

SunGuide[®]:

High Availability and Satellite Configuration



Prepared for:

Florida Department of Transportation
Traffic Engineering and Operations Office
605 Suwannee Street, M.S. 90
Tallahassee, Florida 32399-0450
(850) 410-5600

November 3, 2011

Document Control Panel

File Name:	SunGuide-Satellite.docx		
File Location:	SunGuide CM Repository		
CDRL:			
	Name	Initial	Date
Created By:	Mary K Thornton, SwRI	MKT	06/23/10
Reviewed By:			
Completed By:			

Table of Contents

1. Scope	8
1.1 Document Identification.....	8
1.2 Terminology	8
1.3 Project Overview	8
1.4 Related Documents	9
1.5 Contacts.....	9
2. High Availability	10
2.1 Software-Specific Considerations	10
2.1.1 System Availability	10
2.1.2 Current Version	10
2.1.3 Configuration Files	10
2.1.4 Archived Data Files	11
2.1.5 User-Defined Types.....	11
2.1.6 Center to Center	11
2.2 Sources of Downtime	11
3. Unplanned Downtime.....	12
3.1 Application Server Failure	12
3.1.1 Microsoft Windows Clustering	12
3.1.2 Virtual Servers.....	13
3.2 Database Server Failure	14
3.2.1 Oracle Solutions	14
3.2.1.1 Oracle RAC.....	14
3.2.1.2 Oracle Fail Safe.....	15
3.2.2 Virtual Servers	16
3.3 Data Failure	16
3.3.1 Oracle Flashback	16
3.3.2 Database Backups.....	16
3.3.2.1 Oracle RMAN.....	17
3.3.2.2 Oracle Secure Backup.....	17
3.3.2.3 Oracle Data Pump.....	17
3.3.3 Oracle Automatic Storage Management.....	17
3.3.4 Oracle Data Guard	18

3.3.4.1 Standby Database Type.....	18
3.3.4.2 DataGuard Protection Modes.....	19
3.3.4.3 Oracle Active Data Guard.....	19
3.4 Site Failure	20
3.4.1 ‘Cold’ Backup Site	20
3.4.2 Virtual Servers.....	21
3.4.3 Oracle Data Guard	21
3.5 Communications Failure	22
4. Planned Downtime	22
4.1 System Changes	22
4.1.1 Database Storage System Changes.....	22
4.1.1.1 Oracle Streams.....	22
4.1.1.2 ASM Online Reconfiguration	22
4.1.2 Application Server Changes	22
4.1.2.1 Microsoft Windows Clustering.....	22
4.1.2.2 Virtual Servers	23
4.1.3 Database Server Changes	23
4.1.3.1 Oracle Online Patching.....	23
4.1.3.2 Oracle Fail Safe.....	23
4.1.3.3 Oracle Data Guard	23
4.1.3.4 Oracle Streams	24
4.2 Application/Data Changes.....	24
4.2.1 Satellite or Backup Site	24
4.2.2 Oracle Data Guard	24
4.2.3 Microsoft Clustering.....	24
4.2.4 Online Table Redefinition	25
4.2.5 Oracle Edition-based Redefinition	25
5. Examples	25
5.1 Overview	25
5.2 Orlando-Orange County Expressway Authority (OOCEA).....	25
5.2.1 Application Server Failure.....	27
5.2.2 Database Server Failure.....	27
5.2.3 Data Failure	27
5.2.4 Site Failure.....	27

5.2.5 Communications.....	27
5.2.6 System Changes.....	27
5.2.7 Application/Data Changes	27
5.3 TIMSO	28
5.3.1 Application Server Failure.....	28
5.3.2 Database Server Failure.....	28
5.3.3 Data Failure	28
5.3.4 Site Failure.....	28
5.3.5 Communications.....	28
5.3.6 System Changes.....	28
5.3.7 Application/Data Changes	28
5.4 District 5 RTMC.....	28
5.4.1 Application Server Failure.....	30
5.4.2 Database Server Failure.....	31
5.4.3 Data Failure	31
5.4.4 Site Failure.....	31
5.4.5 Communications.....	31
5.4.6 System Changes.....	31
5.4.7 Application/Data Changes	31
6. Other	31
7. Future Considerations.....	32

1. Scope

The purpose of this document is to provide an analysis of options for installations of the SunGuide application with regards to system availability in the absence of specific system availability requirements.

As a part of contingency planning, system availability requirements may be identified such as redundant real-time mirroring at an alternate site or fail-over capabilities. This document makes an assumption about one of the system availability requirements of a Florida Traffic Management Center running SunGuide. The assumption is that the maximum acceptable recovery time for the SunGuide application is fifteen minutes.

It is recommended that some analysis be done to establish a set of system availability requirements. Some other contingency measures that should be considered are redundant communications paths, lack of single points of failure, enhanced fault tolerance of network components and interfaces, powerful management systems with appropriately sized backup power sources, load balancing, and data mirroring and replication to ensure a uniformly robust system.

1.1 Document Identification

In general, this document presents options available for minimizing downtime of SunGuide software systems. It will present solutions for downtime as a result of both planned and unplanned events. As a part of that discussion, this document addresses downtime caused by site failure or unavailability through the addition of a secondary ‘satellite’ or ‘backup’ site. It includes discussion of how that site may be deployed and managed, from a high-level perspective.

This document is intended for SunGuide ITS Managers and Network Administrators to use as a resource in planning, designing, procuring system equipment and resources, and deploying a highly available SunGuide system. In particular, this document may be especially helpful for those who are intending to deploy a secondary ‘satellite’ or ‘backup’ site for disaster recovery.

1.2 Terminology

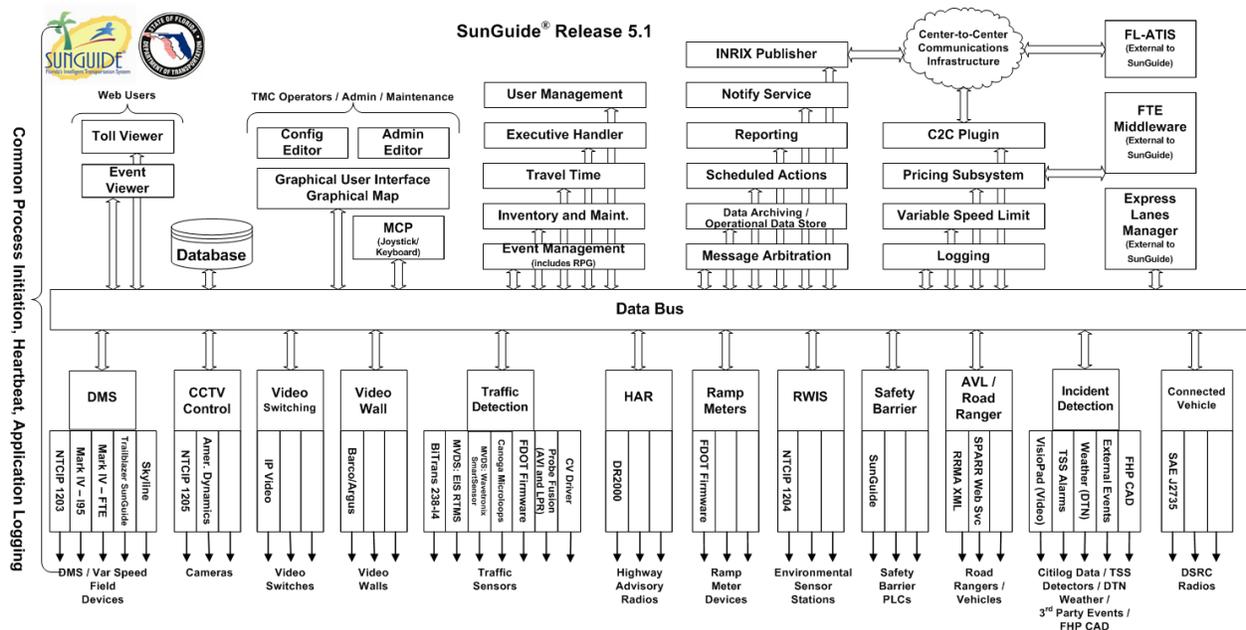
Throughout this document, the terms principal and secondary will denote the intended use of a site or database. If a SunGuide system is set up in such a way that there are two databases or two sites, the term ‘principal’ will indicate the database that is intended to be used for normal operations, regardless of the role it is currently playing. In the same way, the site that is intended to be used for daily operations will be referred to as the principal site.

The term active is used to indicate a component of the system that is operationally in use.

1.3 Project Overview

The Florida Department of Transportation (FDOT) is conducting a program that is developing SunGuide software. The SunGuide software is a set of Intelligent Transportation System (ITS) software that allows the control of roadway devices as well as information exchange across a variety of transportation agencies. The goal of the SunGuide software is to have a common

software base that can be deployed throughout the state of Florida. The SunGuide software development effort was based on ITS software available from the state of Texas. In addition to the reuse of software (along with customization of this software), a number of new software modules are being developed. The following figure provides a graphical view of the software described in this document:



1-1 - High-Level Architectural Concept

1.4 Related Documents

- USAID, *Disaster Recovery Planning Procedures and Guidelines: A Mandatory Reference for ADS Chapter 545*.
- Oracle, *High Availability with Oracle Database 11g Release 2*, September 2009.
- NIST, *Contingency Planning Guide for Information Technology Systems: Recommendations of the National Institute of Standards Technology*.
- Oracle, *Switchover and Failover Best Practices: Oracle Data Guard 10g Release 2*, Dec. 2008
- Oracle, *Oracle Data Guard Concepts and Administration 11g Release 2 (11.2)*, Dec. 2009.
- Oracle, *Oracle Database High Availability Overview*
- SunGuide Project website: <http://sunguide.datasys.swri.edu>.

1.5 Contacts

- Mary Thornton, SwRI Developer, mthornton@swri.org, 407-284-0855
- Robert Heller, SwRI Software Project Manager, rheller@swri.org, 210-522-3824
- Arun Krishnamurthy, FDOT SunGuide Project Manager, Arun.Krishnamurthy@dot.state.fl.us, 850-410-5615

2. High Availability

High availability has to do with operational continuity over measured periods of time. Generally, the term downtime is used to refer to periods when a system is unavailable, or that continuity is disrupted. In the case of Intelligent Transportation Systems, downtime of any kind is always something to be avoided.

Contingency planning should be done to deploy and configure a SunGuide solution that meets the availability needs of the organization. This document presents alternatives and options for system deployment and configuration that are available for maximizing system availability.

2.1 Software-Specific Considerations

When preparing any high availability contingencies for the SunGuide Advanced Traffic Management System, there are a few application-specific things to keep in mind.

2.1.1 System Availability

When system availability requirements and options are discussed in the context of the SunGuide application, they should be defined in terms of operational availability or full system availability. Operational availability refers to the ability of operators/users to have use of the SunGuide application for monitoring, control, and management of traffic systems. However, full system availability includes preserving the integrity and consistency of the historical data that is currently archived in the SunGuide database and on file systems. It is possible to engineer a highly available system in terms of operational availability that includes disruption or loss of historical data. The needs of the deployment must be weighed to determine the acceptability of such loss.

2.1.2 Current Version

Each release of the SunGuide software is accompanied by a Version Description Document (VDD). The VDD should be consulted to determine version-specific deployment requirements that might have changed since the writing of this document. This document is written based on SunGuide version 5.1.0.

2.1.3 Configuration Files

SunGuide makes use of configuration files that reside on file systems and shares within access to the application. Any deployment or redundancy strategy should consider how these files will be synchronized or kept up to date between sites. In the current version of SunGuide, these files include:

- The SunGuide configuration file (config.xml)
- The IP Video switching and Snapshot devices configuration files (IpVideoDevices.xml, SnapshotDevices.xml)
- The SAS configuration files (sequences.xml, schedules.xml)
- The Operator Map configuration file (OMInterface.dll.config)

2.1.4 Archived Data Files

The current SunGuide software provides a feature that allows some data types to be archived, not only to the SunGuide database, but to data files on a designated file system. If historical continuity is to be preserved, these files will need to be synchronized or kept up to date between sites. The location of these files can be determined from a review of the SunGuide configuration file.

2.1.5 User-Defined Types

The Event Management components of the SunGuide database incorporate non-standard user-defined types to implement nested database tables. Conceptually, this marks a divergence from the relational database standard used by the rest of the application to an object-oriented database approach. In practice, the existence of user-defined types complicates the streaming of replicated data from one Oracle database to another and can cause logical-based replication solutions to fail. Solutions and options for this issue are discussed where it pertains to the Oracle software solution.

2.1.6 Center to Center

Many SunGuide installations interact through Center to Center with other SunGuide deployments. Any solution discussed here will need to be analyzed in terms of network connectivity and communications, if this interaction is to be preserved. In addition, interactions with the Florida Advanced Traffic Information System (FL-ATIS) should be considered, where applicable, when developing a highly available SunGuide solution.

2.2 Sources of Downtime

In this document, a distinction is made between planned downtime and unplanned downtime.

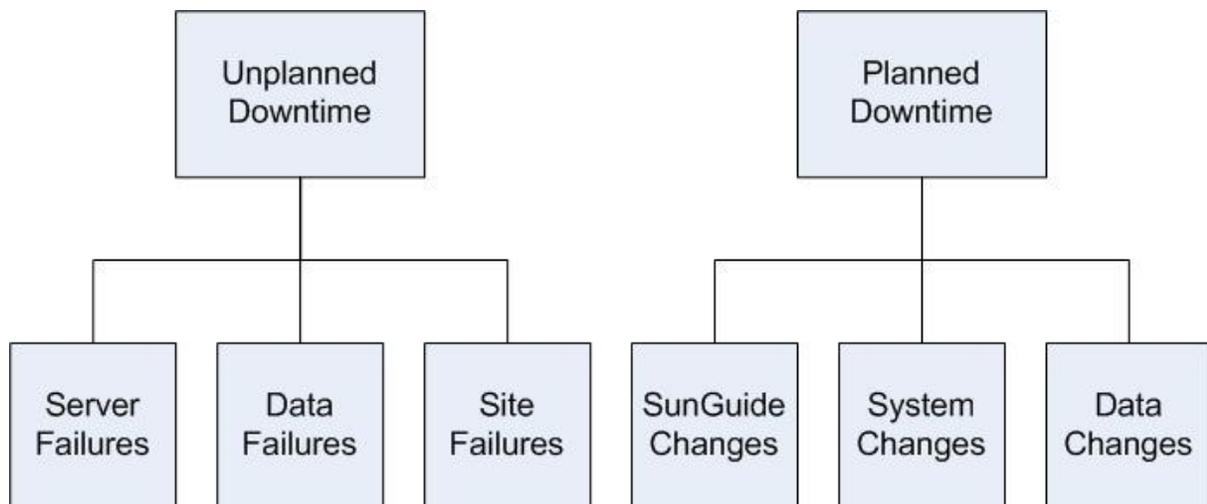


Figure 2.1 – Sources of Downtime

3. Unplanned Downtime

This section discusses each of the causes of unplanned system downtime and presents some existing options for minimizing this downtime. Unplanned downtime events typically arise from some physical event, such as a hardware or software failure or an environmental anomaly.

The NIST refers to contingency planning as a "coordinated strategy involving plans, procedures, and technical measures that enable the recovery of IT systems, operations, and data after a disruption." Because unplanned downtime scenarios possibly could occur, it is recommended that some sort of contingency plan be prepared to handle the scenarios and minimize the system downtime.

3.1 Application Server Failure

An application server failure occurs when the machine hosting a SunGuide application service, either a windows service or a web service, unexpectedly fails, often due to a hardware fault. Hardware faults are essentially unpredictable. Some options for minimizing downtime due to application server failure are presented below.

3.1.1 Microsoft Windows Clustering

In the Microsoft Windows environment in which SunGuide operates, many locations choose to deploy the SunGuide services within a Microsoft Windows Clustering solution. This is a well-understood solution that has successfully served SunGuide deployments for several years. It requires the acquisition of Microsoft Windows Enterprise Server licenses for each of the servers that are to be included in the cluster.

SunGuide services are redundantly installed on a set of two or more servers that are 'clustered' together. Each server is considered a 'node' in the cluster. Related SunGuide service resources run in cluster groups. Each cluster group includes service resources, along with a network hostname and/or IP address resource for each group. That clustered hostname or IP address is used in the SunGuide configuration file to identify which server, or node, currently hosts the service resources. The system may be load balanced by having groups optimally distributed across all the servers in the cluster.

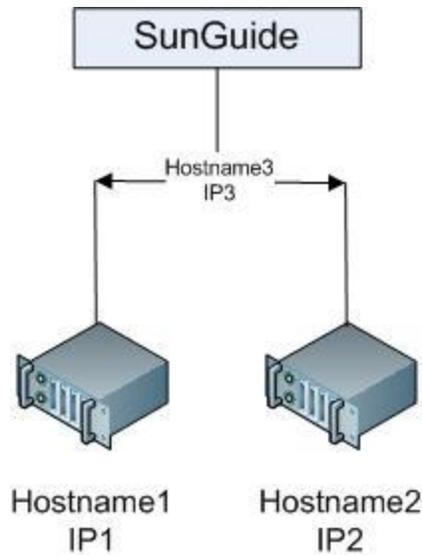


Figure 3.1 – Clustered Application Servers

If one of the services in a cluster group fails, the cluster group fails over to another node, and the clustered hostname and/or IP address moves with it. If the node itself fails, then all the services hosted on that server would automatically failover to one of the other existing application servers. Because the services failover with their hostname or IP address, no SunGuide configuration changes are required to move services from one server to another.

3.1.2 Virtual Server Clustering

Virtual servers are being used at some SunGuide locations to provide high availability of SunGuide services. Platform virtualization is performed on a given hardware platform by host software (like VMWare), which creates a simulated computer environment, a virtual machine (VM), for its guest software. The guest software is not limited to user applications; most hosts allow the execution of complete operating systems. The guest software executes as if it were running directly on the physical hardware.

The virtual servers may be clustered against physical hosts in much the same way that Microsoft Clustering may be used to cluster SunGuide application services. Copies of servers may be kept available and offline to be brought online in the event of a virtual server failure. While this type of clustering does not protect against application failure due to user error, Windows Updates, or other operating system changes, it can be used in conjunction with MS Cluster Service to provide a highly available system that responds well to both server hardware failure and application failure due to server software conflicts and other issues of that nature.

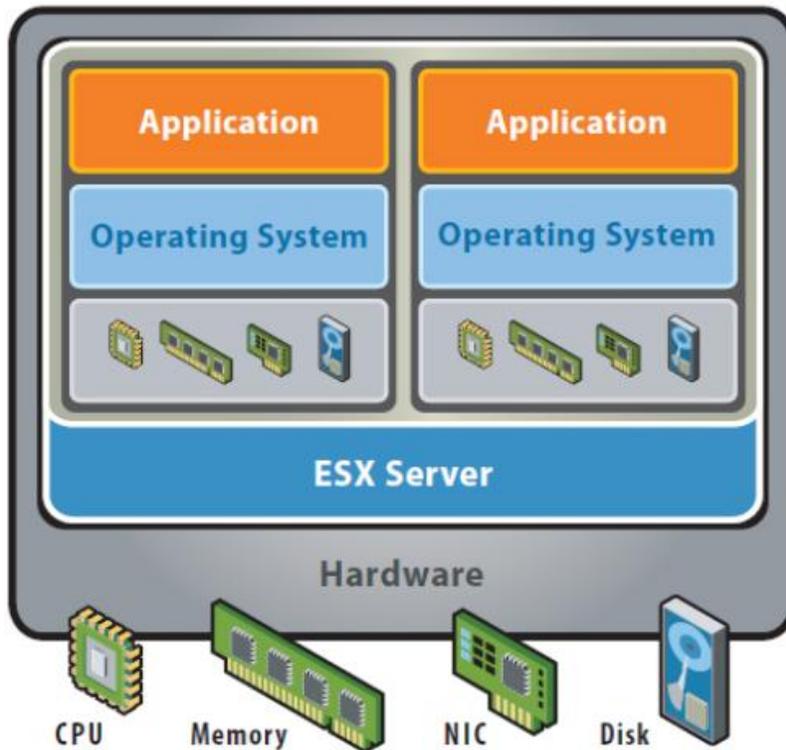


Figure 3.2 – Clustered Virtual Servers

3.2 Database Server Failure

3.2.1 Oracle Solutions

Oracle provides two different solutions for handling database server failure: Oracle Real Application Clusters (RAC) and Oracle Fail Safe.

3.2.1.1 Oracle RAC

Oracle recommends Real Application Clusters (RAC) as a clustering solution that can respond robustly to server failure. The RAC solution has the important advantage of no downtime.

In RAC, a database instance is concurrently serviced by more than one node. Sessions and transactions are balanced between the two nodes. If one node crashes, the second one starts recovering, and only the uncommitted transactions and session state variables on the failed instance are lost. RAC also brings with it an abundance of other useful features like load balancing.

However, this solution requires the purchase of an additional Oracle software license per CPU. RAC is not standard with Oracle Enterprise Edition. Also, the use of RAC requires the use of the included Oracle Clusterware software for clustering the SunGuide application services, instead of the Microsoft Windows Clustering solution. RAC should optimally be deployed on a

RAID-enabled disk subsystem in a network that provides high speed interconnect between the nodes in the cluster.

At this time, there are no known SunGuide deployments using RAC.

3.2.1.2 Oracle Fail Safe

Oracle Fail Safe provides the ability for the SunGuide Oracle database to be clustered using Microsoft Windows Clustering in much the same way as the SunGuide application services. While RAC provides the minimal downtime, Microsoft Clustering solution, used in conjunction with Oracle Fail Safe and combined with shared storage location on a redundant array, is an adequate solution for server failure recovery for the SunGuide system.

Oracle Fail Safe is a well-known and well-understood solution that has served SunGuide deployments for several years. It requires the purchase and retention of Microsoft Windows Enterprise Server 2003 licenses for each of the servers that are to be included in the cluster.

With Oracle Fail Safe, only one node services the database instance at any one time. There is no mechanism available to do load balancing of the database across the nodes. If the hosting node fails, all uncommitted transactions and session state variables currently against the database instance are lost. Oracle FS will start up the database instance on another node. It mounts, recovers, and opens the database within a couple of minutes, because it must wait until the surviving node brings up the database dependencies, such as the database services and listener.

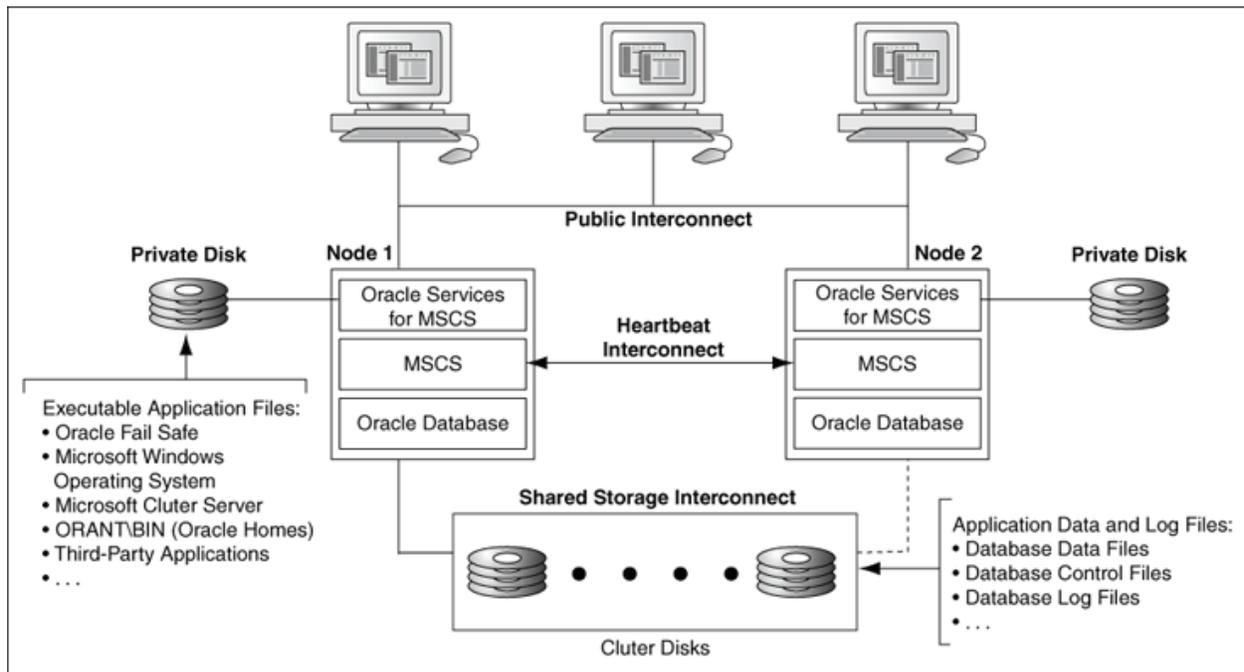


Figure 3.3 – Oracle Fail Safe

While this solution involves some downtime, the downtime should be well within the currently understood acceptable limits.

3.2.2 Virtual Servers

The SunGuide database server may be clustered, in the same way discussed for the SunGuide application servers, across multiple hardware hosts using virtual host software. Database licensing can get expensive using this method, due to the method Oracle uses to license their software by processor. Oracle does provide a virtualization software that bypasses this problem and may be a better solution for those who wish to cluster their Oracle databases on virtual servers. As of the time of this document, this software has not been tested with SunGuide, although the SunGuide database has been hosted successfully by the VMWare product.

3.3 Data Failure

Data failure is the loss, damage, or corruption of business-critical data. Generally, data failure is caused by storage subsystem failure, site failure, human error, database failure, and/or data corruption.

3.3.1 Oracle Flashback

Oracle provides Oracle Flashback for specific data failures. Oracle Flashback operates at the business object level, repairing tables or recovering specific transactions. Flashback Technologies are a unique and rich set of data recovery solutions that enable reversing human errors by selectively and efficiently undoing the effects of a mistake. Correcting an error takes about as long as it took to make it.

The solution is not dependent on the database size. Recovery can take place for a row, transaction, table, or even for the entire database. With Flashback Transaction, a single transaction and all of its dependent transactions can be flashed back with a single PL/SQL operation or by using a wizard to identify and flashback the problem transactions. It relies on undo data and archived redo logs to back out the changes.

The Oracle Flashback solution is recommended in the case of a change being made to the SunGuide configuration (in the database) or data entry that includes a human error. Oracle Flashback can be used to kick that change back out, in the event that the change cannot be reversed through the existing SunGuide tools. This can be a much quicker and more acceptable solution than recovering a backup of the SunGuide database, should an appropriate backup even be available.

3.3.2 Database Backups

Should the data failure fall outside of a single error or issue, then recovering a database backup may be the optimal solution. It has long been understood that database backups are an important part of database administration. "Physical" database backups are file-level copies of the database. This includes shutting down the database and doing a copy of all data, log, and control files and using the Recovery Manager (RMAN) utility to backup the database, either off-line or on-line. Data Pump Exports are "logical" database backups in that they extract only the logical definitions and data from the database to a file or files. It is advisable to use more than one method to backup your database. In addition, all backup and recovery scenarios should be tested carefully. It is better to be safe than sorry. No matter what type of backup is performed, point-in-time recovery is only available when the database is in ARCHIVELOG mode.

3.3.2.1 Oracle RMAN

Oracle's backup and recovery solution (RMAN) can do fine-grained, efficient recovery of individual blocks instead of entire data files. RMAN can also optimally keep track of changed blocks, ensuring that only changed blocks get backed up, thus providing a powerful implicit deduplication capability.

Large databases can be composed of hundreds of files, making backup very challenging. Missing even one critical file can render the entire database backup useless. Worse, incomplete backups go undetected until they are needed in an emergency. Oracle Recovery Manager (RMAN) is the core Oracle Database software component that manages database backup, restore, and recovery processes. RMAN maintains configurable backup and recovery policies and keeps historical records of all database backup and recovery activities. RMAN ensures that all files required to successfully restore and recover a database are included in complete database backups. Furthermore, as part of RMAN backup operations, all data blocks are verified to ensure that corrupt blocks are not propagated into the backup files.

The SunGuide database is recommended to be backed up daily, preferably during a slow traffic period, such as early morning, because, depending on the type of backup, performance of the database may be impacted. The RMAN application may be used, either from the command line or through the Enterprise Manager application. If Oracle Data Guard is deployed, physical backups may be able to be performed from the physical standby database instead of the primary database. This may help reduce the performance hit and eliminate downtime due to these kinds of backups.

3.3.2.2 Oracle Secure Backup

Oracle provides a tape backup management solution for database and file system data, should this be an option that a TMC would like to pursue. Oracle Secure Backup does require an additional licensing cost, as it is not included in the Standard or Enterprise Edition of the Oracle 11g database.

3.3.2.3 Oracle Data Pump

Logical backups may be taken by using the Oracle Data Pump Export utility. It is recommended to do an export of the SunGuide database and copy to an off-site location at least once a week.

3.3.3 Oracle Automatic Storage Management

Traditionally, a Storage Area Network (SAN) is used for file storage that provides storage failure protection.

Oracle database's Automatic Storage Management (ASM) is a RAW device management tool that gives much of the functionality of a file system, and provides storage failure protection. ASM stripes and mirrors everything on multiple disks. When disk failures occur, system downtime is avoided by using the data available on the mirrored disks. If the failed disk is permanently removed from ASM, the underlying data is striped or rebalanced across the remaining disks to continue delivering high performance.

These are pretty much the same features available from the use of a SAN, and one could argue that ASM is unnecessary in most cases. That would be correct. However, to put Oracle files on

disk, you need to use either RAW disk access or some other file system. Many file systems do not support RAC, because the system has to be able to permit multiple servers accessing the same file, even the same block, concurrently. Cluster file systems, which do allow this functionality, are often either very expensive or have limited support available. ASM is a good alternative in this scenario.

A useful feature of ASM is the rolling upgrades available in 11g. Rolling upgrades permit administrators to keep applications online while they upgrade ASM on individual nodes by keeping other nodes in the cluster available during the migration. The ASM instances can run at different software versions until all nodes in the cluster have been upgraded. Also, in release 2 of 11g database, the ASM Cluster File System extends ASM functionality to support non-Oracle db files. This includes application executables, configuration files, trace files, image files, etc.

ASM is incompatible with Oracle Fail Safe, but Oracle Fail Safe is usually used in a situation where the Oracle files are stored on a traditional type of storage disk.

3.3.4 Oracle Data Guard

Oracle Data Guard, a software replication solution available with the Oracle Enterprise Edition license, provides a comprehensive set of services that create, maintain, manage, and monitor one or more duplicate databases to enable production Oracle databases to survive disasters and data corruptions.

In keeping with Oracle Data Guard terminology, the terms ‘primary’ and ‘standby’ are used to indicate the role that a particular database is currently playing. For example, if daily operations are currently transferred to secondary database, then that database becomes the primary database, and the principal database becomes the standby database. Data Guard maintains standby databases as copies of the current primary database.

Then, if the primary database becomes unavailable because of a planned or an unplanned outage, Data Guard can switch any standby database to the primary role, minimizing the downtime associated with the outage. Data Guard can be used with traditional backup, restoration, and cluster techniques to provide a high level of data protection and data availability.

Also, Data Guard provides very fast automatic failover (referred to as fast-start failover) in database failures, node failures, corruption, and media failures. It can be set up in such a way that the failover is transparent to the SunGuide application services.

Oracle Data Guard will validate all redo blocks before transmitting and applying them to the standby database, so no corruptions will be propagated to the standby database. Data Guard also provides a mechanism to allow some operations, like RMAN backups, to be performed against a standby database, instead of against an active database, thereby tying up the database resources.

3.3.4.1 Standby Database Type

To do this, Data Guard transmits redo data from the primary to the standby database server, which is then applied to the standby database.

There are two ways that this data is applied to the standby database. If the standby database is a logical standby database, then SQL Apply is used to apply the Data Definition Language (DDL)

and Data Manipulation Language (DMS) changes. If the standby database is a physical standby, then the changes are applied on the file level.

Because SunGuide has user-defined types, the physical standby database is recommended. However, if a logical standby is required in order to leverage some of the logical standby database features, SwRI can provide some scripts that will provide translation of the user-defined types to standard types for transmission from primary and standby.

A snapshot standby database can be created and used as a test database to test new functionality and new releases. When it is converted back into a physical standby database, all the updates made to the primary database are applied to the snapshot standby and changes made to the snapshot are discarded.

3.3.4.2 DataGuard Protection Modes

To determine the appropriate data protection mode, enterprises need to weigh their business requirements for data protection against user demands for system response time. The following table outlines the suitability of each mode from a risk of data loss perspective. Data loss has a trade-off with performance, however, so the solution with the least data loss will have the greatest impact on the day-to-day performance of the application database.

Mode	Impact on data loss	Transmisson
Maximum Protection	Zero data loss	SYNC
Maximum Availability	Zero data loss – assuming that prior to the failure there was no disruption in synchronous communication while the primary database was committing transactions	SYNC
Maximum Performance	Minimal data loss – as little as a few seconds depending upon network bandwidth	ASync

3.3.4.3 Data Guard Corruption Protection

Data Guard automatically validates the operations before applying them to the duplicate database, which eliminates transmission of most file corruption errors to the standby database. In addition, Data Guard may be configured for a delay in applying changes to the standby database, with no impact on operations of the primary database, to protect the standby database from error-prone operations. For example, if a script is run against the primary database, the delay may be configured to provide time for the script to be run and tested before the change is propagated to the standby database.

3.3.4.4 Oracle Active Data Guard

Like the RAC solution, Oracle Active Data Guard provides additional features that allow access to the Data Guard standby database for queries and testing and such while the Data Guard primary is running. Also like RAC, it requires a separate Oracle license.

One of the most useful features included with Oracle Active Data Guard is the automatic repair of corrupt data blocks. A physical standby database operating in real-time query mode can also be used to repair corrupt data blocks in a primary database. If possible, any corrupt data block encountered when a primary database is accessed is automatically replaced with an uncorrupted copy of that block from a physical standby database operating in real-time query mode.

3.4 Site Failure

Enterprises need to protect their critical data and applications against events that can take an entire data center offline. Natural disasters, power outages, and communications outages are all examples of site failures, by making a datacenter completely unavailable. Frequently updated and tested local and remote backups are good, but restoring them in a site-wide disaster can take more time than the enterprise can afford and the backups may not contain the most up-to-date versions of data. For that reason, enterprises often keep one or more duplicate copies of the production database in physically disparate data centers.

In preparation for site failure, a secondary site may be set up and maintained that can provide disaster recovery relief in the event of a principal site failure. The secondary site may be a 'satellite' SunGuide center, with regular operating hours, that operates against a SunGuide installation at another site (the primary site), but is capable of operating independently of the primary site as well. Or the entire secondary site could be intended only for use for testing or in an emergency. We will use the term 'backup' for these locations.

This document does not specifically address cases where there are multiple satellites that interact with the same principal location, although some of the concepts may be extended to meet the needs of that configuration. It also does not comprehensively address the needs of sites that actively share the same devices or locations while concurrently operating independently on a regular basis, although some solutions for that configuration may be discussed or addressed in this document.

3.4.1 'Cold' Backup Site

SunGuide may be deployed to a secondary site solely for the purpose of providing disaster recovery services for the principal site. The system can be configured and maintained to allow for operations to be transferred to that site in the event of a principal site failure, but contain a unique database that is not kept automatically synchronized with the database at the principal site.

In the event of a principal site failure, operations would be transferred to the backup site. While this option would provide high operational availability, it would not provide historical archive consistency. Any data that is collected by the SunGuide application during operation at the backup site would be unavailable to the principal site and vice-versa.

This option allows for the complete isolation of the secondary site from changes made at the primary site, so there is a degree of safety in this solution. It also allows for the same secondary site to serve as backup for multiple principal sites, because the backup database can be configured as an aggregate of the data in each of the principal sites. But it also requires that all application maintenance, including database changes, be performed in duplicate....once at a primary site and then again on the backup site. If system administrators are incautious and the

solution is not regularly tested, they may find that the secondary site is not fully functional when it is needed. Although any issues can be rectified, this may extend system downtime and impact operational availability.

3.4.2 Virtual Servers

Virtual servers may be distributed across multiple storage arrays in disparate locations. In this configuration, the failure of one site would cause all virtual servers to failover to the second site with mirrored storage, where all services can be started up with minimal downtime.

3.4.3 Oracle Data Guard

Oracle Data Guard can be used, not only as a recovery mechanism for database/storage failure, but also as a tool for keeping a primary and secondary site synchronized and allowing switching between the two when there is a failure. Consider the following figure:

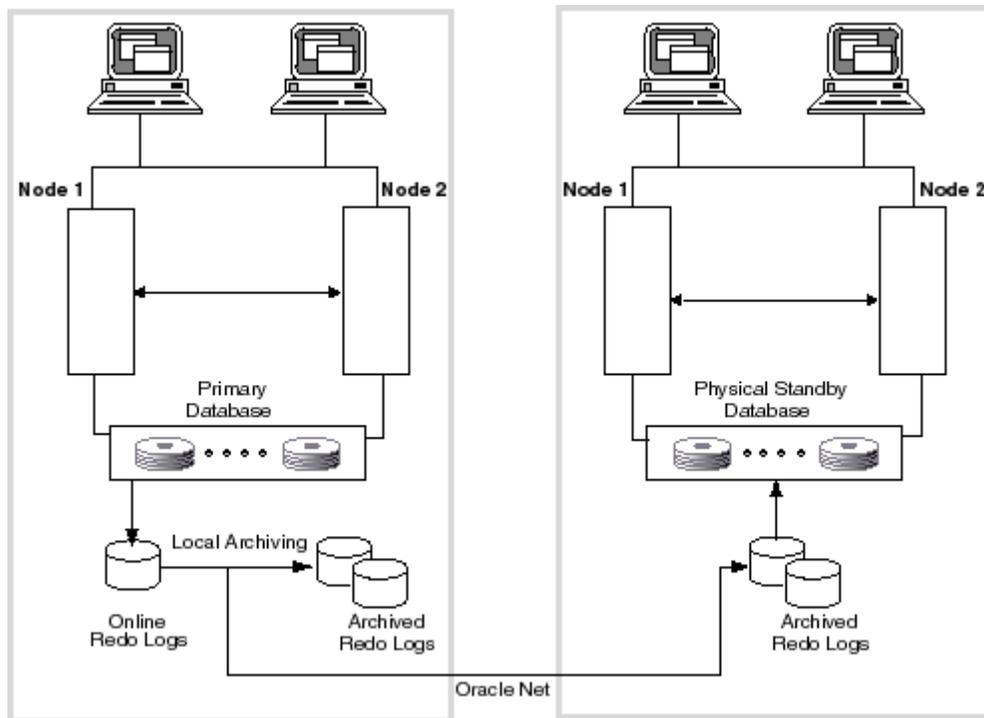


Figure 3.4 – Data Guard Replication

Note that there are not just two databases involved, but two complete SunGuide installations. Should there be a complete site failure at the primary site, then the entirety of SunGuide operations can be transferred to the secondary site.

Note that even using this option, there is some duplication of system maintenance that must take place. Upgrades and configuration file updates must be performed manually on both primary and secondary sites.

3.5 Communications Failure

The contingencies put into place for server and data failure should be sufficient to handle most, if not all, of the downtime scenarios caused by communications failure. For instance, in the case of communications failure to or from a particular server, the downtime may be handled in much the same way as if the server itself failed.

4. Planned Downtime

Typically, planned downtime is a result of maintenance or necessary activities that are disruptive to system operation. These activities include software and hardware upgrades and systemic changes.

4.1 System Changes

This section discusses system changes like database system upgrades and operating system upgrades.

4.1.1 Database Storage System Changes

At times, an Oracle database may ‘outgrow’ its dedicated storage, or there may be another need to migrate it to a different storage location or server. Moving database storage can require quite a bit of downtime. These solutions try to address that downtime.

4.1.1.1 Oracle Streams

A new database with appropriate storage allocation may be created as needed. Oracle Streams may be used to synchronize the new database with the existing active database. When the two databases are in sync, then operations can be transferred to the new database with appropriate storage changes.

4.1.1.2 ASM Online Reconfiguration

Oracle with ASM supports dynamic online system reconfiguration for many components of your Oracle hardware stack. Many configuration changes may be made to a database while the database is up and available. ASM even supports some changes to the file storage configuration of the database. When disks are added or removed from an ASM Diskgroup, Oracle automatically rebalances the data across the new storage configuration while the storage, database, and application remain online. In addition, with RAC, administrators can dynamically add and remove clustered nodes without any disruption to the database or the application.

4.1.2 Application Server Changes

Windows servers require regular maintenance. The software that runs on the servers may require upgrades or patches, or some component of the server may be malfunctioning. These solutions attempt to minimize downtime caused by server maintenance and changes.

4.1.2.1 Microsoft Windows Clustering

SunGuide application services can be migrated to another node while maintenance is being performed on the application server that usually hosts those services. The node is considered

temporarily unavailable. When maintenance on the node is complete, then services will be returned to their original host.

4.1.2.2 Virtual Servers

New duplicate virtual servers may be created, maintenance performed on the virtual servers while 'offline', and then the servers can be placed online with a minimal amount of effort and downtime.

4.1.3 Database Server Changes

As with any other server, the SunGuide database servers require regular maintenance as well. The software that runs on the servers may require upgrades or patches, or some component of the server may be malfunctioning. These solutions attempt to minimize downtime caused by database server maintenance and changes.

4.1.3.1 Oracle Online Patching

Beginning with Oracle Database 11g there is support for online patching for some qualified interim patches. Online patching, which is integrated with OPatch, provides the ability to patch the processes in an Oracle instance without bringing the instance down. Each process associated with the instance checks for patched code at a safe execution point, and then copies the code into its process space.

Oracle also supports the application of patches to the nodes of a Real Application Cluster (RAC) system in a rolling fashion, keeping the database available throughout the patching process. To perform the rolling upgrade, one of the instances is quiesced and patched while the other instance(s) in the cluster continue to service the end users. This process repeats until all instances are patched.

4.1.3.2 Oracle Fail Safe

If the software or system change doesn't require that it be done on the server actively hosting the Oracle database, then the change can be applied to an inactive node in the Oracle Fail Safe cluster first, then the database may be moved to the inactive node and the active node can be modified.

4.1.3.3 Oracle Data Guard

Oracle's Data Guard provides for rolling Oracle software upgrades, operating system upgrades, and other server software changes needed on the active Oracle database host server. Utilizing Oracle's Data Guard SQL Apply technology, administrators can apply database patchsets, major release upgrades, and cluster upgrades with near-zero downtime to the end users. The process begins configuring Data Guard to convert the physical standby to a logical standby and get it synchronized with the production database. Once the Data Guard configuration is complete, the administrator will pause the synchronization and all redo data will be queued. The standby database is upgraded, brought back online, and Data Guard is reactivated.

All queued redo data will be propagated to and applied on the standby to ensure no data loss occurs between the two databases. The standby and production databases can remain in this mixed mode until testing on the logical standby database confirms that the upgrade completed

successfully. At this point, the switchover can occur resulting in a database role reversal – the standby database is now servicing the production workload and the production database is ready to be upgraded, following a process symmetric to the one described. Finally, a second switchover can be initiated and the original production system resumes accepting production traffic.

4.1.3.4 Oracle Streams

Also available for upgrading or changing the active Oracle database host server, Oracle Streams is similar to using Data Guard in SQL Apply mode. Oracle Streams may be used to set up a temporary replica of the Oracle database. This logical copy configured and maintained using Oracle Streams is called a replica, rather than a standby, because it provides many capabilities that are beyond the scope of a normal definition of a standby database.

A database or system upgrade can then be performed on the replica database while the original primary database is available online. Oracle Streams is used to apply the changes made at the primary database to the replica database, and the replica is made available for SunGuide operations.

4.2 Application/Data Changes

This includes application upgrades and configuration changes that normally require system downtime. In particular, the SunGuide application upgrade is a regularly occurring event that can cause downtime.

4.2.1 Satellite or Backup Site

If a secondary ‘satellite’ or ‘backup’ site is available, all operations can be transferred temporarily to the satellite while system maintenance is being performed at the principal site.

4.2.2 Oracle Data Guard

Oracle Data Guard can be used to minimize system downtime caused by the database upgrade portion of a SunGuide application upgrade. The database scripts may be run against the standby database, and then a switchover performed so that the standby database is active. At this point, the primary database may be upgraded.

Depending on the other solutions chosen and available, this could radically improve system availability.

4.2.3 Microsoft Clustering

At this time, cluster SunGuide application server upgrades may be managed similar to application server failures by migrating services off of nodes and upgrading one node at a time. A copy of the configuration file is then upgraded and ready to be switched with the live configuration file. When enough servers have been upgraded to host all of the services temporarily, then the services can be moved to the upgraded nodes, the config file switched out, and the database upgraded. If Oracle Data Guard or a similar solution is deployed, then the database can simply be switched over, and operations will continue within minutes. If not, then downtime is minimized to only the time it takes to run the application upgrade scripts against the SunGuide database.

4.2.4 Online Table Redefinition

Online Table Redefinition is a feature of Oracle that allows changes to a table structure while continuing to support an online production system. It uses the DBMS_REDEFINITION package. Administrators using this feature enable end users to access the original table, including insert/update/delete operations, while the maintenance process modifies an interim copy of the table. The interim table is routinely synchronized with the original table and once the maintenance procedures are complete, the administrator performs the final synchronization and activates the newly structured table. This may be a feature worth looking into to include in future database upgrade scripts for the SunGuide application.

4.2.5 Oracle Edition-based Redefinition

Oracle Edition-based Redefinition is a feature that may also prove useful for doing online application upgrades in the future. Using this feature, code changes are installed in a new Edition in the database. The data changes to the database are made safely, writing only to new columns or new tables not seen by the old Edition. A trigger propagates the data changes made in the old edition into the new edition's structure. This feature would require application changes that would introduce an editioning view in front of every table and changing current table and view access to access those editions, and then editions-enabling the intended users.

5. Examples

This section provides some real-life examples of some of the configurations discussed in this document.

5.1 Overview

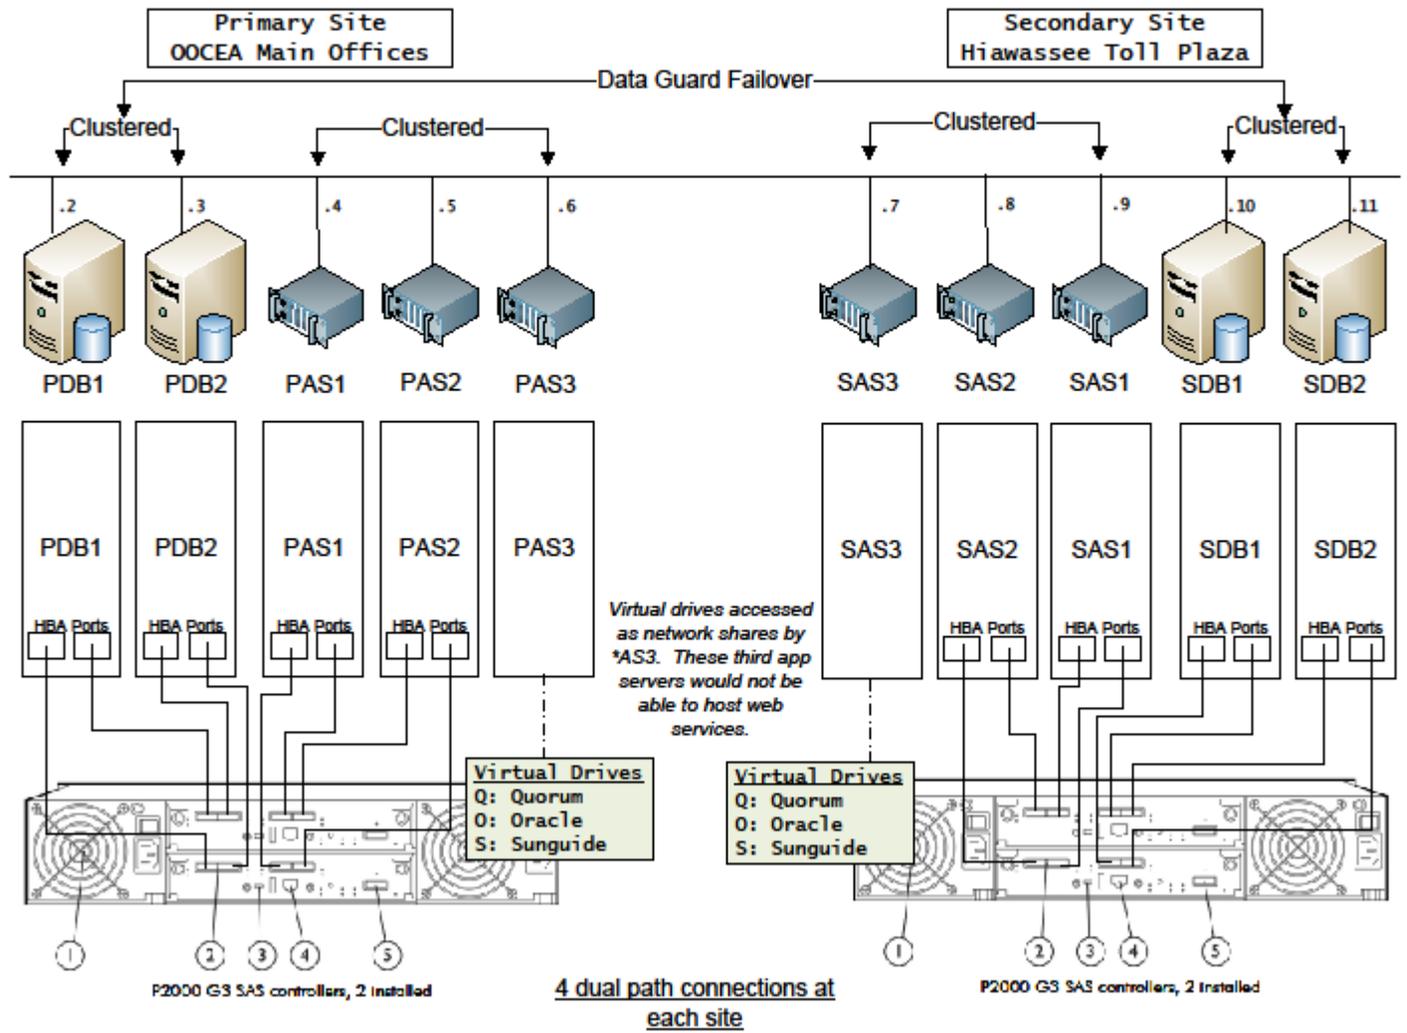
TBD. To include a table that discusses the pros and cons in the context of cost, time, maintenance, maximum downtime, etc. of each of the examples.

5.2 Olando-Orange County Expressway Authority (OOCEA)

The OOCEA chose to deploy a solution for the SunGuide application that utilizes the Microsoft Clustering, Oracle Fail Safe or Oracle Real Application Clusters (RAC), and Oracle Data Guard software to present a redundant and highly available system across multiple sites with separate hardware and storage.

OOCEA has SunGuide installed at two physically disparate sites. The primary site is located in the OOCEA Headquarters building and the secondary site is to be located at the Hiawasse Toll Plaza location. Below is a diagram that depicts the OOCEA deployment configuration.

Title:	OOCEA Server Configuration
Drawn by:	SWRI
Date:	4 Nov 2010



5.2.1 Application Server Failure

Microsoft Windows Clustering is used to minimize downtime in the event of an application server failure.

5.2.2 Database Server Failure

Oracle Fail Safe and Microsoft Windows Clustering are used to minimize downtime in the event of a database server failure at both sites.

5.2.3 Data Failure

Oracle backups, both logical and physical, are created on a regular basis. Oracle Flashback may be used if necessary to roll back the database to a prior point in time. A SAN is used for file storage at both principal and secondary sites. In addition, Data Guard is used to provide a synchronized standby copy of the primary database at the secondary site. Oracle validates the operations transmitted from the primary database and eliminates most opportunities for database corruption to be transferred from site to site.

5.2.4 Site Failure

Oracle Data Guard is used to provide a synchronized standby copy of the primary database at the secondary site. All operations may be failed over to a fully functional, up-to-date, SunGuide system at the Hiawassee location. A separate set of application servers are clustered at the secondary site and ready to be used. Changes made to the SunGuide application services must be made at both locations.

5.2.5 Communications

The OOCEA SunGuide system is set up with separate domains at each SunGuide site. The servers at the principal site are joined to a principal domain, and the servers at the secondary site are joined to a secondary domain. Each domain has a unique set of sub-net IP addresses that are valid for that domain. There is a trust relationship established between the two domains.

This network configuration limits their failover capabilities somewhat, because the application servers in the primary domain cannot access the database in the secondary domain without reconfiguring their

The OOCEA publishes data through Center to Center to both District 5 RTMC and FL-ATIS. Therefore, network communications must be open from both the primary SunGuide system and the secondary SunGuide system to District 5 and FL-ATIS systems.

5.2.6 System Changes

System changes may be handled by updating the secondary site, then transferring operations to the secondary site while the principal site is being updated.

5.2.7 Application/Data Changes

Application changes may be handled by updating the secondary site, then transferring operations to the secondary site while the principal site is being updated.

5.3 TIMSO

At this time, District 4 has implemented this solution. They have two production SunGuide installations, the Broward RTMC and the Palm Beach Vista Center TMC, which have redundant hardware and also utilize Microsoft Windows Clustering and Oracle Fail Safe for server failure scenarios. Each TMC has a distinct SunGuide system, with varying device and location configuration, and individual system configurations.

5.3.1 Application Server Failure

Microsoft Windows Clustering is used to minimize downtime in the event of an application server failure.

5.3.2 Database Server Failure

Oracle Fail Safe and Microsoft Windows Clustering are used to minimize downtime in the event of a database server failure at the primary site.

5.3.3 Data Failure

Redundant hardware is used at the principal site.

5.3.4 Site Failure

A cold backup disaster recovery site has been created in Fort Pierce, FL that is called the Traffic Incident Management Support Office (TIMSO). This location provides a recovery solution for both Broward RTMC and Palm Beach Vista Center TMC. The TIMSO SunGuide system is expected to retain a distinct and unique database that will include devices for all five counties of the two existing District Four ITS deployments. The goal in creating the TIMSO site is to have a fully operational SunGuide that can be used to control any device in the event of a catastrophic failure at any other site. There is no redundancy built into this secondary site, but it is only expected to be used for short periods of time.

5.3.5 Communications

The TIMSO deployment is not currently set up to interact with FL-ATIS or any other district through Center to Center. It is independent and isolated from the primary sites that it supports.

5.3.6 System Changes

Microsoft Clustering may be leveraged to minimize downtime during system upgrades and other maintenance.

5.3.7 Application/Data Changes

Microsoft Clustering may be leveraged to minimize downtime during system upgrades and other maintenance.

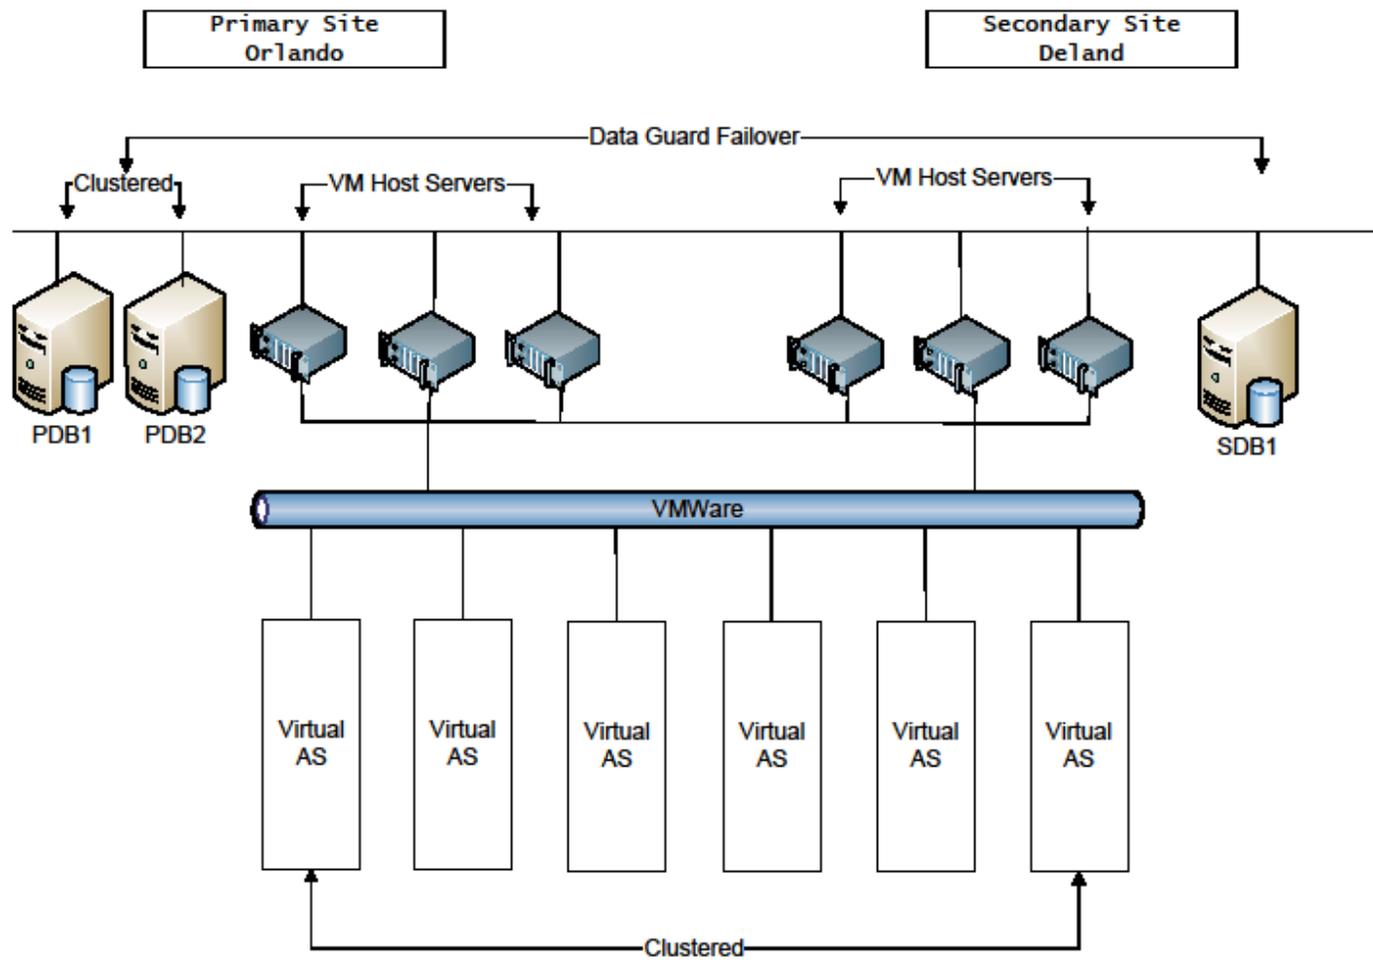
5.4 District 5 RTMC

District 5 chose to deploy a solution for the SunGuide application that utilizes the Microsoft Clustering, Oracle Fail Safe, VMWare, and Oracle Data Guard software to present a redundant

and highly available system across multiple sites with separate hardware and storage that is connected and mirrored across the network.

The District 5 SunGuide system is configured so that all of the SunGuide application servers are hosted by VMWare on host servers in physically disparate locations, Orlando and Deland. For the purpose of this document, Orlando will be considered the primary site and Deland the secondary site. The SunGuide database server is hosted on a two node cluster of hardware servers in Orland, and replicated in Deland on a single hardware server. Below is a diagram that depicts the District 5 deployment configuration.

Title: District 5 Server Configuration
Drawn by: SWRI
Date: 3 Nov 2011



5.4.1

5.4.2 Application Server Failure

Microsoft Windows Clustering is used to minimize downtime in the event of an application server failure. In addition, VMWare is used to help maximize system availability in the event of a server hardware failure.

5.4.3 Database Server Failure

Oracle Fail Safe and Microsoft Windows Clustering are used to minimize downtime in the event of a database server failure at both sites.

5.4.4 Data Failure

Oracle Flashback may be used if necessary to roll back the database to a prior point in time. A SAN is used for file storage at both principal and secondary sites. In addition, Data Guard is used to provide a synchronized standby copy of the primary database at the secondary site. Oracle validates the operations transmitted from the primary database and eliminates most opportunities for database corruption to be transferred from site to site.

5.4.5 Site Failure

Oracle Data Guard is used to provide a synchronized standby copy of the primary database at the secondary site. All operations may be failed over to a fully functional, up-to-date, SunGuide system at the Deland location. The servers actively hosting SunGuide application services may be failed over to the secondary site through the VMWare software and seamlessly operate against the secondary database.

5.4.6 Communications

District 5 operates a single domain, single subnet network between Orlando and Deland. Because of this, the virtual IP addresses may be transferred to the Deland servers with no disruption of operation.

District 5 publishes data through Center to Center to both OOCEA and FL-ATIS. Therefore, network communications must remain open from the District 5 network to these entities for this capability to remain online.

5.4.7 System Changes

Microsoft Clustering and virtual clustering may both be leveraged to minimize downtime during system upgrades and other maintenance.

5.4.8 Application/Data Changes

Microsoft Clustering and virtual clustering may both be leveraged to minimize downtime during system upgrades and other maintenance.

6. Other

While we have presented three complete deployment solutions in this document, there are many more similar configurations that can be created using the components of the solutions discussed

here. As discussed before, it is recommended that system availability requirements be determined and then an analysis of those requirements, data reporting requirements, system uptime thresholds, budget constraints, etc. be performed to determine the ideal solution for any SunGuide deployment.

7. Future Considerations

Rolling Application Upgrades using Edition-based Redefinition should be evaluated and cost estimate performed for FDOT.