents will

remain read only to be used for restore, and the images can be duplicated/copied to a new location. Once the data has been copied off and is no longer needed, the storage can be reconfigured to use in a new OST Optimized Duplication configuration. The inherent Oracle ZFS Storage Appliance replication features cannot be used to reverse replication direction. There is no current method in the OST framework to register these changes.

## Conclusion

The Oracle ZFS Storage Appliance implementation of the Symantec NetBackup OST Optimized Duplication feature is a powerful tool for quickly copying backup images from one Oracle ZFS Storage Appliance to another. It offloads the action from the NetBackup media server and reduces network bandwidth utilization by half. By using the provided guidelines, the solution will run seamlessly in a simple or complex enterprise backup environment.

## References

See the following resources for additional information relating to the products covered in this document.

References to Sun ZFS Storage Appliance, Sun ZFS Storage 7000, and ZFS Storage Appliance all refer to the same family of Oracle ZFS Storage Appliance products. Some cited documentation may still carry these legacy naming conventions.

- Oracle ZFS Storage Appliance Documentation Library, including Installation, Analytics, Customer Service, and Administration guides:  
<http://www.oracle.com/technetwork/documentation/oracle-unified-ss-193371.html>
- The *Oracle ZFS Storage Appliance Administration Guide* is also available through the Oracle ZFS Storage Appliance help context.  
The Help function in Oracle ZFS Storage Appliance can be accessed through the browser user interface.
- Oracle Support Center  
<http://www.oracle.com/support>
- Patches and updates downloads from My Oracle Support (MOS)  
(search under Oracle ZFS Storage Software Patches)
- Oracle ZFS Storage Appliance Plug-ins  
<http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html>
- Oracle Storage Product Information  
<http://www.oracle.com/us/products/storage/overview/index.html>
- Oracle ZFS Storage Appliance Technical White Papers and Solution Briefs  
<http://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/index.html>
- Symantec NetBackup  
<http://www.symantec.com/netbackup>



Implementing Oracle ZFS Storage Appliance  
OST Optimized Duplication in a Symantec  
NetBackup Data Protection Environment  
September 2014, v1.1  
Application Integration Engineering  
Author: Eric Polednik

Oracle Corporation  
World Headquarters  
500 Oracle Parkway  
Redwood Shores, CA 94065  
U.S.A.

Worldwide Inquiries:  
Phone: +1.650.506.7000  
Fax: +1.650.506.7200

[oracle.com](http://oracle.com)



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2014, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0611

**Hardware and Software, Engineered to Work Together**