

Oracle Optimized Solution for Secure Disaster Recovery

Highest Application Availability with Oracle SuperCluster


ORACLE WHITE PAPER | OCTOBER 2015





Table of Contents

Executive Overview	1
Introduction	2
Disaster Recovery Planning	2
Creating More-Secure Disaster Recovery Environments	3
Disaster Recovery for Oracle SuperCluster	3
Oracle SuperCluster Overview	3
Disaster Recovery Strategy for Oracle SuperCluster	4
Example Solution for Oracle SuperCluster Disaster Recovery	5
Disaster Recovery for Applications	5
Oracle ZFS Storage Appliance	5
Remote Replication	6
Pools, Projects, and Shares	7
ZFS Replication Modes	7
Disaster Recovery Guidelines	7
Other Replication Tools	9
Disaster Recovery for Databases	9
Data Guard and Oracle Active Data Guard	10
Data Guard Implementation Overview	10
Data Guard Best Practices	12
Oracle GoldenGate	13
Oracle GoldenGate Architecture	13



Deploying Oracle GoldenGate 11g for Disaster Recovery	14
Database Migration	14
Non-Oracle Database Replication Tools	15
Database Recommended Use Cases	15
Complementary Technologies	16
Oracle Recovery Manager	16
Zero Data Loss Recovery Appliance	16
Oracle Solaris Cluster	17
Oracle Solaris Cluster Geographic Edition	17
Oracle Clusterware	18
Integrated System Monitoring	18
Private and Hybrid Cloud Configurations	19
Best Practices for a Secure Disaster Recovery Implementation	19
Security Technical Implementation Guides	19
Component-Level Security Recommendations	20
Example Best-Practices Implementation	22
Implementation Overview	22
Disaster Recovery Setup	22
Disaster Recovery Testing	24
Summary	25
References	26



Executive Overview

Oracle SuperCluster is Oracle's most powerful and scalable engineered system. This pretested and optimized system includes integrated server, storage, networking, and software. Components include Oracle's powerful SPARC servers; Oracle Exadata Storage Servers and Oracle ZFS Storage Appliance; high-speed, low-latency InfiniBand fabric; and the Oracle Solaris operating system with built-in virtualization. The Oracle SuperCluster platform is designed from the ground up for high availability. Hardware components have no single point of failure, and there is end-to-end software high availability. However, in addition to built-in high availability, enterprise deployments need further disaster recovery strategies for protection from unforeseen disasters and natural calamities.

The typical disaster recovery solution involves setting up a standby site at a geographically different location from the production site. All data—application data, configuration data, metadata, and all database information—are replicated to the standby site on a periodic or continual basis. In the event of a disaster, activity can transfer to the standby site for continued operation.

Oracle Optimized Solution for Secure Disaster Recovery uses components from Oracle's end-to-end hardware and software technology stack—including Oracle ZFS Storage Appliance, Oracle's Zero Data Loss Recovery Appliance, Oracle GoldenGate and Data Guard—to provide next-generation data protection. This solution uses the replication technology of Oracle ZFS Storage Appliance for protection of middle-tier applications and components running on the cluster. Additionally, Data Guard or Oracle GoldenGate are used to provide disaster recovery for databases that are part of Oracle SuperCluster deployments. Oracle Solaris Cluster Geographic Edition and Oracle Enterprise Manager provide management of the entire disaster recovery solution. Third-party replication tools are also supported, if necessary, to provide integration of legacy and non-Oracle databases¹ and applications.

This technical paper provides an overview of the best practices and implementation strategies for advanced, efficient disaster recovery on Oracle SuperCluster.

¹ In this paper, the term *non-Oracle databases* refers to all databases other than Oracle Database.



Introduction

Disaster recovery planning requires careful consideration, with special attention given to the Recovery Point Objective (RPO) and Recovery Time Objective (RTO) of business applications. The RTO determines how quickly applications and databases must be made available after a failure occurs; the RPO is the maximum amount of time for which data might be lost if a major incident occurs. For example, an organization might determine that it is acceptable to lose one hour of data (RPO) and that application services must be back online within two hours (RTO). Data synchronization points must also be specified, to enable data backups to correctly correlate to each other and achieve a fully consistent recovery point.

Oracle provides flexible disaster recovery alternatives that consider RPO and RTO as well as varying software releases, the nature of data (either structured or unstructured), and other special customer needs.

Disaster Recovery Planning


Planning for disaster recovery starts with determining the acceptable RTO and RPO for the applications and services provided by a given IT deployment. Determining the RTO and RPO depends on many factors, including cost, existing infrastructure capabilities, requirements for compliance with government regulation, and other business objectives. The number of standby sites, the physical distance between sites, and the need for synchronous or asynchronous communications are all important considerations in disaster recovery planning.

A complete discussion of disaster recovery planning and objectives is beyond the scope of this paper. However, the establishment of a standby site (or multiple standby sites) at a location that is geographically distant from the production site is common to virtually all disaster recovery solutions. Natural catastrophes such as fire and flood, and other disasters such as sabotage or human error, can render an entire data center site unusable. In these situations, operations can be configured to continue at a geographically distant site that is unaffected by the event.

The remote standby site hosts a redundant application tier and a synchronized standby database. The standby site might be symmetric, with an equal number of services and resources compared to the production site. Alternatively, an asymmetric standby site, with fewer services and resources, can be configured. All data from the primary production site—including application and database data, as well as configuration data and metadata—is replicated to the standby site. This replication can be scheduled on a periodic or continual basis, depending on business requirements. In the event of catastrophic failure of the primary site, operation can be quickly switched over to the backup site.

The standby site can be configured in a passive mode; it is started when the primary site becomes unavailable. This deployment model is referred to as an *active-passive model*. It is also possible to configure the standby site for operations such as reporting, testing, or other business functions. This deployment model, referred to as an *active-active model*, eliminates idle redundancy and provides better utilization of cluster resources. The choice of active-passive or active-active configuration depends on business requirements. Some deployments might require active-passive solutions for compliance or other business reasons. Many organizations choose an active-active configuration to better utilize standby cluster resources and achieve a higher return on investment of resources.

Planning for the standby site should take into consideration the processing and storage requirements for both the primary business operations and any other functions that are planned while the site is in standby mode. If a standby site is actively used for functions such as development, testing, or reporting, these activities might need to be temporarily reduced or suspended during emergency failover when primary business operations are transitioned to



the standby site, depending on available server resources. Consolidation of multiple environments might require additional storage to ensure the standby site can be fully utilized during a failover event. Additional Exadata Storage Expansion Racks from Oracle can be configured at the standby site to facilitate any extra capacity that is required.

Creating More-Secure Disaster Recovery Environments

The following steps can help create a more-secure disaster recovery environment:

- » **Simplify the infrastructure.** Most disaster recovery environments are based on a complex infrastructure, making implementation and management complicated. This complexity increases the risk of security vulnerabilities. A disaster recovery implementation as a whole is only as secure as its most vulnerable component, and it can be challenging to securely configure the myriad interacting components and products in a heterogeneous system. Oracle Optimized Solutions simplify disaster recovery implementations through the use of consolidation and virtualization technologies. Oracle also offers security guidelines and recommendations, and many Oracle components have security built-in by default.
- » **Reduce implementation flaws.** Secure software is important but not sufficient by itself. Most security vulnerabilities arise from flawed implementation and architecture, including improper configuration and access control, lack of patch management, unencrypted communications, and inadequate security policies and processes. Based on current security best practices, Oracle Optimized Solutions provide proven and tested architecture recommendations for increased disaster recovery solution protection.
- » **Eliminate performance and cost penalties.** Many security processes, such as on-the-fly encryption/decryption, can have a significant negative impact on the performance and cost of a disaster recovery solution. Oracle Optimized Solutions leverage Oracle's SPARC-based systems, which offer high-performance security using cryptographic instruction accelerators that are directly integrated into the processor cores. By providing wire-speed security capabilities, Oracle systems eliminate the performance and cost penalties typically associated with real-time, secure computing.

Disaster Recovery for Oracle SuperCluster

The following sections provide an overview of Oracle SuperCluster and the recommended disaster recovery strategy to provide the highest application availability.

Oracle SuperCluster Overview

Oracle SuperCluster is a multipurpose engineered system that has been designed, tested, and integrated to run mission-critical enterprise applications and rapidly deploy cloud services while delivering extreme efficiency, cost savings, and performance. Preconfigured with Oracle's SPARC servers, Oracle Exadata Storage Servers, Oracle ZFS Storage Appliance, InfiniBand technology, and Oracle Solaris, the Oracle SuperCluster platform is delivered fully tested and ready to deploy. This system is well suited for multitier enterprise applications with web, database, and application components. Oracle SuperCluster is designed to host the entire Oracle software solution stack, as well as third-party applications and customer-developed software, all within a single rack enclosure.

Oracle SuperCluster combines highly available and scalable technologies with industry-standard hardware, and it is architected from the ground up with end-to-end high availability. Oracle Real Application Clusters (Oracle RAC), a feature of Oracle Database 11g, and Oracle Clusterware provide high availability and failover capabilities for the database. Oracle Solaris Cluster provides high availability for applications. Key hardware components, including the SPARC servers, Oracle ZFS Storage Appliance, and Oracle Exadata Storage Servers, are configured with no single point of failure. Data availability is delivered through features such as memory mirroring and extended-ECC memory, as well as through the data protection features of Oracle ZFS Storage Appliance (such as dual controllers and ECC memory) and the Oracle Solaris ZFS file system (such as self-healing, triple-RAID parity, and triple-mirroring).

Applications that run on the SPARC servers run in either a Database Domain or an Application Domain. A Database Domain is dedicated to running Oracle Database 11g Release 2 (or later) using Oracle Exadata Storage Servers for database storage.

Disaster Recovery Strategy for Oracle SuperCluster

Although the Oracle SuperCluster platform is designed for high availability, enterprise deployments need protection from unforeseen disasters and natural calamities. Oracle SuperCluster uses a combination of technologies to provide disaster recovery support for applications and databases deployed on this platform (see Figure 1). Oracle Active Data Guard and Oracle GoldenGate replication are the best-practices recommendation for database content; ZFS replication is recommended for applications and unstructured data; and Oracle Solaris Cluster Geographic Edition and Oracle Enterprise Manager are recommended for the management of the entire disaster recovery solution.

In the Oracle SuperCluster platform, applications and unstructured data (that is, non-database data) reside in shared file systems on Oracle ZFS Storage Appliance. This data can include applications such as Oracle enterprise applications, third-party applications, and any custom applications running on the Oracle SuperCluster platform. Disaster recovery strategies for this data utilize the remote replication features of Oracle ZFS Storage Appliance. By maintaining a replica of the primary data at a remote site, disaster recovery time is dramatically reduced compared to traditional offline backup architectures. The Oracle ZFS Storage Appliance contained in Oracle SuperCluster includes a 1 Gigabit Ethernet (GbE) port that is reserved for replication purposes; no additional hardware is required.

Replication and cloning are separately licensed features, but these licenses are included with the Oracle ZFS Storage Appliance that is internal to Oracle SuperCluster. However, licenses must be purchased for each external Oracle ZFS Storage Appliances at the local and remote sites.

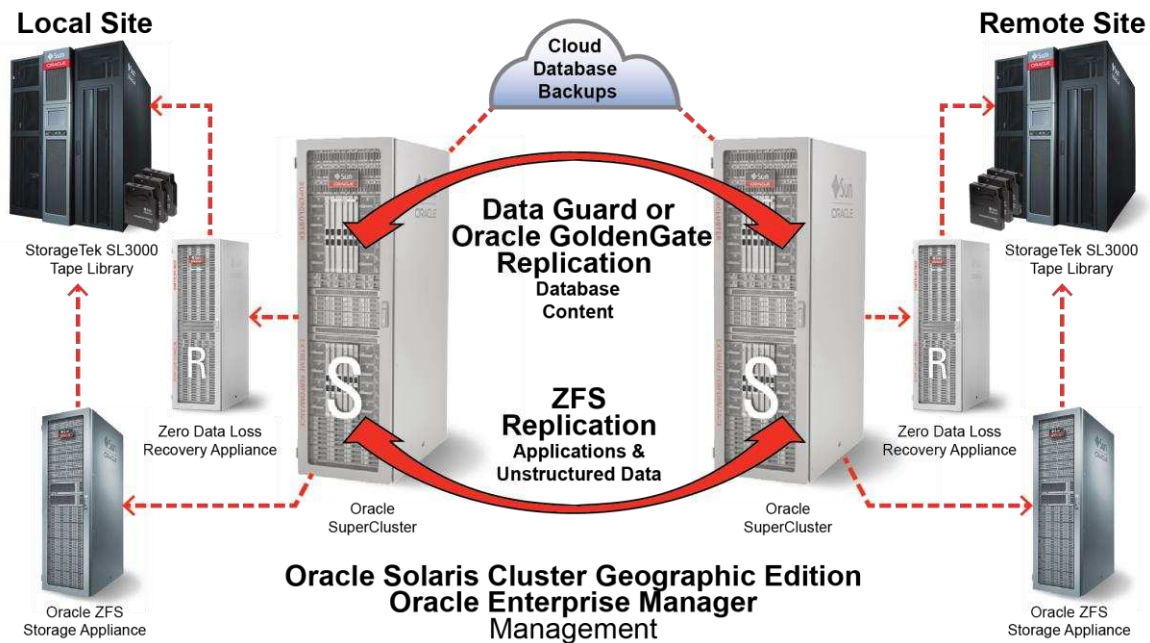



Figure 1. Oracle SuperCluster can be deployed as a part of an effective disaster recovery topology.



Oracle SuperCluster provides a choice of databases: all databases that run on Oracle's SPARC servers and Oracle Solaris 11, including Oracle Database and non-Oracle database solutions, are supported. Oracle Database 11g Release 2 (or later) instances run in the Database Domain of Oracle SuperCluster and have access to Oracle Exadata Storage Servers for database storage. Other earlier Oracle Database versions and non-Oracle databases run in the Application Domain of Oracle SuperCluster and do not have access to Oracle Exadata Storage Servers. These databases can use Oracle ZFS Storage Appliance or external Fibre Channel SAN storage as the repository for data.

Data Guard and Oracle Active Data Guard are the standard recommendations for disaster recovery for Oracle Database. Oracle GoldenGate should be used for disaster recovery for non-Oracle database environments, for replication across heterogeneous Oracle Database releases, or for bidirectional replication when both replicas must be open in read-write mode at the same time. Databases using Oracle ZFS Storage Appliance can also use ZFS replication. In addition, non-Oracle database replication tools are also supported for legacy implementations (but are not recommended as a best-practices solution).

Database replication typically uses the same 10 GbE ports in Oracle SuperCluster that are used for users and applications. However, separate ports can be configured, if necessary, to meet performance or other business requirements.

In the event of complete site failures, additional wide area network (WAN) hardware might be needed to provide business continuity. To maintain availability, users must be redirected to the standby site. A WAN traffic manager can be used to execute a Domain Name Server (DNS) failover—either manually or automatically—to redirect users to the application tier at the standby site while a database failover transitions the standby database to the primary production role. Please see Oracle Database high availability best practices documentation for information on automating complete site failover.

Example Solution for Oracle SuperCluster Disaster Recovery

An example disaster recovery implementation for Oracle E-Business Suite on Oracle SuperCluster is described in My Oracle Support Note 1558827.1, *Oracle E-Business Suite R12.1.3 Disaster Recovery: Implementation Guide on Oracle SuperCluster*. Although this implementation guide is specific to Oracle E-Business Suite, the general principles and methodology are applicable to other applications that run on Oracle SuperCluster.


The example implementation uses the recommended best-practices combination of Oracle Active Data Guard (for the database), Oracle ZFS Storage Appliance replication (for applications), and Oracle Solaris Cluster (for management). This example implementation for disaster recovery is summarized later in this paper in the section "Example Best-Practices Implementation" on page 22.

Disaster Recovery for Applications

Applications and unstructured data (non-database data) reside in shared file systems on Oracle ZFS Storage Appliance. The recommended best practice is to use the remote replication features of Oracle ZFS Storage Appliance to maintain a copy of this data at a remote site. Oracle SuperCluster also integrates with other non-Oracle replication tools, providing legacy support for existing deployments that use other replication solutions.

Oracle ZFS Storage Appliance

The Oracle SuperCluster platform is preconfigured with a dual-controller Oracle ZFS Storage Appliance. These systems feature a common, easy-to-use management interface and have the industry's most comprehensive analytics. Oracle ZFS Storage Appliance includes an intelligent hybrid storage pool designed to automatically



optimize performance. The storage utilization suite features data deduplication and compression to improve storage efficiency. Oracle ZFS Storage Appliance also provides replication, making this storage system an outstanding target for Oracle SuperCluster backup/recovery and disaster recovery strategies.

In addition to providing disaster recovery protection, Oracle ZFS Storage Appliance can make the disaster recovery site more productive—actually contributing to an organization's productivity, instead of just sitting there waiting for a disaster. Oracle ZFS Storage Appliance's snapshot and clone features at the disaster recovery site can create database instances that can be used for test, development, and reporting functions. Using the disaster recovery site for these functions offloads the main production site so it can focus exclusively on transaction processing, improving service levels to the business.

Remote Replication


Oracle ZFS Storage Appliance supports snapshot-based replication of projects and shares from a source appliance to any number of target appliances. A snapshot is a view of a file system at a particular point in time, including both data and metadata. Replication can be performed manually, on a schedule, or continuously, depending on business requirements.

In a disaster recovery strategy, replication can be used to mirror Oracle ZFS Storage Appliances. In the event of a disaster that impacts service of the primary appliance, administrators activate service at the remote disaster recovery site. The remote site then takes over operation using the most recently replicated data. When the primary site has been restored, data that changed while the disaster recovery site was in service can be migrated back to the primary site and normal service can be restored. Such scenarios are fully testable before a disaster occurs.

The remote replication feature in Oracle ZFS Storage Appliance has several important properties:

- » **Snapshot-based.** The replication subsystem takes a snapshot as part of each update operation. For full updates, the entire project contents are sent to the snapshot. For incremental updates, only the changes since the last replication snapshot for the same action are sent.
- » **Block-level.** Each update operation traverses the file system at the block level and sends the appropriate file system data and metadata to the target.
- » **Asynchronous.** Because the replication function takes snapshots and then sends them, data is necessarily committed to stable storage before replication begins sending the snapshots. Continuous replication effectively sends continuous streams of file system changes, but this process is still synchronous with respect to NAS and SAN clients.
- » **Includes metadata.** The underlying replication stream serializes both user data and Oracle Solaris ZFS metadata, including most properties configured on the Shares screen. These properties can be modified on the target after the first replication update is complete (though not all changes take effect until the replication connection is severed). For example, this capability to modify properties allows sharing over NFS to a set of hosts that is different from that on the source.
- » **Secure.** The replication control protocol used among Oracle ZFS Storage Appliances is secured with SSL. Data can optionally be protected with SSL as well. For additional security, appliances can replicate only to/from other appliances after an initial manual authentication process.
- » **Protocol independent.** Oracle ZFS Storage Appliance supports both file-based (CIFS and NFS) and block-based (Fibre Channel, iSCSI, and iSER) storage volumes. The replication mechanism is protocol independent.
- » **Includes cloning and replication licenses.** Oracle SuperCluster includes licensing for ZFS cloning and replication, eliminating any additional expense for these capabilities.

Because replication is asynchronous, the system does not have to wait for data to be saved at the replication site. This approach has the advantage of processing writes much faster at the primary location. In addition, this technique allows for replication over much longer distances. The link between the storage systems can have a lower



bandwidth (not every write has to be replicated, only the state of the system at certain points in time) and higher latency (because writes don't need to be confirmed at both sites at once). The obvious disadvantage is that in case of a failure on the primary system, data loss is probable (and almost guaranteed) because of the small delay in storing replication data to the replication site. The secondary system will always be missing any data that has been written to the master but not yet stored in the replica. Performance is greatly increased; but if local storage is lost, the remote storage is not guaranteed to have a current copy of the data and the most recent data might be lost.

Pools, Projects, and Shares

Oracle ZFS Storage Appliance uses storage pools and projects to organize data.

- » The storage pool (similar to a volume group) is created over a set of physical disks. File systems are then created over the storage pool. On the Oracle SuperCluster platform, the storage pool is configured with a mirrored disk layout by default. It is recommended to use the mirrored disk layout for increased fault tolerance and improved read performance.
- » All file systems and LUNs are grouped into projects. A project can be considered a *consistency group*. A project defines a common administrative control point for managing shares. All shares within a project can share common settings, and quotas can be enforced at the project level in addition to the share level. Projects can also be used solely for grouping logically related shares together, so their common attributes (such as accumulated space) can be accessed from a single point.
- » Shares are file systems and LUNs that are exported over supported data protocols to clients of the appliance. Exported file systems can be accessed over CIFS, NFS, HTTP/WebDav, and FTP. LUNs export block-based volumes and can be accessed over iSCSI, Fibre Channel, and iSER. The project/share is a unique identifier for a share within a pool. Multiple projects can contain shares with the same name, but a single project cannot contain shares with the same name.

ZFS Replication Modes

Oracle ZFS Storage Appliance remote replication supports three different modes: on-demand, scheduled, and continuous. The replication process is the same for each mode; the only difference is the time interval between replications. The replication mode can be changed at any time to support different and changing business requirements.

- » **On-demand.** Replication is triggered manually by the user at any time.
- » **Scheduled.** Replication is automatically executed according to a predetermined schedule. Schedules can be defined at the granularity of half-hourly, hourly, daily, weekly, and monthly.
- » **Continuous.** The replication process is automatically executed continuously. As soon as one replication update is complete, a subsequent update is started. This way, the changes are transmitted as soon as possible.

Disaster Recovery Guidelines

The following disaster recovery guidelines are recommended for application (that is, non-database) data on the Oracle SuperCluster platform.

Replication Mode Guidelines

Business processes, such as RTO, RPO, and service-level agreement (SLA), should be considered in deciding the mode of replication. The rate of change, latency, bandwidth, and number of projects to replicate all influence the decision-making process. Table 1 lists the replication mode options for Oracle ZFS Storage Appliance.

TABLE 1. REPLICATION MODE GUIDELINES

Mode	Requirements	Comment
Continuous	Near-real-time protection (RPO/RTO < few minutes) is needed.	Updates are sent as fast as network bandwidth permits.
Scheduled	A longer RPO/RTO is permitted or there is insufficient bandwidth for continuous replication.	This mode reduces network traffic while preserving consistent and timely copies of the primary data set.
On-Demand	Data needs to be in a specific state before replication can occur.	This mode can use automated scripting to trigger on-demand replication.

The following list provides more details on each mode:

» **Continuous replication** of the project is an appropriate choice for technical and operational requirements that require near-real-time protection of data at the remote site, such as when there is an RPO and RTO of less than a few minutes. Updates to the source data set will be sent to the target site as fast as the network permits in this case.

Oracle ZFS Storage Appliance systems use asynchronous communication between the source and the target to ensure that network latency does not slow production operations. This technology cannot guarantee that updates to the source will be present at the target site after a loss of the source site; however, the image of the project at the target site is guaranteed to be write-order consistent as of the time of the most recently completed data transfer.

» **Scheduled replication** provides a good alternative to make the best use of available resources when available replication network bandwidth is insufficient for continuous replication, or when technical and operational requirements allow for a longer RPO and RTO. With scheduled replication, Oracle ZFS Storage Appliance periodically replicates a point-in-time image (snapshot) of the source project to the remote site. This reduces network traffic while preserving consistent and timely copies of the primary data set.


» **On-demand replication** is designed for applications that need to put data into a specific state before the replication can occur. For example, a replica of a cold or suspended database can be produced every time the database is shut down by integrating a call to trigger an on-demand replication update in the database shutdown or suspend scripts. On-demand replication updates can be triggered from arbitrary locations in the application-processing stack through the automated scripting language of the Oracle ZFS Storage Appliance command-line interface.

Project-Level Replication versus Share-Level Replication

Oracle ZFS Storage Appliance enables remote replication to be configured on both the project and share level. By default, the shares in a project inherit the configuration of the parent project. Inheriting the configuration not only means that the share is replicated to the same target on the same schedule with the same options as its parent project, but also that the share is replicated in the same stream using the same project-level snapshots as other shares inheriting the project's configuration. This capability is important for applications that require data consistency among multiple shares.

Overriding the configuration means that a share is not replicated with any project-level actions, though it may be replicated with its own share-level actions that include the project. It is not possible to override part of the project's replication configuration and inherit the rest.

More precisely, the replication configuration of a project and its shares defines some number of replication groups, each of which is replicated with a single stream using snapshots taken simultaneously. All groups contain the project itself (which essentially just includes its properties). One project-level group includes all shares inheriting the



replication configuration of the parent project. Any shares that override the project's configuration form a new group consisting of only the project and shares themselves.

Oracle strongly recommends that project-level and share-level replication be avoided within the same project, because it can lead to surprising results (particularly when reversing the direction of replication). In addition, project-level replication is required if Oracle Solaris Cluster Geographic Edition is used in the configuration.

Other Remote Replication Considerations

- » Synchronous mode is not supported, so a zero data loss (ZDL) requirement cannot be met. However, the continuous replication mode can provide an alternative that offers minimal data loss in the event of a disaster.
- » The write ordering and write consistency are maintained at the granularity of the replicated component. The write ordering is preserved within the share if the replication is set at the share level. However, the write ordering is not preserved across the shares if more than one share is replicated. The write ordering at the target for all the shares in the project is preserved if the replication happens at the project level. The write ordering is not preserved across the projects. Refer to the administration guide at oracle.com/technetwork/documentation/oracle-unified-ss-193371.html for details.
- » The target site should be configured with sufficient storage capacity. Before initiating replication, the target site should be verified to make sure that it has enough storage capacity to store the replica. (The target site is not automatically verified for the space requirement when the replication is established.)

Other Replication Tools

Although using the replication capabilities of Oracle ZFS Storage Appliance is generally recommended as a best practice for protecting unstructured application data in Oracle SuperCluster, other replication tools are also supported. Examples include products such as Oracle's Pillar Axiom system and replication tools such as Hitachi's Replication Manager and EMC Replication Manager. These tools can provide replication of data at a remote location for backup and disaster recovery purposes.


Support for these tools enables existing deployments that use these products for backup of applications, zones, or other unstructured data to run unchanged on Oracle SuperCluster. In addition to providing support for legacy deployments, support for these products provides a migration path to a future disaster recovery solution featuring Oracle ZFS Storage Appliance replication, if desired.

A full discussion of other replication tools is beyond the scope of this paper. Please contact your Oracle representative for more information.

Disaster Recovery for Databases

Multiple releases of Oracle Database and non-Oracle databases are supported for deployment on the Oracle SuperCluster platform. Oracle Database 11g Release 2 (or later) instances use Oracle Exadata Storage Servers for database storage. Earlier Oracle Database versions and other legacy non-Oracle databases do not have access to Oracle Exadata Storage Servers. Instead, these databases run in an Application Domain in Oracle SuperCluster and use the shared storage on Oracle ZFS Storage Appliance or other external storage hardware.

Data Guard (a feature included in Oracle Database, Enterprise Edition) is the general recommendation for disaster recovery for all Oracle Database versions prior to Oracle Database 11g; Oracle Active Data Guard is the general recommendation for Oracle Database 11g onward. Oracle GoldenGate is recommended for disaster recovery for all non-Oracle databases, or for heterogeneous Oracle Database configurations where primary and replica databases operate at different releases or run on different hardware architectures. Oracle GoldenGate is also recommended for configurations with bidirectional replication when both database replicas are simultaneously open in read-write mode.



Oracle SuperCluster also integrates with non-Oracle database replication tools, providing legacy support for existing deployments that use other hardware or software replication solutions.

Data Guard and Oracle Active Data Guard

Data Guard is included with Oracle Database, Enterprise Edition 11g and is used to maintain availability in the event unexpected outages impact the production database. Data Guard provides the management, monitoring, and automation software to create and maintain one or more synchronized copies (*standby databases*) of a production database (*primary database*). These standby databases protect the primary database in the event of failures, corruption, errors, and disasters, and can be used to minimize downtime during planned maintenance.

Data Guard's native integration with Oracle Database enables the highest level of data protection and performance. Corruption detection ensures that data is logically and physically consistent before it is applied to a standby database: Data Guard automatically repairs physical block corruption detected at either the primary or standby database using a good copy of the block retrieved from the other database. Data Guard is a lightweight, network-efficient, Oracle Database-aware process that transmits a database redo (a small fraction of the total write volume at a production database) directly from the memory of the primary database to all remote standby databases. Another Data Guard process running on the standby site receives the redo, validates that there is no corruption, and applies the changes to the standby database. In this manner Data Guard enforces strong isolation between the primary copy and the disaster recovery copy while providing the fastest, most reliable replication possible.

Data Guard supports both synchronous (zero data loss) and asynchronous (near-zero data loss) configurations. Administrators can use either manual or automatic failover to quickly transition a standby database to the production role.

Oracle Active Data Guard extends basic Data Guard functionality by allowing read-only access to a synchronized replica (physical standby) database. Changes transmitted from the primary database are continuously applied, while read-only access to the standby is permitted. Using the standby database to offload queries, provide reporting, or perform backups while also providing disaster recovery protection puts otherwise idle resources to work—increasing performance and providing an increased return on investment.

Data Guard Implementation Overview

Data Guard creates and maintains one or more standby databases. A standby database is initially created from a backup copy of the primary database. As users commit transactions to the primary database, Oracle Database generates redo records and writes them to a local online log file. Simultaneously, Data Guard transport services automatically transmit redo records directly from the primary log buffer to the standby databases(s), where the information is written to a standby redo log file (SRL).

Figure 2 shows a high-level overview of the Data Guard implementation. Redo data is transmitted, either synchronously or asynchronously, from the primary database to the remote replica as it is generated (1). At the remote replica, this redo data is used to update the standby database files (2). The primary database process updates the primary database files, independently of Data Guard (3). Data Guard provides automatic outage resolution (4), resynchronizing the standby database after any outages of the network or the standby database. Redo information archived at the primary database is used for this resynchronization.

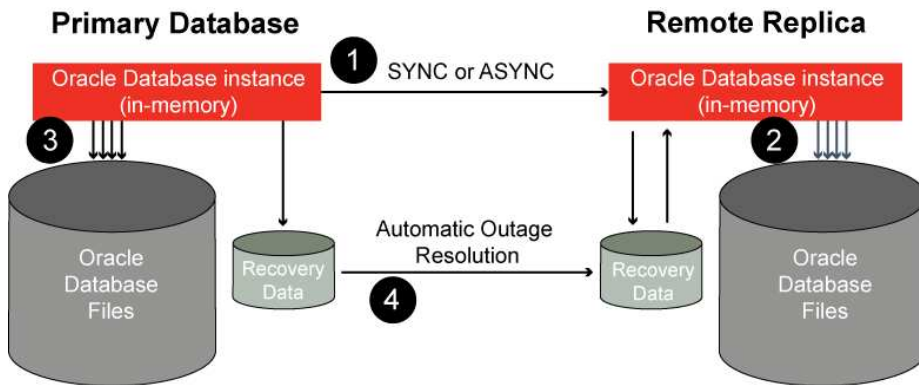


Figure 2. Data Guard provides remote database replication.

Redo Apply Feature of Data Guard and Oracle Active Data Guard

Data Guard and Oracle Active Data Guard use *Redo Apply* to maintain a synchronized copy of the production database.

A physical standby database is a physically identical copy of the primary database, with on-disk database structures that are identical to the primary database on a block-for-block basis. The database schema, including indexes, are the same. A physical standby database is kept synchronized with the primary database through Redo Apply, which uses media recovery to apply changes to a standby database that is open in read-only mode (Oracle Active Data Guard). Redo Apply maintains a block-for-block, exact replica of the primary database, ensuring that data is protected at all times.

Asynchronous Versus Synchronous Redo Transport

Redo information can be transmitted either synchronously or asynchronously.

- » **Synchronous redo transport (SYNC)** requires that the primary database wait for confirmation from a standby database that a redo has been received and written to disk (a standby redo log file) before it will acknowledge commit success to the application. This provides a guarantee of zero data loss in the event of any single failure, up to and including a complete site failure.
- » **Asynchronous redo transport (ASYNC)** avoids any impact to primary database performance by having the primary database acknowledge commit success to the application without waiting for acknowledgment that a redo has been received by the standby database. The performance benefit of ASYNC, however, is accompanied by the potential exposure for a small amount of data loss, because there can be no guarantee that at any moment in time all of a redo for committed transactions has been received by the standby database.

Data Protection Modes

Data Guard provides three modes of data protection: maximum protection, maximum availability, and maximum performance (see Table 2). Each mode uses a specific redo transport method and has rules that govern behavior if the primary database loses contact with the standby.

TABLE 2. DATA GUARD PROTECTION MODES

Mode	Data Loss	Transport	If No Acknowledgement from the Standby Database, Do This:
Maximum Protection	Zero data loss, double failure protection*	SYNC	The primary signals commit success to the application only after acknowledgement is received from a standby database that a redo for that transaction has been hardened to disk. *Double failure protection exists if multiple standbys are configured.
Maximum Availability	Zero data loss, single failure protection	SYNC	The primary signals commit success to the application only after acknowledgement is received from a standby database or after NET_TIMEOUT threshold period expires—whichever occurs first.
Maximum Performance	Potential for minimal data loss	ASync	The primary never waits for standby acknowledgment to signal commit success to the application.

Maximum protection mode ensures no data loss will occur if the primary database fails. To ensure that data loss cannot occur, the primary database will shut down (rather than continue processing transactions) if it cannot write its redo stream to at least one synchronized standby database. Maximum availability mode provides the highest level of protection possible without compromising availability: this mode ensures that no data loss occurs if the primary database fails, but only if there is not a second failure. The default protection mode is maximum performance. This protection mode offers slightly less data protection than maximum availability mode and has a minimal impact on primary database performance. (See [Oracle Data Guard Concepts and Administration 11g Release 2](#) for more details.)

Role Management Services

Data Guard role management services enable the primary and standby roles of the databases in a Data Guard configuration to be changed. For disaster recovery purposes, a failover transition can be initiated in response to a failure of the primary database. In this event, the standby database is transitioned to the primary role. The original primary database is then removed from the Data Guard configuration.

Data Guard Best Practices

The following Data Guard best practices are recommended for Oracle SuperCluster disaster recovery.

Redo Apply

Physical standby databases provide the best disaster recovery protection for Oracle Database. Therefore, configuration of a physical standby using the Redo Apply synchronization method is recommended as a best practice for Oracle SuperCluster disaster recovery.

Redo Apply is the simplest, fastest, and most reliable method of maintaining an independent, synchronized replica of a primary database. A physical standby database applies the redo received from its primary database using the managed recovery process (MRP), an extension of standard media recovery that is used by every Oracle Database instance. The MRP controls the highly parallel recovery mechanism native in the Oracle Solaris kernel.

A number of Redo Apply performance enhancements have been implemented to take specific advantage of the superior I/O characteristics of Oracle Exadata Storage Servers. In general, Redo Apply performance should be sufficient for most workloads using default settings. If, however, the standby database is unable to keep pace with the rate of primary database redo generation, see the [best practices for Oracle Maximum Availability Architecture](#) for tuning media recovery.

Redo Transport and Protection Mode

Data Guard synchronous redo transport with maximum availability mode is recommended for applications that have a zero-data-loss RPO. Maximum availability with SYNC is always recommended for ideal data protection if the round-trip time (RTT) between the primary and standby databases is less than 5 milliseconds. Higher RTT latency might still be acceptable for applications that are not as sensitive to the impact of SYNC latency. Performance testing is always recommended when deploying maximum availability mode.

Data Guard asynchronous redo transport with maximum performance mode is recommended when there is no zero-data-loss requirement or when the performance impact of RTT latency is too great to use maximum availability.

Oracle GoldenGate

Oracle GoldenGate is the best-practice recommendation for disaster recovery for non-Oracle database environments, for replication across heterogeneous Oracle Database releases, or for bidirectional replication when both replicas must be open in read-write mode at the same time.

Oracle GoldenGate is a comprehensive software package for enabling the replication of data in heterogeneous data environments. This high-performance software platform provides real-time capture, routing, transformation, and delivery of transactional data while imposing minimum system and network overhead. The software offers log-based, bidirectional replication and enables critical systems to support 24/7 operations. A typical environment would include capture, pump, and delivery processes.

Furthermore, Oracle GoldenGate enables the following:

- » Migration from other database platforms (for example, DB/2) to Oracle SuperCluster, while incurring minimal downtime
- » Active-active database instances for data distribution and continuous availability, minimal to zero downtime during planned (or unplanned) outages for disaster recovery, system migrations, upgrades, and maintenance
- » Real-time data warehousing or database consolidation on Oracle SuperCluster, from various sources including heterogeneous databases
- » Data capture from OLTP applications running on Oracle SuperCluster to support further downstream consumption such as SOA-type integration

Oracle GoldenGate Architecture

Oracle GoldenGate provides real-time, log-based change data capture and delivery between heterogeneous systems. Using this technology, the software enables a cost-effective and low-impact real-time data integration and continuous availability solution.

Oracle GoldenGate moves committed transactions with transaction integrity and minimal overhead on existing infrastructure. The architecture supports multiple data replication topologies such as one-to-many, many-to-many, cascading, and bidirectional. Its wide variety of use cases includes real-time business intelligence; query offloading; zero-downtime upgrades and migrations; and active-active databases for data distribution, data synchronization, and high availability.

The high-level Oracle GoldenGate architecture, shown in Figure 3, consists of three decoupled modules that facilitate the movement of transactional data to a target system. At any point before applying the data to the target system, Oracle GoldenGate can be used to execute a number of built-in functions, such as filtering and transformations.

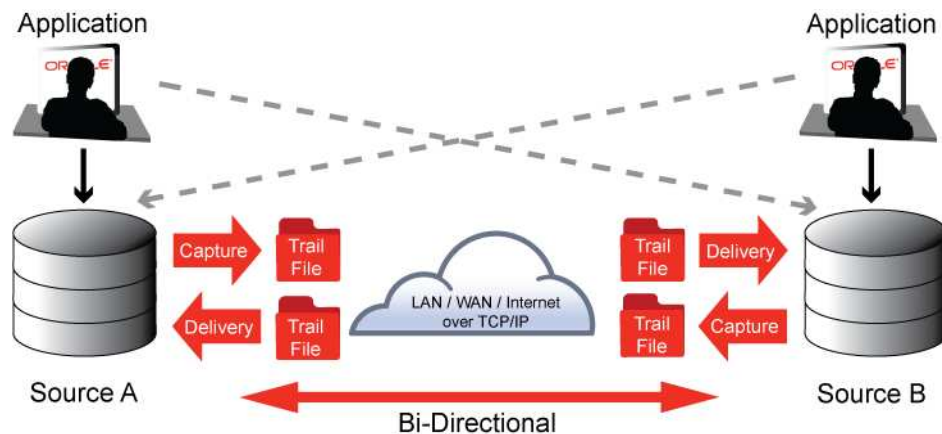


Figure 3. High-level Oracle GoldenGate architecture.

The Oracle GoldenGate modules do the following:

- » **Capture.** Oracle GoldenGate software captures changed data operations committed in the database transaction logs in a nonintrusive, high-performance, low-overhead implementation. The Capture module moves only committed transactions, which reduces infrastructure load and also eliminates potential data inconsistencies. Further optimization is achieved through transaction grouping and optional compression features.
- » **Trail Files.** The Trail Files contain the database operations for the changed data. Information is stored in a platform-independent data format.
- » **Delivery.** The Delivery module takes the changed data from the latest Trail File and applies it to the target database. Transactions are applied in the same order in which they were committed, for consistency and transactional integrity. Oracle GoldenGate can use a variety of transport protocols, and it can compress and encrypt changed data prior to routing. Transactional data can be delivered via Open Database Connectivity–compliant databases or through a specialized adapter to a Java Message Service message queue or topic.

Deploying Oracle GoldenGate 11g for Disaster Recovery

When configured for disaster recovery and data protection, Oracle GoldenGate provides a continuous availability solution that significantly improves recovery time for mission-critical systems. The disaster recovery and data protection configuration in Oracle GoldenGate complements Oracle Active Data Guard by offering continuous availability via active-active bidirectional database synchronization for non-Oracle databases, and for environments that require replication between different operating systems and Oracle Database versions. Oracle GoldenGate delivers up-to-the-second data to the backup system and enables immediate switchover to the new system if an outage occurs. It also immediately initiates real-time data capture from the standby database to update the primary system, once it is online, with any new data processed by the standby system.

Database Migration

Oracle GoldenGate supports an active-passive bidirectional configuration, in which Oracle GoldenGate replicates data from an active primary database to a full replica database on a live standby system that is ready for failover during planned and unplanned outages. This provides the ability to migrate to a database deployment on Oracle SuperCluster allowing the new system to work in tandem until testing is completed and a switchover is planned. Using Oracle GoldenGate for database migration is most applicable when reduced downtime is a requirement and Data Guard cannot be used for the database migration.

Non-Oracle Database Replication Tools

Although Oracle Active Data Guard and Oracle GoldenGate are generally recommended as a best practice for protecting database storage in Oracle SuperCluster, other non-Oracle database replication tools are also supported. Examples include SAP Sybase Replication Server and high availability disaster recovery (HADR) features in DB2. In addition, storage replication tools, such as Oracle's Pillar Axiom system, and non-Oracle replication tools, such as Hitachi's Replication Manager and EMC Replication Manager, can also provide data replication at a remote location for backup and disaster recovery purposes.

Support for these tools enables existing deployments that use non-Oracle databases and backup solutions to run unchanged on Oracle SuperCluster. In addition to providing support for legacy deployments, support for these products provides a migration path to a future disaster recovery solution featuring Oracle GoldenGate, if desired.

A full discussion of the available non-Oracle replication tools is beyond the scope of this paper. Please contact your Oracle representative for more information.

Database Recommended Use Cases

Oracle Active Data Guard, Oracle GoldenGate, non-Oracle database replication tools, or a combination of these can be used for database disaster recovery. Table 3 summarizes these options.

TABLE 3. COMPARISON OF DATABASE REPLICATION OPTIONS

Replication Requirement	Oracle Active Data Guard	Oracle GoldenGate	Non-Oracle Tools
Data protection/data availability/disaster recovery	✓	✓	
Database rolling upgrades	✓	✓	✓
Cross-platform migrations	✓	✓	
Zero downtime application upgrades		✓	
Active-active multimaster		✓	
Data integration		✓	
Many-to-one replication		✓	✓
Ability to replicate data subsets and transformations		✓	
Non-Oracle databases		✓	✓


The following disaster recovery use cases are recommended for databases running on Oracle SuperCluster.

» **Simple, Full Oracle Database Protection: Oracle Active Data Guard**

Oracle Active Data Guard is the recommended solution for complete replication of one Oracle Database instance to another. This approach is easy to implement and is suited for configurations that do not require schema or data subsets, target-side write capability, or heterogeneous combinations. Oracle Active Data Guard works with any application—custom or packaged—using any data types, as long as the databases are all Oracle Database, the platforms share a similar architecture, and the entire databases are replicated.

» **Flexible, Heterogeneous Database Protection: Oracle GoldenGate**

Oracle GoldenGate is recommended for all scenarios not covered by Oracle Active Data Guard. This includes any heterogeneous database or platform combination, any schema subsets, and active-active configurations. Note



that active-active configurations usually require data conflicts to be managed by the application, so such an architecture is more suitable for custom applications.

» **Combining Oracle Active Data Guard and Oracle GoldenGate**

Oracle Active Data Guard and Oracle GoldenGate offer additional advantages when used together. For example, a centralized global manufacturing database can be protected using an Oracle Active Data Guard physical standby, set up with Data Guard fast-start failover with synchronous redo transport, ensuring zero data loss and integrated failover of applications in the event of an outage at the primary data center. At the same time, using Oracle GoldenGate, it is possible to set up bidirectional replication configurations from this central database to smaller regional databases supporting local manufacturing operations. These can be non-Oracle databases, and they could also be configured in a hardware and operating system platform that is different from that of the central database. Enabling such a fully active, globally distributed and highly available configuration is one of the unique value propositions of implementing Oracle GoldenGate together with Oracle Active Data Guard.

With a wide array of continuous availability, disaster tolerance/recovery, and data integration/migration scenarios, the combination of Oracle Active Data Guard and Oracle GoldenGate provides a modular foundation that easily scales to address high-volume, low-impact data integration and replication challenges faced by enterprises of various sizes and complexities. In addition, Oracle SuperCluster also integrates with non-Oracle database replication tools, providing support and a migration path for existing legacy database deployments that currently use these tools.

Complementary Technologies

The following complementary technologies—Oracle Recovery Manager (Oracle RMAN), Oracle's Zero Data Loss Recovery Appliance, Oracle Solaris Cluster, Oracle Clusterware, and integrated system monitoring—can play a key part in disaster recovery strategies for Oracle SuperCluster.

Oracle Recovery Manager


Oracle SuperCluster works with Oracle RMAN to enable efficient Oracle Database backup and recovery. All existing Oracle RMAN scripts work unchanged in the Oracle SuperCluster environment. Oracle RMAN is designed to work closely with the server, providing block-level corruption detection during backup and restore. Oracle RMAN optimizes performance and space consumption during backup with file multiplexing and backup set compression.

Although the Oracle ZFS Storage Appliance that is internal to Oracle SuperCluster could be used by Oracle RMAN to back up Oracle Database, using an external Oracle ZFS Storage Appliance or Zero Data Loss Recovery Appliance is a recommended best practice. Backups inherently consume bandwidth and latency. If the internal storage appliance is used, the response time of other executables reading and writing to this storage will be impacted during backup and recovery.

Zero Data Loss Recovery Appliance

Oracle's Zero Data Loss Recovery Appliance—an engineered system for database backup that eliminates data loss exposure without impacting the performance of production environments—is an option that should be considered when traditional backup and recovery approaches are not sufficient to meet enterprise requirements. Compute, network, and storage are integrated into a massively scalable appliance with a cloud-scale architecture that provides fully automated database backup and recovery for multiple databases.

Featuring an incremental-forever backup strategy, the recovery appliance provides minimal-impact backups. The databases send only changes, and all backup and tape processing is offloaded from the production servers to the



appliance for improved system performance. Real-time database redo block information is transmitted, eliminating potential data loss and providing instant protection of new transactions. Database recoverability is improved with end-to-end reliability, visibility, and control of the database as a whole, rather than as a disjoint set of files.

The recovery appliance features secure replication to help protect against disasters such as site outages or regional disasters. Backups on a local recovery appliance can be easily and quickly replicated via secure transport to a remote recovery appliance. Flexible replication topologies are supported to match a data center's requirements. For example, replication can be set up in a simple one-way topology, or two recovery appliances can be set up to replicate each other, or a central recovery appliance can be used for replication from multiple satellite recovery appliances. In all topologies only changed blocks are replicated, minimizing WAN network usage.

Use of secure replication to a recovery appliance can help speed recovery times in the event of an outage. If the local recovery appliance is not available, restore operations can run directly from the remote recovery appliance without first staging the data locally.

Oracle's Zero Data Loss Recovery Appliance is a complementary technology to other backup and recovery options such as Oracle ZFS Storage Appliance and Oracle Active Data Guard. For example, an enterprise backup and recovery solution could use Zero Data Loss Recovery Appliance to provide a centralized backup service for all databases, use the snapshot and cloning capabilities of Oracle ZFS Storage Appliance for applications and other unstructured data, and use Oracle Active Data Guard to provide fast failover capabilities for critical databases.

Oracle Solaris Cluster

Oracle Solaris Cluster provides high availability with failover protection and helps automate failover procedures for applications and virtualized workloads that run on Oracle SuperCluster in traditional or cloud-based deployments. Although Oracle SuperCluster is designed with full redundancy at the hardware level, Oracle Solaris Cluster offers high availability for today's complex solution stacks, with failover protection from the application layer through the storage layer, including specific integration with Oracle ZFS Storage Appliance for NFS I/O fencing and lock release.


To limit outages due to single points of failure, mission-critical services can be run in clustered physical servers that efficiently and smoothly take over the services from failing nodes with minimal interruption to data services. Oracle Solaris Cluster offers built-in support for Oracle Database and other Oracle applications, with solution-specific failure detection and automatic recovery. A web-based user interface offers centralized management and access to status and configuration capabilities.

Oracle Solaris Cluster handles failover between Oracle Solaris Zones and Application Domains within Oracle SuperCluster. Tightly coupled with Oracle Solaris, Oracle Solaris Cluster detects failures without delay (zero-second delay) and provides much faster failure notification, application failover, and reconfiguration time. Applications run in an Oracle Solaris Cluster environment without modification. By coordinating dependencies across the entire solution stack, Oracle Solaris Cluster helps provide consistent failover and recovery capabilities for complex deployments.

Oracle Solaris Cluster Geographic Edition

Oracle Solaris Cluster Geographic Edition software, a layered extension to the Oracle Solaris Cluster software, supports multiple clusters that are separated by long distances. The clusters can be global clusters, zone clusters, or a combination of both. By using a secondary cluster with duplicated application configuration and replicated data, Oracle Solaris Cluster Geographic Edition enables a cluster to tolerate a disaster that disables the primary location.

The software provides a suite of tools to configure and manage geographically separated clusters, and it provides an automated (not automatic) mechanism for migrating services to a secondary site. Using Oracle Solaris Cluster Geographic Edition software, a set of clusters is configured. The primary cluster provides application services under



normal operation; a second cluster is configured to take over the primary cluster services if a disaster occurs. The Oracle Solaris Cluster Geographic Edition software manages configuration, data replication, and heartbeat monitoring between the local and remote clusters.

Oracle Solaris Cluster Geographic Edition software supports several options for software replication, including Data Guard and non-Oracle replication solutions. The Data Guard broker and fast-start failover at the database tier complement cluster failover and enable complete failover automation from one site to another. Oracle Solaris Cluster Geographic Edition software also has specific integration with Oracle ZFS Storage Appliance to automate Oracle ZFS Storage Appliance replication.

Oracle Solaris Cluster Geographic Edition 4.2 includes new features that are relevant for Oracle SuperCluster disaster recovery:

- » **Disaster recovery orchestration.** Orchestrated disaster recovery support enables Oracle Solaris Cluster to manage the automated and synchronized recovery of multiple applications and their respective replication solutions across multiple sites. A service constructed out of multiple tiers, possibly on multiple clusters, can be managed as a unit. This feature reduces risk and provides fast and reliable disaster recovery for multitiered services.
- » **Data Guard replication control for a remote database.** A local Oracle Solaris Cluster Geographic Edition protection group can be created to use the Data Guard broker to control the Data Guard replication of a database instance on a remote system. This feature leverages remote database connectivity, a standard feature of Oracle Database, and the Oracle Solaris Cluster HA for Oracle External Proxy. Using this feature, the disaster recovery of all tiers can be orchestrated, even if the database tier is on a system that isn't running Oracle Solaris Cluster.

Oracle Clusterware

For database failover support, Oracle Clusterware can be used. Oracle Clusterware is portable cluster software that allows the clustering of independent servers so that they cooperate as a single system. Oracle Clusterware is an independent server pool infrastructure, which is fully integrated with Oracle RAC, capable of protecting data access in a failover cluster.

There are APIs to register an application and instruct Oracle Clusterware regarding the way an application is managed in a clustered environment. The APIs are used to register the Oracle GoldenGate Manager process as an application managed through Oracle Clusterware. The process should then be configured to automatically start or restart other Oracle GoldenGate processes. Similarly, Data Guard works seamlessly with Oracle Clusterware.

Integrated System Monitoring

Oracle SuperCluster provides comprehensive monitoring and notifications to enable administrators to proactively detect and respond to problems affecting hardware and software components. With direct connectivity to the hardware components of Oracle SuperCluster, Oracle Enterprise Manager Ops Center can alert administrators to hardware-related faults and log service requests automatically through integration with Oracle Auto Service Requests for immediate review by Oracle Customer Support. Problems that would have required a combination of database, system, and storage administrators to detect them in traditional systems can now be diagnosed in minutes because of integrated systems monitoring for the entire Oracle SuperCluster platform.

- » **Oracle Configuration Manager** collects configuration information and uploads it to a management repository. This configuration data provides valuable information to customer support representatives, and can help reduce the resolution time for support issues and provide proactive problem avoidance.
- » **Oracle Enterprise Manager Ops Center 12c** helps IT staff understand and manage every architectural layer—from bare metal to operating systems and applications. It provides a centralized interface for physical and virtual machine lifecycle management, from power-on to decommissioning. In addition, it offers IT administrators a

unique insight into the user experience, business transactions, and business services, helping administrators to quickly detect changes in system health and troubleshoot issues across the entire environment.

Private and Hybrid Cloud Configurations

Oracle Optimized Solution for Secure Disaster Recovery supports both private and hybrid cloud configurations. In a hybrid cloud configuration, existing production databases remain on premises and standby databases used for disaster recovery are deployed on Oracle Cloud. Oracle Cloud offers a great alternative for hosting standby databases for customers who prefer not to deal with the cost or complexity of establishing and managing a remote data center.

Oracle Database Cloud Service can be used to deploy disaster recovery services for on-premises database systems. In this configuration, a Data Guard standby database is instantiated in the Oracle Database Cloud Service. Once instantiated, Data Guard maintains synchronization between the primary database on premises and the standby database in the cloud. If there is a complete site outage, the production applications and databases can fully run in Oracle Cloud. The standby database can also be used during planned maintenance, as well as during unplanned outages.

Customers can choose to deploy either a Data Guard or Oracle Active Data Guard standby database in the cloud, depending on their requirements. New cloud tools, including one-click automated backup with point-in-time recovery and one-click patching and upgrades, provide simple and fast management. Databases can be provisioned and ready for use in minutes, running on a dedicated virtual machine with preinstalled database software. Administrators have full administrative access to manage their databases, providing the same control as in a private cloud configuration.

For more information on Oracle Database Cloud Service, see cloud.oracle.com/database. For general information on Oracle Cloud, see oracle.com/cloud.

Best Practices for a Secure Disaster Recovery Implementation

Disaster recovery systems cannot rely solely on perimeter security. A combination of system-wide security measures and best practices—including the rule of least privilege, strong authentication, access control, encryption, auditing, disabling of unnecessary services, antimalware protections, and configuring system services for enhanced security—should also be implemented for secure operations.

Oracle highly recommends leveraging existing recommendations and guidelines from product security guides, Center for Internet Security (CIS) benchmarks, ISACA publications, and Department of Defense (DoD) Security Technical Implementation Guides (STIGs) when designing a disaster recovery environment.

Security Technical Implementation Guides

STIGs are continually updated and currently available for many Oracle products. A list of STIGs relevant to this solution is shown in Table 4.

TABLE 4. EXAMPLES OF RELEVANT STIGS

STIG	Location
Oracle Solaris	iase.disa.mil/stiqs/os/unix-linux/Pages/solaris.aspx
Oracle Database 11g Release 2	iasecontent.disa.mil/stiqs/zip/Apr2015/U_Oracle_Database_11-2g_V1R3_STIG.zip

Oracle Integrated Lights Out Manager	iase.disa.mil/stiqs/app-security/database/Pages/exadata_lights.aspx
Oracle Exadata Storage Server	iase.disa.mil/stiqs/app-security/database/Pages/exadata_storage.aspx
Oracle's Sun Datacenter InfiniBand Switch 36	iase.disa.mil/stiqs/app-security/database/Pages/exadata_infiniband.aspx
Oracle ZFS Storage Appliance	iase.disa.mil/stiqs/app-security/database/Pages/exadata_zfs.aspx
Oracle WebLogic Server 12c	iase.disa.mil/stiqs/Documents/u_oracle_weblogic_server_12c_v1r1_stig.zip
DoD Secure Telecommunications	iase.disa.mil/stiqs/net_perimeter/telecommunications/Pages/index.aspx
Oracle Linux 6 Manual STIG	iasecontent.disa.mil/stiqs/zip/Apr2015/U_Oracle_Linux_6_V1R2_STIG.zip
Storage Area Network (SAN)	iase.disa.mil/stiqs/Documents/u_storage_area_network_v2r2_stig.zip

For more STIGs, please see the website iase.disa.mil/stiqs/Pages/index.aspx.

Component-Level Security Recommendations

Oracle recommends the following component-level security guidelines.

- » **Change system default passwords.** Using known vendor-provided default passwords is a common way cyber criminals gain unauthorized access to infrastructure components. Changing all default passwords to stronger, custom passwords is a mandatory step during infrastructure deployment.
- » **Keep component patching current.** Ensure that all components are using the most recent firmware and software versions to the extent possible. This tactic ensures that each component is protected by the latest security patches and vulnerability fixes.
- » **Leverage isolated, purpose-based network interfaces.** Network interfaces, virtual or physical, should be used to separate architectural tiers, such as client access and management. In addition, consider using network interfaces to separate tiers within a multitier architecture. This enables per-tier security policy monitoring and enforcement mechanisms including network, application, and database firewalls as well as intrusion detection and prevention systems.
- » **Enable encrypted network communications.** Ensure all endpoints use encrypted network-based communications, including secure protocols, algorithms, and key lengths. For Oracle WebLogic, use the UCrypto provider to ensure that cryptography leverages the hardware assist capabilities of the SPARC platform.
- » **Enable encrypted data-at-rest protections.**
 - » Use encrypted swap, /tmp, and ZFS datasets for any locations that could potentially house sensitive or regulated data. This automatically takes advantage of cryptographic acceleration in Oracle Solaris.
 - » Use tape drive encryption to protect data that must leave the data center for off-site storage.
 - » For databases, use Transparent Data Encryption (TDE) to protect tablespaces that might store sensitive or regulated data. TDE automatically takes advantage of cryptographic acceleration in Oracle Solaris on SPARC systems.
- » **Secure the database.** Refer to Oracle Optimized Solution for Secure Oracle Database security best practices and recommendations.
- » **Deploy application services in Oracle Solaris non-global zones.** Deploying applications within Oracle Solaris non-global zones has several security advantages, such as kernel root kit prevention, prevention of direct memory and device access, and improved control over security configuration (via `zonecfg(1M)`). This approach also enables higher assurance auditing, because audit data is not stored in the Oracle Solaris non-global zone, but rather in the Oracle Solaris global zone.
- » **Implement a baseline auditing policy.** Use audit logs and reports to track user activity—including individual transactions and changes to the system—and to flag events that fall out of normal parameters. These should be implemented at both the Oracle Solaris and database levels. The baseline security audit policy should include login/logout activity, administrative actions, and security actions, as well as specific command executions for

Oracle Solaris. This tactic enables auditing of a core set of security-critical actions without overburdening the system or database.

- » **Follow the rule of least privilege.** Increase access control by granting only those privileges that a given individual needs. This should be implemented at both the enterprise resource planning (ERP) system level and the infrastructure level.
- » **Use strong authentication.** Many intellectual property attacks use stolen credentials. Implementing strong authentication methods, such as Kerberos, RADIUS, and SSL, can help prevent unauthorized access.
- » **Leverage role-based access control.** As the number of applications and users increases, user-based identity management can quickly become time consuming and labor intensive for IT staff. Consequentially, many users are granted inappropriate authorities. Though it requires increased efforts during the design and implementation phases, role-based access control (RBAC) is a popular option for low-maintenance, scalable access control, and it can help alleviate the burden of identity management.

Table 5 lists Oracle SuperCluster security recommendations. A full list of relevant component security recommendations is shown in Table 6.

TABLE 5. ORACLE SUPERCLUSTER SECURITY RECOMMENDATIONS

Title	Location
"Best Practices for Securely Deploying the SPARC SuperCluster T4-4"	oracle.com/technetwork/articles/servers-storage-admin/supercluster-security-1723872.html
"SPARC SuperCluster T4-4 Platform Security Principles and Capabilities"	oracle.com/us/products/servers-storage/servers/sparc-enterprise/supercluster/supercluster-t4-4/ssc-security-pac-1716580.pdf
"Oracle SuperCluster T5-8 Security Technical Implementation Guide (STIG) Validation and Best Practices on the Database Servers"	oracle.com/technetwork/server-storage/hardware-solutions/stig-sparc-supercluster-1841833.pdf
"Secure Database Consolidation Using the Oracle SuperCluster T5-8 Platform"	oracle.com/technetwork/server-storage/sun-sparc-enterprise/documentation/o13-053-securedb-osc-t5-8-1990064.pdf

TABLE 6. EXAMPLES OF COMPONENT SECURITY RECOMMENDATIONS

Resource	Location
<i>Oracle Solaris 11 Security Guidelines</i>	docs.oracle.com/cd/E36784_01/html/E36837/index.html
<i>Oracle Solaris 11.2 Security Compliance Guide</i>	docs.oracle.com/cd/E36784_01/pdf/E39067.pdf
"Secure Deployment of Oracle VM Server for SPARC"	oracle.com/technetwork/articles/systems-hardware-architecture/secure-ovm-sparc-deployment-294062.pdf
<i>Oracle Solaris Cluster Security Guide</i>	docs.oracle.com/cd/E39579_01/html/E39649/index.html
"User Authentication on the Solaris OS: Part 1"	oracle.com/technetwork/server-storage/solaris/user-auth-solaris1-138094.html
<i>Oracle ILOM Security Guide</i>	docs.oracle.com/cd/E37444_01/html/E37451/index.html
<i>Database Advanced Security Administrator's Guide</i>	docs.oracle.com/cd/E11882_01/network.112/e40393/toc.htm
"Oracle Database 12c Security and Compliance"	oracle.com/technetwork/database/security/security-compliance-wp-12c-1896112.pdf
"Best Practices for Deploying Encryption and Managing Its Keys on the Oracle ZFS Storage Appliance"	oracle.com/technetwork/server-storage/sun-unified-storage/documentation/encryption-keymgr-1126-2373254.pdf
<i>Securing the Network in Oracle Solaris 11.2</i>	docs.oracle.com/cd/E36784_01/html/E36838/index.html

<i>Securing Users and Processes in Oracle Solaris 11.2</i>	docs.oracle.com/cd/E36784_01/html/E37123/index.html
<i>Securing Systems and Attached Devices in Oracle Solaris 11.2</i>	docs.oracle.com/cd/E36784_01/html/E37121/index.html
<i>Securing Files and Verifying File Integrity in Oracle Solaris 11.2</i>	docs.oracle.com/cd/E36784_01/html/E37122/index.html
<i>Managing Encryption and Certificates in Oracle Solaris 11.2</i>	docs.oracle.com/cd/E36784_01/html/E37124/index.html
<i>Developer's Guide to Oracle Solaris 11 Security</i>	docs.oracle.com/cd/E36784_01/html/E36855/index.html
"Configuring Oracle GoldenGate Security"	docs.oracle.com/goldengate/1212/gg-winux/GWUAD/wu_security.htm#GWUAD354
"Managing Security for Backup Networks"	docs.oracle.com/cd/E26569_01/doc.104/e21477/network_security.htm#OBINS277

Example Best-Practices Implementation

This section summarizes an example best-practices implementation of disaster recovery for Oracle E-Business Suite running on Oracle SuperCluster. For complete details, refer to My Oracle Support Note 1558827.1, *Oracle E-Business Suite R12.1.3 Disaster Recovery: Implementation Guide on Oracle SuperCluster*. The step-by-step instructions in this note assume an Oracle E-Business Suite deployment, but the methodology is applicable to other applications running on Oracle SuperCluster. Information in this note is relevant to Oracle's SPARC SuperCluster T4-4 and all newer Oracle SuperCluster models.

Implementation Overview

The example implementation assumes that Oracle E-Business Suite 12.1.3 has been installed on two geographically separated sites for disaster recovery protection. The implementation guide provides step-by-step directions for switching over Oracle E-Business Suite from the primary site to the secondary site, simulating the disaster recovery behavior that would occur in the event of a failure. The guide also demonstrates the continuous replication and reverse replication features of Oracle ZFS Storage Appliance.

Data Guard is used to duplicate the database content. A physical standby database using the maximum availability mode is configured at a secondary site. Oracle RMAN is used to create the initial standby database. The replication features of Oracle ZFS Storage Appliance are used to provide continuous replication of the concurrent manager log and out files. Oracle Solaris Cluster and the Data Guard broker provide management capabilities to facilitate the switchover from one site to the other.

Disaster Recovery Setup

The following steps are used for disaster recovery setup. Refer to the My Oracle Support Note for complete details.

» Installation at primary site

Oracle E-Business Suite is first installed at the primary site. Two zone clusters are created on an Application Domain running Oracle Solaris, and Oracle E-Business Suite is installed on nodes in these zone clusters (see Figure 4). The database tier uses Oracle RAC running on a Database Domain.

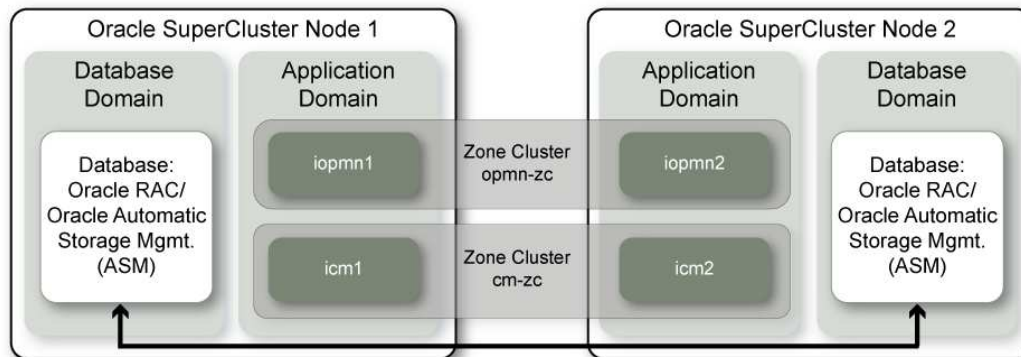


Figure 4. Logical architecture of Oracle E-Business Suite implementation on Oracle SuperCluster.

» Replication at secondary site

Oracle E-Business Suite is installed and configured at the secondary site, in the same manner as at the primary site. Then, Oracle RMAN is used to create a physical standby database at the secondary site.

» Oracle SuperCluster configuration

Oracle SuperCluster is used to manage the switchover between sites. The necessary files (`AutoConfig`, `listener.ora`, and `tnsnames.ora`) are configured for the Oracle Clusterware agent for the Oracle E-Business Suite database.

» Oracle ZFS Storage Appliance configuration at primary site

The continuous replication feature of Oracle ZFS Storage Appliance is used to synchronize the concurrent manager log and out files and continuously replicate those files from the primary to the standby. To provide for this capability, a new Oracle ZFS Storage Appliance project is created and presented as a share. Additionally, the new project is configured for Oracle Solaris Cluster fencing.

» Oracle ZFS Storage Appliance replication at secondary site

The remote replication services of Oracle ZFS Storage Appliance are used to replicate files from the primary site to the secondary site. A replication action of “Continuous” is specified to configure continuous replication.

» Switchover configuration

Oracle Solaris Cluster and the Data Guard broker are used to set up the switchover configuration. The cluster resources, created as part of the installation process, are confirmed to be set up correctly. Then Data Guard is configured to use the Data Guard broker to specify the primary database and the physical standby database (at the secondary location).

After manually verifying Data Guard role transitions with Data Guard broker and manually verifying Oracle ZFS Storage Appliance replication to change the replication direction back and forth, configure Oracle Solaris Cluster Geographic Edition to manage the entire environment for both Oracle SuperCluster units (see Figure 1Figure 5):

- » First configure one Oracle Solaris Cluster Geographic Edition partnership between the zone clusters for Concurrent Manager, and then configure a second partnership between the zone clusters for Oracle Process Manager and Notification.
- » Use the first partnership to configure an Oracle Solaris Cluster Geographic Edition protection group to manage the Data Guard configuration in the database domains and a second protection group to manage the Oracle ZFS Storage Appliance replication and the resource groups for the Concurrent Manager.

- » Then use the second partnership to configure a Solaris Cluster Geographic Edition protection group to manage the resource groups for Oracle Process Manager and Notification. An action script is provided to this protection group to activate an external name services update to remap the Oracle E-Business Suite application entry point host name to the appropriate IP address corresponding to the site.
- » Finally configure an Oracle Solaris Cluster Geographic Edition multigroup to orchestrate the three protection groups.

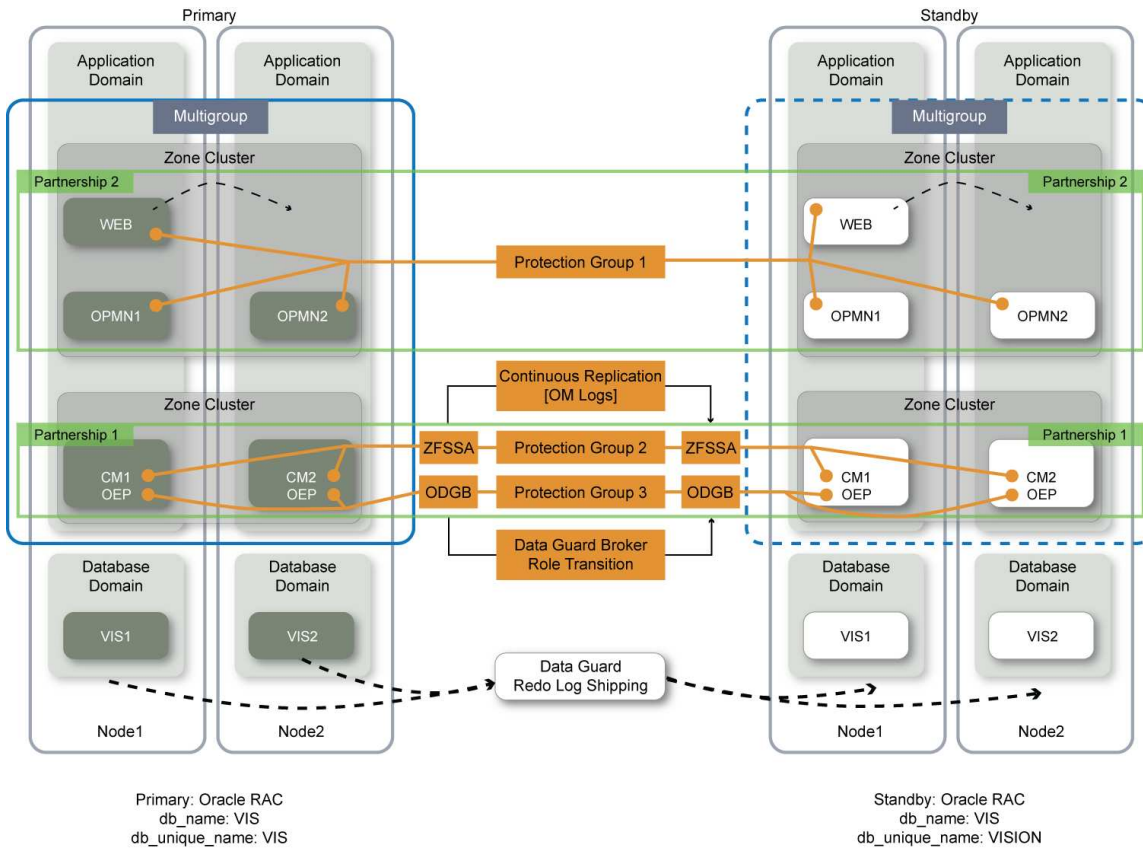


Figure 5. Example configuration for disaster recovery switchover testing.

Disaster Recovery Testing

At this point, Oracle E-Business Suite is installed on both the primary and secondary sites, and the primary database has been replicated to a physical standby database at the secondary site. Data Guard is providing replication of the database, and the Oracle E-Business Suite logs are being continuously replicated to the secondary site using Oracle ZFS Storage Appliance replication.

As described next, to verify the configuration, first perform a switchover from the primary to the secondary site. After confirming correct operation at the secondary site, perform a switchover back to the primary site.

» Switchover testing: primary to secondary

To test the switchover capability, first log in to the primary site and invoke an Oracle Solaris Cluster Geographic Edition multigroup switchover command to switch services from the primary site to the secondary site. This operation automatically orchestrates the stopping of the Oracle E-Business Suite services on the primary site, reverses the Data Guard role of the databases and the Oracle ZFS Storage Appliance replication direction, and

finally starts the Oracle E-Business Suite services on the secondary site. At that point Oracle E-Business Suite users can log in to the same application entry point, but they are now serviced from the secondary site.

» **Switchover testing: secondary to primary**

Then, test the switchover from the secondary site back to the primary site. Similar to the previous section, this involves initiating a switchover from the secondary site to the primary site. This is again achieved by invoking an Oracle Solaris Cluster Geographic Edition multigroup switchover command to switch services from the secondary site to the primary site.

» **Takeover testing: primary to secondary**

To test the takeover capability, first shut down (even powering off) the primary site. Then log on to the secondary site and invoke an Oracle Solaris Cluster Geographic Edition multigroup takeover command to bring up services on the secondary site. This operation automatically orchestrates the change of Data Guard role to the database on the secondary site and the Oracle ZFS Storage Appliance replication to make the replicated project active on this site. It then starts the Oracle E-Business Suite services on the secondary site. At that point Oracle E-Business Suite users can login to the same application entry point, but they are now serviced from the secondary site.

Summary

Planning for protection from disasters and other catastrophic events is essential for businesses and organizations, and most mission-critical enterprise deployments configure a remote standby site for this purpose. In the event of a disaster, activity can transfer to the standby site for continued operation. Oracle SuperCluster supports a range of options to provide the disaster recovery solution that best meets an organization’s needs.

Table 7 summarizes the components used in a typical Oracle SuperCluster disaster recovery solution.

TABLE 7. ORACLE SUPERCLUSTER DISASTER RECOVERY COMPONENTS

Category	Product	Business Need or Deployment Characteristic
Applications and Unstructured Data	Oracle ZFS Storage Appliance	Recommended best practice for remote replication
	Non-Oracle replication tools	Legacy deployments using non-Oracle tools
Database	Data Guard/Oracle Active Data Guard	Recommended best practice for disaster recovery for very large Oracle Database environments
	Oracle GoldenGate	Recommended best practice for disaster recovery for non-Oracle database environments, heterogeneous Oracle environments, or to implement Oracle configurations that use bidirectional replication
	Non-Oracle database replication tools	Legacy deployments using non-Oracle tools
Management	Oracle Solaris Cluster Geographic Edition	Management and automation of application failover
	Oracle Clusterware	Database failover management
	Oracle RMAN	Database backup and migration

Oracle ZFS Storage Appliance replication technology is recommended for disaster protection of middle-tier applications and components running on the cluster. Additionally, Oracle Active Data Guard or Oracle GoldenGate are recommended to provide disaster recovery for databases that are part of Oracle SuperCluster deployments. Non-Oracle replication tools are also supported, providing integration with legacy and other non-Oracle databases. Complementary technologies, such as Oracle Solaris Cluster and Oracle Clusterware, can also be used to help automate the failover and recovery process.

To get the latest information on Oracle Optimized Solution for Secure Disaster Recovery, please see My Oracle Support Note [1558852.1](#).

Disaster recovery planning is a critical component for ensuring continued business operations in the event of a disaster. There are many complex factors to consider, and each Oracle SuperCluster implementation has its own unique business requirements. Please contact your Oracle representative or Oracle Consulting for more information on architecting an Oracle SuperCluster disaster recovery solution that follows best practices and helps provide the necessary levels of data protection.

References

For more information, visit the web resources listed in Table 8.

TABLE 8. WEB RESOURCES FOR FURTHER INFORMATION

Description	Web Resource URL
Oracle Maximum Availability Architecture	oracle.com/maa
Oracle Optimized Solutions	oracle.com/optimizedsolutions
Oracle SuperCluster	oracle.com/supercluster
Oracle ZFS Storage Appliances	oracle.com/zfsstorage
Zero Data Loss Recovery Appliance	oracle.com/engineered-systems/zero-data-loss-recovery-appliance/
Oracle Active Data Guard	oracle.com/us/products/database/options/active-data-guard/overview/index.html
Oracle GoldenGate	oracle.com/goldengate
Oracle Database	oracle.com/database
Oracle Solaris	oracle.com/solaris
Oracle Optimized Solution for Oracle SuperCluster Disaster Recovery Oracle Support Document 1558852.1	https://support.oracle.com/epmos/faces/DocumentDisplay?id=1558852.1







Oracle Corporation, World Headquarters

500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries

Phone: +1.650.506.7000
Fax: +1.650.506.7200

CONNECT WITH US

-  blogs.oracle.com/oracle
-  facebook.com/oracle
-  twitter.com/oracle
-  oracle.com

Integrated Cloud Applications & Platform Services

Copyright © 2015, 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 of The Open Group. 0615

Oracle Optimized Solution for Secure Disaster Recovery: Highest Application Availability with Oracle SuperCluster
October 2015
Author: Dean Halbeisen