An Oracle Technical White Paper
September 2011

The Art of Data Replication

ORACLE

Introduction

This technical paper describes how to implement an efficient data replication process for your business. It begins by explaining the different kinds of risk a natural disaster can pose to your business. It then explains the technical concepts you need to understand before you can select the best solution. Finally, it describes the best way to identify your requirements, select your replication method and technology, and ensure your data's consistency.

A loss of vital business data can cripple a company or even put it out of business. As more customers become aware of the value of their business data, securing that data is being seen as paramount to running a successful business. Keeping copies of the data at a physically separate location is a common approach to safeguarding data from loss through a natural disaster or fire.

Data replication is a technique that can be used to copy data over a computer network to one or more remote locations. However, the use of data replication adds a level of complexity to an IT storage design. Before embarking on designing a data replication-based architecture, it is important to clearly understand terminology and concepts related to replication.

Gathering requirements for the design should start at the business-process level with a review of the contingency strategies and requirements for risk management as defined in the business continuity plan. Limiting the focus to just IT-related requirements presents the risk of defining a solution that does not meet business requirements and creating a false sense of security.

It is also helpful to gather information about the lifecycle of the data that will be replicated. Replicating all the data that has been stored over the last ten years will be unnecessarily costly if only data from the last year is relevant to the day-to-day operation of the business.

Understanding all aspects of the replication process will help you successfully determine design criteria and select the type of replication that best meets your requirements.

Understanding Concepts Related to Mitigating Business Risk

The terms *business continuity* and *disaster recovery* are often used interchangeably in the context of data replication and casually inserted into advertising and marketing material for storage products. However, each term has a distinct meaning, and understanding them is essential when translating requirements into data replication solutions. This section explains these terms and other concepts related to mitigating business risk.

Business Continuity

Business continuity refers to what needs to be done by an organization to ensure that essential business functions can continue during and after a disaster, including preventing the interruption of mission-critical services and reestablishing full functioning as quickly as possible. Risks to the business must be identified and risk mitigation strategies must be put in place to prevent business processes disruption.

Prevention and noninterruption are key objectives for business continuity, so the focus is on business continuation rather than business (or disaster) recovery. The first step to a good business continuity strategy is the creation of a business continuity plan supported by a commitment of assigned resources and budgeted funds to both develop and maintain the plan.

Disaster Recovery

Disaster recovery refers to activities and programs designed to get a business back up and running after an unexpected interruption of services. A disaster recovery plan describes the response to be made to a disruptive event to restore crucial business functions. The recovery process extends beyond the IT infrastructure to include the immediate response, restoration or relocation of assets such as servers and storage systems, reestablishment of asset functionality, data recovery and synchronization, and finally, restoration of business functions.

For disaster recovery, the focus is on how to react to an unexpected disruption, in contrast to business continuity, which focuses on prevention of disruptions and continuation of essential business functions. A disaster recovery solution is typically concerned with recovering the functionality of assets after a disruptive event, while a business continuity strategy looks more broadly at how business processes as well as assets can be maintained through a potentially disruptive event. Thus, a disaster recovery plan is often a part of an overall business continuity strategy.

Availability Versus Reliability

In IT infrastructure design, availability is often given more emphasis than the other two RAS components: reliability and serviceability. However, a focus on availability can create a false sense of security with respect to data integrity. Guaranteed access to data 99.999 percent of the time (the widely accepted "five 9s uptime" requirement) does not guarantee that data is consistent or uncorrupted.

Data loss due to human error or to an application bug quickly spreads through data sets when data sets are replicated between sites. Therefore, when creating a data storage architecture, it is important to consider aspects of data reliability, such as guaranteeing data consistency, along with strategies for mitigating the risk of data corruption.

Planning for Business Continuity

The business continuity planning process is similar to that of any project plan methodology. The following are typical steps for developing a business continuity plan:

- *Project Initiation Phase* – Includes defining objectives and assumptions
- *Functional Requirements Phase* – Includes gathering facts, identifying alternatives, and eliciting decisions by management
- *Design and Development Phase* – Includes designing the plan
- *Implementation Phase* – Includes rolling out the plan
- *Testing and Exercising Phase* – Includes conducting a post-implementation review of the plan
- *Maintenance and Updating Phase* – Includes keeping the plan up to date and aligned to business processes changes, changes in roles and responsibilities, asset changes, and other ongoing adjustments

More information about business continuity planning can be found at the Website of the global certification and training organization DRI International.

A substantial effort is required, in terms of budget and resources, to complete all the steps for creating and maintaining a business continuity plan—a process that is well beyond the scope of a typical IT consultant and or sales person.

Oracle offers assistance in completing the Functional Requirements Phase, including conducting a business impact analysis and a risk assessment study. These activities cover a broader scope than just the IT infrastructure. The business impact analysis identifies requirements for recovery point objectives (RPO) and recovery time objectives (RTO) for essential business processes in a company. (For more information about RPO and RTO, see the sections "Recovery Point Objective" and "Recovery Time Objective" below.)

Caution: A request for an IT infrastructure disaster recovery solution should always be a part of a broader disaster recovery planning effort. For example, if a business continuity plan is not in place, it will be difficult to determine realistic values for RPO and RTO. Engagement with higher-level management outside the IT department might be required to obtain business continuity information. A commitment to a realistic budget for the IT infrastructure portion of the disaster recovery solution is also desirable.

The following two examples illustrate the need to look at a disaster recovery solution from a business continuity perspective.

Case 1

Although many people move to Florida to enjoy its warm climate, Florida is also, unfortunately, the frequent landing point for hurricanes moving in from the Caribbean. In this case, a hurricane caused severe flooding as a secondary effect. A company had invested in creating a disaster recovery site well out of the flood zone, with well-defined IT procedures and equipment in place. However, the company could not get its site up and running because the employees who were to supposed execute the disaster recovery activities were dealing with the impact of the flooding on their own homes and taking care of their families. The company had focused on setting up an IT site as part of its business continuity plan, but it neglected to take non-IT resourcing factors into account in the disaster recovery plan.

Issues faced by other companies during that event included lack of access to a fuel supply for power generators, communication infrastructure failures interfering with coordination of disaster recovery activities, and much more.

Case 2

As part of a business continuity assessment project, a business continuity consultant interviewed operational department heads to identify risks and threats to their business operation and measures that could be taken to lower the risks. One concern identified was how outstanding invoices would be tracked if, for example, a major building was destroyed by fire.

In this case, the IT department proposed that a synchronous replication of the database be implemented at a disaster recovery site with the ability to print invoices and payment reminders at the site. However, when the consultant asked the financial department for its requirements, the department proposed a much simpler solution that included a printed list of outstanding customer invoices and a stack of forms that could be manually processed with data from the printed list. This temporary solution was sufficient to ensure that money would continue to come into the business for at least the two weeks following an event. This case shows that to avoid disconnects between the IT department and the departments running the business processes, it is important to verify IT requirements against requirements at the business level.

When developing a business continuity plan, it is important to consider mitigation strategies that address conditions resulting from a disruptive incident or disaster that could impact staff availability; access to your company premises; flow of supplies; access to electricity, water, and gas; security of information. Figure 1 shows several ways the effects of flooding could affect the efficacy of a business continuity plan.

Figure 1. Flooding in the UK several years ago caused conditions that could impact business continuity.

Developing an Information Lifecycle Plan and Policies

An information lifecycle policy and plan are important complements to a business continuity plan. The information lifecycle plan classifies data according to its value to the business, ranging from noncritical to mission critical, and it defines a lifecycle policy for each classification.

When developing an information lifecycle management (ILM) plan, it is important to identify regulatory compliance requirements, because they can affect how particular data is classified. The process of developing a plan can also reveal that some data might not need to be kept in primary storage and replicated to a secondary site. Figure 2 shows how the cost to implement a disaster recovery solution is related to how data is classified and the type of solution required to protect each data classification.

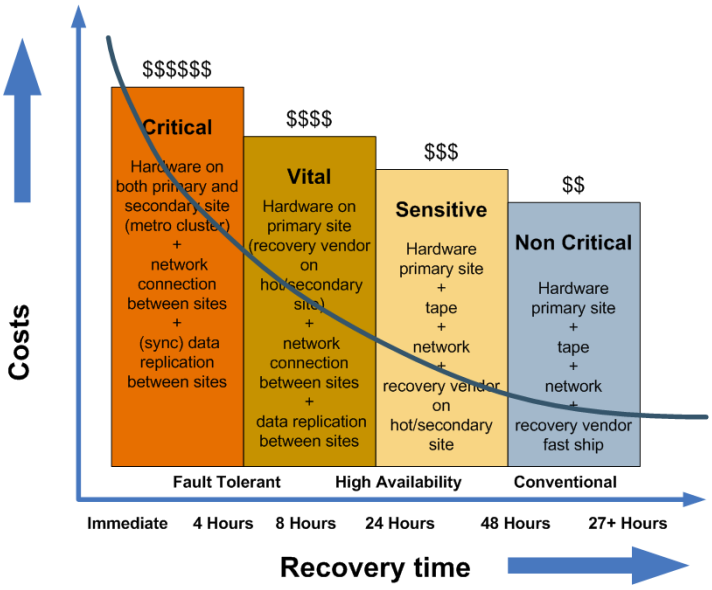

Figure 2. Data classification and the type of disaster recovery solution impact the cost to implement and maintain a disaster recovery solution.

Collecting Requirements for a Data Storage Architecture

Too often, IT architecture is designed around servers, with storage bolted on. An approach to an optimally designed storage infrastructure starts with the development of a logical diagram that shows all data instances, the dependencies between them, and which applications use each instance.

Once the logical diagram has been reviewed and accepted, requirements for each of the data instances (such as size, expected growth, I/O access profile, and importance to the business) can be gathered. Business continuity RPO and RTO requirements should be identified for each data instance and it should be confirmed that they match the RPO and RTO requirements of the business processes. (For more information about RPO and RTO, see the sections "Recovery Point Objective" and "Recovery Time Objective" below.)

Each data instance should be assessed to determine the impact of corruption or permanent loss. From this information, redundancy and recovery requirements can then be derived. These requirements determine the RAID level used (RAID 1 versus RAID 5, for example), how recovery from corruption of data is accomplished, whether data is restored from backups or from snapshots, and whether a fully copied instance of the data is to be kept at another site.

The logical diagram also clarifies whether storage space can be shared by data instances with the same requirements or whether provisions need to be made for different applications or data instances. Once these requirements are clear, you are ready to design the storage architecture.

Understanding Concepts Related to Data Replication

A business continuity strategy often includes the replication of business data to an off-site location. The distance between sites and the method of data replication are determined based on the results of the business impact analysis and risk assessment activities described in the previous section. The goal is to find the right "insurance policy" to address risks that are identified as potential threats to the integrity of business processes. For example, if data is to be replicated to remote sites over long-distance connections, a Fibre Channel-based SAN might be the preferred option for a storage architecture.

Brief explanations of key concepts and the methodologies available for data replication are described in the following sections. Additional information can be found in the references listed in the "Resources" section.

Methods of Data Replication

Replication of data over distance can be implemented using several different technologies. These technologies are described in this section along with the advantages and disadvantages of each.

Host-based Replication

In host-based replication, the replication process is driven by a host process inserted into the data I/O driver stack.
Advantages: Host-based replication is storage agnostic.
Disadvantages: Replication is limited to a single pair of servers or, at best, several servers running the same version of operating system. An additional license might be required for each server added.

SAN Network-based Replication

In SAN network-based replication, the replication process is incorporated into the data interconnect infrastructure, which, in most cases, is a SAN.
Advantages: SAN network-based replication is storage agnostic and server-architecture agnostic. When combined with mirroring or snapshot functionality, SAN network-based replication can be easily integrated into and managed in an application/database environment.
Disadvantages: SAN network-based replication adds latency to the primary data path. When an application writes data to primary storage, the data is also written to the remote site, adding delay to the write operation. Thus, the SAN infrastructure does not act as a transparent layer between the server and the primary storage, which can make diagnosing faults more difficult.

Storage Subsystem-based Replication

In storage subsystem-based replication, the replication process is part of the array management control and virtualization software or firmware.
Advantages: Storage subsystem-based replication is server-architecture agnostic.
Disadvantage: A license is required for each storage subsystem and the cost is often based on the subsystem capacity.

Types of Replication

Replication can be performed in either synchronous or asynchronous mode. Synchronous replication guarantees continuous data integrity during the replication process with no extra risk of data loss. However, the impact on the performance of the application can be significant and is highly dependent on the distance between the sites. Asynchronous replication overcomes the performance limitations of synchronous replication, but some data loss must be accepted. The primary factors influencing the amount of data loss are the rate of change (ROC) of data on the primary site, the link speed, and the distance between sites. (For more information about ROC, see the "Rate of Change" section below.)

Synchronous Replication

When synchronous replication is used, an update made to a data volume at the primary site is synchronously replicated to a data volume at a secondary site. This guarantees that the secondary site has an identical copy of the data at all times.

The disadvantage of synchronous replication is that a write I/O operation acknowledgement is sent to the application only after the write I/O operation is acknowledged by the storage subsystem at both the primary and the secondary site. Before responding to the application, the storage subsystem must wait for the secondary subsystem I/O process to complete, resulting in an increased response time to the application. Thus, performance with synchronous replication is highly impacted by factors such as link latency and link bandwidth. Deployment is only practical when the secondary site is located close to the primary site.

When evaluating the use of synchronous replication, an important consideration is the behavior of the storage subsystem when the connection between the primary and secondary subsystem is temporarily disrupted. For more details, see the "Recovery Point Objective" section below.

Synchronous replication does not provide protection against data corruption and loss of data due to human errors. Snapshot technology must be used with synchronous replication to provide full protection against both loss of access to data (protected by replication) and loss of data due to data corruption (protected by creating snapshots).

Asynchronous Replication

In an asynchronous mode of operation, I/O operations are written to the primary storage system and then sent to one or more remote storage systems at some later point in time. Due to the time lag, data on the remote systems is not always an exact mirror of the data at the primary site. This mode is ideal for disk-to-disk backup or taking a snapshot of data for offline processes, such as testing or business planning. The time lag enables data to be replicated over lower-bandwidth networks, but it does not provide the same level of protection as synchronous replication.

Asynchronous replication is less sensitive to distance and link transmission speeds. However, because replication might be delayed, data can be lost if a communication failure occurs followed by a primary site outage.

Considerations when sizing link requirements and configuring the storage subsystem include the following:

- Incoming ROC (for more information about ROC, see the "Rate of Change" section below)
- Speed of the replication process in the storage subsystem (how quickly data can be pushed out to the replication link)
- Line speed and line latency
- Size of the replication buffer (queue space) in the storage subsystem; the buffer has to be large enough to cope with peak ROCs

Rate of Change

The rate of change (ROC) is the amount of data that is changed over time on the data volumes that are to be replicated to a second location. ROC can be expressed as a peak value or as an average over a period, such as a day.

ROC peak values are used for sizing links used for synchronous replication, while ROC average values are used for sizing links used for asynchronous replication. For asynchronous replication, it is also important to verify that the storage subsystem is able to buffer peak ROCs in a queue. If not, the storage subsystem will likely throttle down the write speed to the server (and, thus, to the application).

The ROC is used with RPO requirements when determining the speed of a replication link. For more details, see the "Recovery Point Objective" section below.

It is risky to use only server or application ROC information to determine the required inter-site *link* speed for replication. A one-to-one relationship might not always exist between the amount of data written by the application to the storage subsystem and the amount data replicated from the primary storage subsystem to the storage subsystem at the secondary site. Implementations of storage subsystem replication nearly always use some type of mechanism to track data blocks that have been updated on the primary volume (that is, dirty blocks). So, in addition to active synchronous replication, data at the secondary site is being updated by the dirty-block mechanism. The resolution of the mechanism determines the minimum amount of data that is replicated when one disk sector is changed on the primary storage subsystem.

For example, assume that the storage subsystem uses a bitmap for dirty-block marking, which uses 1 MB per LUN, and the LUN is 2 TB in size. This means that when a change is made to data in primary storage, the minimum size of the data block that will be tracked by the dirty-block marking mechanism is 2 TB divided by 1 MB, or 2 MB. If a true random workload is generating 512-byte block updates, in the worst case, a replication load would be generated that is a factor of 2000 higher than the ROC generated by the application. Although this is an extreme example, it shows the importance of taking into account other considerations besides server and application ROC information when determining the required link speed for a replication implementation.

Recovery Point Objective

The recovery point objective (RPO) is the amount of time that a consistent copy of data at the secondary site can lag behind the current data at the primary site. In other words, this is the maximum acceptable time between sequential backups of data from the primary site to the secondary site. When a business process needs to be restarted, the RPO reflects the point to which data will be rolled back as a result of the recovery process.

The desired RPO is determined by first establishing how much data or how many transactions it is acceptable to lose if a business process is disrupted. The amount of acceptable data loss (in MB), the ROC, and the replication speed can then be used to determine the actual RPO. See the "Selecting the Interconnect Technology" section for more information about the interaction between ROC and RPO.

If the lag time between current data at the primary site and recovered data must be kept to a minimum (small RPO value), synchronous data replication might be required.

Recovery Time Objective

The recovery time objective (RTO) is the amount of time in which essential business processes must be up and running again after a disaster. The IT component of the RTO is the time it takes to recover the most recent consistent copy of data.

In a replication environment, a key concern is the integrity of the replicated data. A data set can be recovered only if the data in the data set is consistent at the time that the disruptive event occurs. The order in which data changes are applied to volumes at the secondary site is an important factor in maintaining the integrity of replicated data. For more details about write order consistency (WOC), see the next section, "Write Order Consistency."

In a tape-based backup environment, an important factor is the time it takes to recover data from tapes. Recovery time can be significantly reduced by keeping an up-to-date copy of data available and online at a separate location.

Write Order Consistency

An important objective of a replication solution is to ensure the integrity of the data replicated to the secondary site. When a site failover occurs and the secondary site becomes active, applications must be able to access a consistent set of data at the secondary site.

Two considerations for data integrity are the consistency of data on each volume and the consistency of data across a group of volumes that form a database. For both, the order in which data changes are applied to the secondary volume(s) is crucial if data integrity is to be maintained during a primary site failover.

For a file system, the integrity of data must be maintained as files are created and deleted. For example, if inode structures are not updated in the same sequence in which they were issued by the host, the file system will be completely corrupted.

In a transactional environment where multiple volumes are used for a single data set, the order in which writes are applied is particularly important. Often replication solutions offer the capability to establish one or more write order consistency (WOC) groups to ensure data consistency across all volumes in a database. The WOC group capability guarantees that all writes to replication sets (pairs of volumes) in a consistency group are applied in the same sequence at the secondary site as the application generates them on the primary site.

For any replication product used in the solution, it is important to determine which replication states are used (such as synchronizing, suspend, or replicating) and in which of these states the volumes are in a WOC state and in which they are not.

Replication products from different vendors and on different platforms behave differently. Product information is often not clear about how WOC is handled. Table 1 shows typical replication states used in replication products and in which of these states WOC might be in question in a particular product. In the table, gray cells indicate states for which you might need to check with the vendor to find out how WOC is handled before you implement the product in your solution.

TABLE 1. WOC IN VARIOUS REPLICATION STATES
(GRAY CELLS INDICATE STATES IN WHICH WOC MIGHT BE HANDLED DIFFERENTLY BY DIFFERENT VENDORS)

REPLICATION MODE	REPLICATING	SUSPEND STATE	SYNCHRONIZING
Synchronous	WOC	Possible WOC when queuing is used. No queuing implies no WOC.	Possible WOC when queuing is used. No queuing implies no WOC.
Asynchronous	Possible WOC when queuing is used. No queuing often implies no WOC.	No WOC unless journaling is used.	No WOC unless synchronizing from a journal.
Asynchronous with WOC group	WOC	No WOC unless journaling is used.	No WOC unless journaling is used.

When a volume at the primary site is not in a replicating state, changes on the volume can be tracked using some method of logging. The simplest method is the use of a bitmap, or scoreboard, to track changed blocks. The disadvantage of this solution is that no time information is kept, so information about the order in which the changes occurred is lost. Thus, changes cannot be applied on the secondary volume in the same time sequence in which they occurred on the primary volume.

Time-based tracking can be implemented using a queuing mechanism or journaling method on the primary side. It is important to check the size of the queue and whether it is fixed or can be specified in the setup. When the queue is full, the replication continues in scoreboard (bitmap) mode and WOC is lost from that point forward. Thus, the optimal queue size is dependent on the ROC and the RPO.

It's important to understand how a component malfunction, such as an array controller failure, that occurs during the synchronization process will impact the integrity of the replication data set. As an example, the volume shown in Figure 3 is being used in a transactional application. While the replication set was in the suspend state, three transactions took place. Now the replication set has moved to the synchronizing state. The scoreboard on which the changes have been recorded is scanned in a logical sequence of block order numbers. The result is that the blocks are applied to the secondary volume of the replication set out of order with respect to time. If the primary site were to go down during this synchronization process, the secondary volume would be corrupted and all data at the secondary site would be lost.

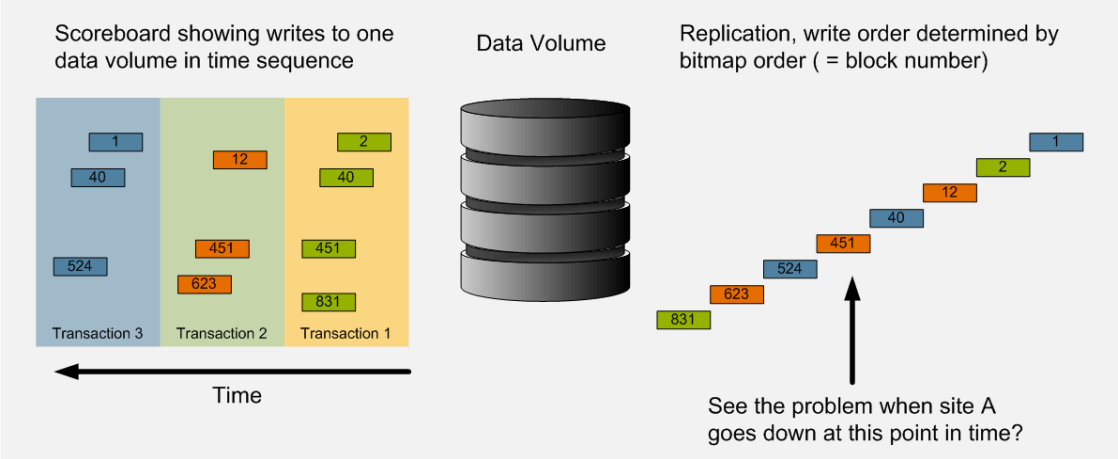

Figure 3. Not maintaining write order consistency in a transactional environment can cause data corruption.

Snapshots can be used in a replication design to ensure the availability of a consistent set of data at all times. For example, if a synchronous replication set using a scoreboard to track changes has been running in suspend mode for some time, the scoreboard might have recorded a number of changed blocks during the period during in which the connection was lost to the secondary site. The creation of a snapshot of the volumes at the secondary site before the replication process is restarted ensures a consistent set of data will be available to fall back to if, for some reason, the primary site goes down during the synchronization process resulting in a corrupted set of data at the secondary site.

Performance-Related Characteristics of Replication Links

Replication solutions primarily use one of two types of physical links: Fibre Channel (FC), which are glass, or IP-based, which are copper. FC links are used primarily for synchronous replication over shorter distances. They offer high-bandwidth, low-latency connections and are used for distances up to 50 km.

Deploying FC links over longer distances is expensive and difficult to implement. To utilize the full capacity of a long-distance FC link, bandwidth is often shared by multiplexing several logical FC communication channels on one physical pair of fibre strands using coarse/conventional wavelength division multiplexing (CWDM) or dense wavelength division multiplexing (DWDM). Multiplexing might be a security concern in some cases.

IP-based links use copper lines and the standard IP network traffic protocol. IP-based links are typically used for asynchronous replication over longer distances, such as intercontinental links between sites. Replication data often travels over public networks. Special switches that convert FC-over-Fibre protocol to FC-over-IP protocol are used as a bridge between internal SAN FC networks and the public IP network.

Line Bandwidth

The line bandwidth (also referred to as line speed) is expressed either in bits per second or bytes per second. Taking protocol overhead into account, a rule of thumb is that ten bits are needed to transmit one data byte. So a 2-gigabit link can transmit a maximum of 200 MB of data per second. However, line speed alone is not an accurate reflection of the maximum replication rate that can be achieved. Other factors must also be taken into account, such a line latency (described in the next section, "Line Latency").

Line Latency

Line latency is related to the propagation speed of the signal over the cable. The longer the cable, the longer it takes for the bits to arrive at the receiving site. Once data is received, an acknowledgement is sent back. The line latency is the time it takes for a data packet to propagate over the link to its destination and for the acknowledgement to arrive back to the sender.

Latency depends on the speed of the transmission medium (for example, copper wire, optical fiber, or radio waves) and delays in transmission by devices along the way (for example, routers and modems). Low latency indicates high network efficiency.

Another indicator of latency is the velocity of propagation (VoP) or velocity factor (VF). The VF characterizes the speed at which an electrical signal passes through a medium. It is the ratio of the signal transmission speed to the speed of light in a vacuum (300,000 km/s) expressed as a number between 0 and 1 (or a percentage).

Table 2 shows some typical velocity factors taken from the *The ARRL Handbook for Radio Communications*. The transmission speed of a signal in a glass fibre link is 66 percent of the speed of light in a vacuum, or 200, 000 km/s.

TABLE 2. TYPICAL VELOCITY FACTORS

VF%	TRANSMISSION LINE
95	Ladder line
82	Twin-lead
79	Coaxial cable (foamed polyethylene dielectric)
66	Coaxial cable (solid polyethylene dielectric)
66	Glass fibre

The transport protocol to send data over an FC cable consists of a number of protocol layers. At the transport layer (FC-2), data is packaged in frames. The distance and the speed of transmission in glass (200,000 km/s) determine the time it takes for a data frame to travel the length of the cable. The latency of the line is the round trip comprising the time to transmit a frame of data and to receive an acknowledgement back. Line latency is often expressed in milliseconds.

Most storage subsystems use SCSI as the upper-layer protocol to exchange data between the initiator and the target using FC-2 as the transport layer. Each SCSI command consists of a command phase and a data phase. Each phase ends with an acknowledgement from the target, so each phase requires that at least one frame of data be sent in each direction at the FC-2 level. A complete execution of a SCSI command requires, at a minimum, two round trips of frames. Therefore, the data latency at the SCSI level is twice the line latency at the FC-2 level.

For example, if the time for a data frame to traverse a 100 km link in one direction is 500 μ, the line latency of the link is 1 ms (one round trip). When the SCSI protocol is used, the time to complete the transmission of a frame of data over the link, or the data latency, is 2 ms.

Latency is closely related to throughput, another fundamental measure of network performance. Latency measures the amount of time between the start of an action and its completion, while throughput is the total number of such actions that occur in a given amount of time. Utilities that are commonly used to measure latency and identify latency issues are `ping` and `traceroute`.

Limitations Imposed by Distance on Replication Links

In an extended SAN configuration, each FC logical layer (FC0 through FC4) of the interconnection (data path) between the application and the physical storage imposes limitations.

Physical Limitations

In the FC physical and transport layers (FC0 through FC2), the degradation of the signal in the cable, the speed of light, and the cable bandwidth are factors that limit the practical distance for a link. With the use of DWDM optical technology, distances up to 300 km can be covered without the need to regenerate the signal along the way. The speed of light in a glass cable is 200,000 km/s. This results in a 5 μs/km delay or 1 ms of latency for every 100 km.

The speed of the transceivers at each end of a fibre cable and the physical cable characteristics determine the transfer speed of the link. However, the transfer speed is not the same as the replication speed that can be achieved over the link. The replication source would need to be able to generate enough I/O operations to fully utilize the available bandwidth. Assuming a 200-MB/s fibre link and an 8-KB SCSI I/O data block, the source would need to generate 200/8 * 1000 = 25,000 I/O operations per second (IOPS) to do this! In addition, the FC nodes would need enough buffer credits to fully utilize the link capacity. (Buffer credits are described in the next section, "Protocol Limitations.")

Protocol Limitations

The FC protocol, like many other protocols, uses a form of flow control to avoid loss of data packets at the receiver side. The FC protocol uses a credit-based algorithm to control the flow of packages. When a session is started between an ingress port and an egress port, a buffer-to-buffer credit count is established. The number of buffer credits is equal to the number of available receiver buffers at the ingress port. To prevent buffer overruns that would result in frames being dropped, the egress port can send only as many frames as the number of available buffer credits before waiting for an acknowledgement from the ingress port.

It is important to have enough buffer credits to be able to utilize the line bandwidth. The longer the distance, the more frames can be in transit on the line. The ideal situation is to have as many buffer credits on the ingress port as the maximum number of frames in transit. The following formula can be used to calculate the number of buffer credits given the distance and transmission speed of the line:

$$BB-C = \frac{d * ls * td}{(10 * fs) + 1}$$

`BB-C` is the number of buffer-to-buffer credits, `d` is the distance of the line (km), `ls` is the line speed (bits/s), `td` is the transmission delay (s/km), `10` is the number of bits per frame, and `fs` is the frame size (bytes).

As described in the previous "Line Latency" section, each execution of a SCSI command contributes a latency of at least two round trips over the link. This latency is independent of the number of buffer credits used on the transport layer. Thus, the bandwidth of a link can be fully utilized only by executing SCSI commands in parallel using multiple I/O threads. This explains why the I/O performance of single-threaded applications over long distances is so sensitive to latency.

Replication over IP Links

For longer distances (> 500 km), replication over an IP WAN is the most cost-effective solution. FC-to-IP routers take care of the conversion between the FC and IP protocol layers.

The IP protocol is fundamentally different from the FC protocol. The FC protocol is a deterministic, lossless transport protocol with low latency, while the TCP/IP protocol is characterized by unpredictable latency due to dynamic routing and no guaranteed delivery of packets. These characteristics impose extra requirements on the firmware in the routers. Features that make up for TCP/IP protocol limitations must be implemented, such as a mechanism for flow control and a way to guarantee a specified quality of service (QoS). Technologies that implement such features should be selected for a design based on a proven track record.

When designing a replication solution that uses the TCP/IP protocol, make sure that any links provided by an external telecommunications provider deliver guaranteed QoS.

Application Limitations

The impact of using an extended SAN on an application depends on the I/O profile of the application. For example, when using synchronous replication, the performance of an application with a single-threaded sequential write pattern will be significantly impacted due to the increase in latency introduced by longer distances. Assume, for example, that an application typically sees a 1-ms response time from the local storage subsystem. When a 100-km, synchronously-replicated remote storage system is used, a 1-ms round trip delay is added. The result is a 100-percent impact on application performance.

At the other end of the spectrum, an application with a random write pattern with multiple threads might quickly reach a 5-ms response time locally. The 1-ms latency due to the distance to the remote storage system is much less significant in this situation.

It is essential to understand the I/O behavior of each application that is involved in a replication project and the impact on performance of latency introduced by distance before designing the solution. Oracle's StorageTek Workload Analysis Tool (SWAT) is an excellent tool to use to determine application I/O characteristics. It can be downloaded at the My Oracle Support Website (access requires support contract).

How It All Fits Together

Now that the concepts and technologies involved and requirements such as RPO and RTO have been explained, it is time to see how this all plays together.

Identifying Requirements

Many factors need to be considered before starting to design a data availability solution based on data replication. The first step is to determine the business requirements and translate them into IT availability requirements as they apply to disaster recovery.

Start by identifying key factors that will impact the design, including the following:

- Physical limitations, such as site distance and the type of transport links available
- Financial limitations, including the available IT budget and costs of components such as links
- Business limitations, particularly the RPO and RTO

Figure 4 shows the key factors to be determined in the initial project phase.

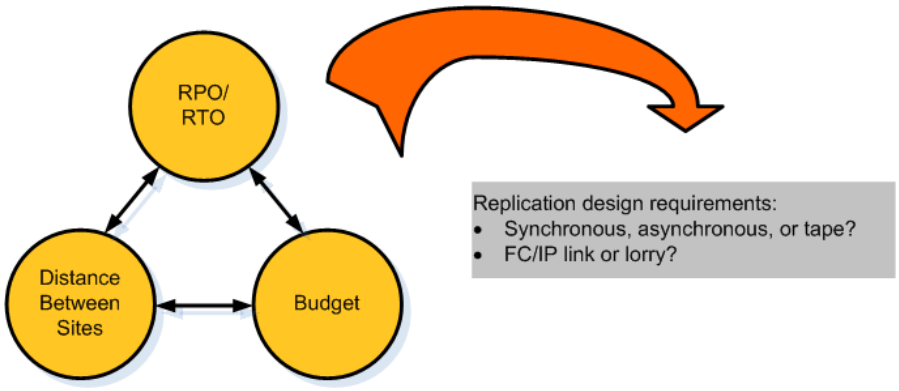

Figure 4. Key factors that impact a replication design should be determined in the initial project phase.

Next, identify the data sets used by applications and gather requirements for each data set. Data is not always stored in a single data repository. Different applications or multiple instances of a single application might use their own data sets. Identify where multiple copies (mirrors) of a data set, or versioning using snapshot technology, will be needed to support reliability in the design. For large designs, it also helps to identify the type of storage entities to be used, such as JBODs with host-based RAID and ZFS volume management or controller-based arrays.

Obtain information about RPO and RTO requirements from the business impact analysis. In general, the more rigorous the RPO and RTO requirements, the stricter the requirements will be for any IT components used in the design. As a result, costs will be higher, as shown in Figure 2 on page 7. On the other hand, if the RPO and RPO requirements are expressed in terms of multiple hours to days, using tapes might be the preferred solution.

The RPO typically describes the point in time at which data must be recoverable after a disaster. For example, an RPO of 1 hour indicates that the recovered data set must have been created no more than one hour previous to the disruption that resulted from the disaster. When the RPO is expressed in time, it does not tell the specific amount of data or the number of transactions that are acceptable to be lost. On the other hand, in transactional environments, RPO might be expressed as the number of transactions. In this case, the time-window information is absent.

A more restrictive (lower) RPO points toward a synchronous replication process as the preferred solution. A lower RPO also puts additional requirements on the link speed and influences the choice of transport technology. See the "Selecting the Interconnect Technology" section below for more details.

The RTO indicates how soon data must be back online and available for business operations. Thus, the RTO must be considered when determining the speed at which recovery operations need to be completed.

Obtaining information from an information lifecycle management (ILM) policy also helps in determining the type of solution to be used. When data is classified in different categories, such as critical, vital, sensitive, or noncritical, each class of data will have its own set of availability requirements. It might be feasible to put different availability architectures in place for different classifications of data as a way to minimize the cost of the total solution. The ILM policy might also specify that data that is not frequently used be archived, thus reducing the amount of data to be replicated.

It is also important to understand how a disaster is identified, who is responsible for declaring a disaster, what procedures are to be executed, and whether they are to be executed automatically or manually. This information should be contained in the business continuity plan. Any IT disaster recovery procedure must be consistent with this process.

Selecting the Replication Method

Now that the budget, distance, and RPO and RTO requirements are scoped, the next step is to determine the type of replication method to be used: synchronous or asynchronous.

Synchronous replication is recommended when consistent and identical data sets are required at the primary and secondary sites. A synchronous replication solution might not be practical if the distance between sites is greater than 70 to 100 km or data links have limited bandwidth.

The following are factors to consider when developing design criteria for a synchronous replication solution:

- The replication process and the transfer speed of the replication link must be able to handle application peak loads as determined by the peak ROC. For more details, see the "Rate of Change" section.
- When no data loss can be tolerated, factors such as line latency and the transfer speed of the replication link are key considerations in the design criteria. For more details, see the "Line Latency" and "Limitations Imposed by Distance on Replication Links" sections.
- Careful consideration must be given to recovery from error scenarios, such as a temporary loss of a link or loss of an array controller.
- One of the biggest challenges is the implementation and management of failover procedures, which must address the issue of preserving data consistency at the secondary site if the primary site goes down while synchronization is taking place. For more details, see the "Ensuring Data Consistency" section below.

Asynchronous replication is recommended if the distance between sites is greater than 70 to 100 km or replication links have limited bandwidth. The following are factors to consider when developing design criteria for an asynchronous replication solution:

- The average ROC must be less than the link bandwidth, as shown in Figure 5.

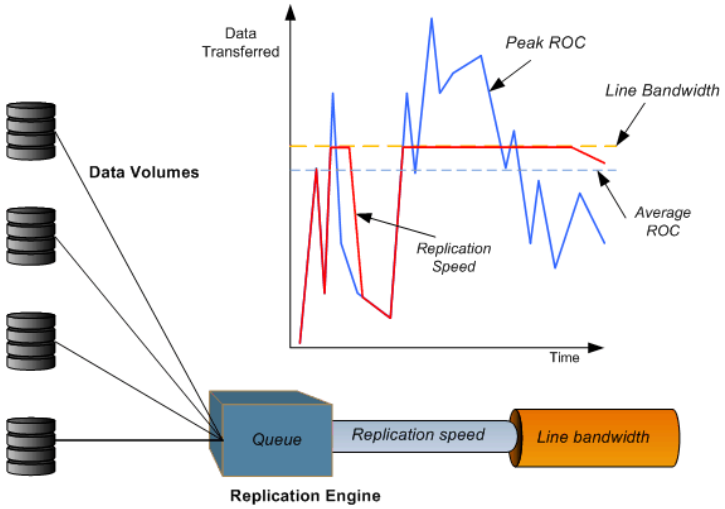

Figure 5. In asynchronous replication, the average ROC must be less than the line bandwidth.

- The RPO is a key consideration in determining the amount of data loss that can be tolerated. It might be helpful to consider RPO in combination with ROC. For more information, see the "Selecting the Interconnect Technology" section below.
- When defining and implementing failover and failback procedures, it is important to understand how the consistency of data is affected at the secondary site in each state of replication. For more details, see the "Ensuring Data Consistency" section below.

Selecting the Interconnect Technology

Once you have selected the replication method and have an idea what technology will be used for the inter-site link, you need to verify that the proposed interconnect technology will accommodate the I/O requirements of the application.

It is important to observe application I/O characteristics over a long-enough period to ensure reoccurring peak events, such as the creation of monthly and quarterly sales reports, have been captured. The ROC provides an important indicator of the expected usage of the storage system by the application and, thus, the bandwidth required for the inter-site link.

For synchronous replication, the maximum peak value of the ROC is used to determine the required line bandwidth so that a limitation in line bandwidth doesn't result in the throttling down of the I/O speed of an application.

For asynchronous replication, a combination of the RPO and the average ROC is used to determine the required line bandwidth. The maximum lag time between the primary volume and the secondary volume (the actual RPO) is the amount of time the ROC exceeds the line bandwidth of the data link, as shown in Figure 6. The actual RPO must be less than the specified RPO.

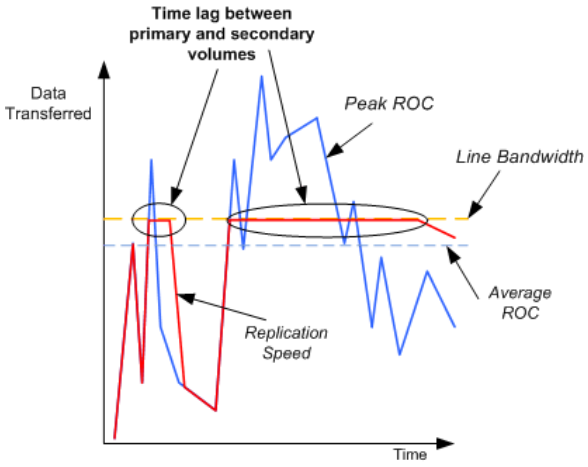

Figure 6. Lag time is affected by line bandwidth.

It is important to understand the impact on the behavior of the replication engine when the size of the block of data sent by the application to the primary storage subsystem differs from the size of the block sent by the primary storage subsystem to the secondary storage subsystem. For more details, see the section *Rate of Change* above for a description of how the replication load is affected by updates made by a dirty block mechanism.

Caution: In this section, it is assumed that the application(s), in combination with the storage subsystem at the primary site, can generate enough I/O operations to come close to saturating the link to the secondary site. When this is true, the ROC is a good indicator of the line speed required. If, however, the application environment is a transaction processing system (TPS) with a tendency toward single-threading, it is unlikely that the link bandwidth will be a limitation. For more details, see the previous "Physical" section. (Note that the link speed is expressed as maximum bandwidth, not as IOPS.)

Oracle's SWAT tool can be used to generate I/O traces over periods during which the application generates high I/O activity. The performance graphs produced by SWAT are helpful in obtaining an understanding of the application I/O profile. Also, the captured traces can be used to simulate the application I/O in a lab environment to test a proposed data availability architecture or determine the performance impact of the replication process on the application.

Ensuring Data Consistency

Careful consideration must be given in the design of the replication solution to the consistency of data at the secondary site. Behavior with respect to data consistency is different for synchronous and asynchronous replication and, also, for different replication products that might be incorporated into a solution. A well-designed solution must address how to mitigate the risks of data inconsistency for the following:

- Each volume to be replicated
- Each set of volumes that have dependencies on one another, such as an application database

Data inconsistency becomes a risk when application data changes are applied to the secondary volume in a sequence that is different from the sequence in which they are applied to the primary volume.

The replication product used should have the ability to queue writes in the replication engine without losing time-order information. A product that simply uses a scoreboard bitmap mechanism to keep track of changed blocks cannot guarantee data consistency when replicating the changes recorded in the scoreboard bitmap.

Let's look into this in a bit more detail. On a conceptual level, the following states can be distinguished for a replication process:

- *Replicating state* – Normal operational state. Data written to the storage subsystem is replicated to the secondary system.
- *Synchronizing state* – Entered after replication has been suspended for any reason and the replication process has been restarted at the primary site. Any changes queued at the primary site are now replicated to the secondary site. Once the secondary site has caught up with the pending changes, the replicating state is reentered.
- *Suspended state* – Entered when the replication process has been stopped, either manually or as a result of a link failure. Data changes are tracked at the primary site during this state.
- *Failed state* – Entered when the storage subsystem has lost the ability to track changes in data between the primary site and the secondary site. User intervention is needed to determine how to recover from this state. In most cases, a complete resynchronization of the data to the secondary site is needed.

When a replication set is in the replicating state, all changes on the primary volume are replicated to the secondary volume. For synchronous replication, changes are applied at the secondary site in the same order in which they occur in the primary volume, but this is not necessarily true for asynchronous replication.

Figure 7 shows the states for a synchronous replication when bitmap score boarding is used to track changes, and it describes the events that typically cause each state transition. The data in the volume(s) on the secondary site are always in a consistent state during the replicating state. If the state changes to suspend mode, the secondary volume is still in a consistent state.

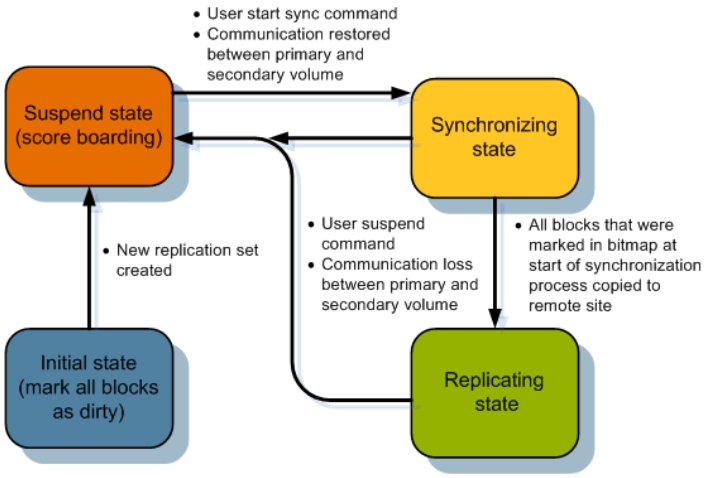

Figure 7. Replication states for synchronous replication use a bitmap to track changes.

However, if changes are tracked at the primary site using a bitmap, from the moment the state changes to synchronizing until the transition to the replicating state is complete, the secondary volume is not in a consistent state. If the primary site goes down during this synchronization period, the application will not be able to use the data volume(s) at the secondary site because the data might not be in a consistent state. (Note that this implies that it is not a good idea to allow the replication to be restarted automatically by the replication engine.) One way to avoid this situation is to create a snapshot at the secondary site before starting the resynchronization.

When the replication design implements a queue in addition to score boarding, write order is maintained for data in the queue. Although the state changes become slightly more complex, as shown in Figure 8, data consistency is maintained during the replicating state, while data is replicated from the queue.

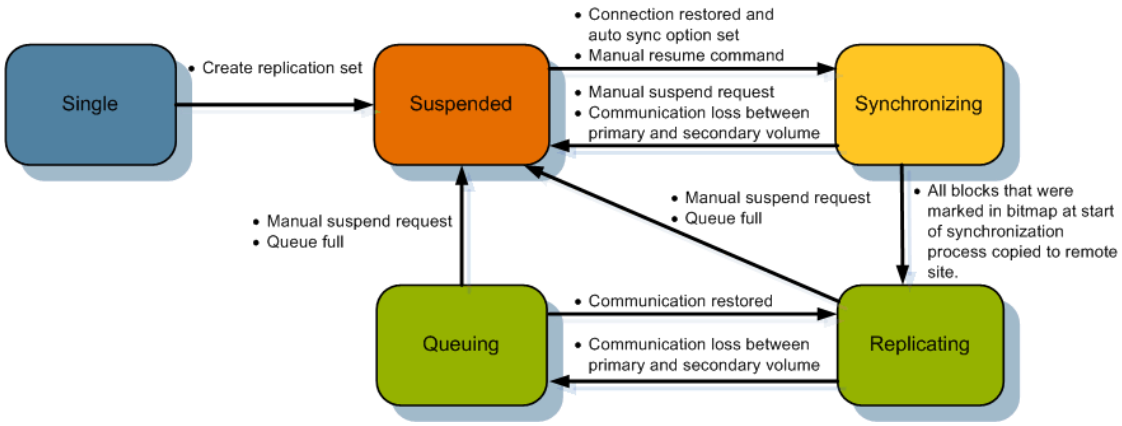

Figure 8. When replication queues are used for replication states for synchronous replication, write order is maintained.

This solution has a number of advantages. Small disruptions in communication on the links between the sites do not cause the volumes to go into the suspend state, so no user intervention and monitoring is needed to recover from these events. Also, the risk of a volume on the secondary site being in an inconsistent state is reduced. The size of the queue is a determining factor in minimizing this risk.

The state diagrams in Figure 7 and Figure 8 also apply to asynchronous replication, but with one significant difference. When the replication product is not using queueing functionality, data volumes on the secondary site are never in a write-consistent state unless special measures have been put in place. The handling of WOC in an asynchronous replication is discussed in the next section, "Handling WOC in Asynchronous Mode."

Handling WOC in Asynchronous Mode

The way WOC is handled in an asynchronous replication depends on the replication product. In an asynchronous replication, the changes applied to the secondary volume lag behind the changes applied to the primary volume. The amount of lag depends on the ROC, the number of volumes to which data is being replicated, and the line speed.

One consideration is how the replication engine tracks the difference between data changes on the primary volume and the secondary volume. Ideally, the replication engine has knowledge of when a change was made and is able to replicate the blocks at the secondary volume in the same time sequence in which they were applied to the primary volume. To do this, the replication engine must have access to a queue that keeps track of each changed block and the order in which it was changed. If multiple changes are made to a block over time, each change needs to be replicated to the secondary volume in sequence to guarantee data consistency.

Another consideration is the order in which changes are made to multiple volumes at the primary site as a result of a single database transaction and, subsequently, to multiple volumes at the secondary site. To maintain data consistency, replication products allow volumes to be assigned to a WOC group. The replication engine then sends the changes to the volumes in the WOC group in a single replication stream, thus ensuring the order of changes is guaranteed across the group. Note, however, that this approach can have a negative impact on performance when changes are replicated to a WOC group at the secondary site.

A third consideration is whether snapshots will be used to reduce the risk of data loss if the primary site goes down during the synchronization process, when data might not be in a consistent state at the secondary site. Factors to consider are the type of replication used (synchronous versus asynchronous) and the way a replication product maintains WOC during the various state transitions in the replication process. It is helpful to understand how the states of a replication process are implemented in the replication product being considered and how that replication process recovers from failure scenarios.

If the replication solution is to be deployed in a transactional environment, a set of scripts must be designed to quiesce the database and replicate all outstanding changes to the secondary site before creating a snapshot at the secondary site.

Another option is to combine replication with the use of mirrors on the primary site. After quiescing the database, a mirror instance can be split off. The replication product must be capable of using a mirror member as a primary replication volume, and it must have the ability to split and rejoin the mirror member, triggering a complete resynchronization of the volume to the secondary site.

Conclusion

To help ensure business continuity after an unexpected disruption due to a disaster, many organizations incorporate into their disaster recovery strategy the replication of vital business data to one or more remote locations. The success of such a strategy requires an efficient replication process that meets the business continuity requirements of the organization.

This paper provided the information needed to determine design criteria and select an appropriate method of replication for a replication architecture that meets these requirements. First, the terminology and concepts associated with business risk mitigation and data replication were explained. Then important factors to consider when identifying customer requirements were discussed along with how to use these requirements to design the data replication architecture, with emphasis placed on the importance of approaching the solution from a business-process point of view rather than from a more-limited IT-technology point of view. Finally, consideration was given to ways to mitigate the risk of data inconsistency at the secondary site.

Resources

The following resources were referenced earlier in this document:

- DRI International Website: www.drii.org
- *The ARRL Handbook for Radio Communications*: http://www.arrl.org/arrl-handbook-2011
- Oracle's StorageTek Workload Analysis Tool (SWAT), which is available at the My Oracle Support Website (access requires support contract): https://support.oracle.com

Here are some additional resources:

- *Oracle Real Application Clusters on Extended Distance Clusters* by Erik Peterson http://www.oracle.com/technetwork/database/enterprise-edition/extendedrac10gr2-131186.pdf
- *Oracle Solaris Cluster Essentials* by Tim Read (ISBN-13: 978-0132486224) http://www.amazon.com/Oracle-Solaris-Cluster-Essentials-Administration/dp/0132486229

White Paper Title
September 2011, Version 1.1
Author: Peter Brouwer

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

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com

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

Hardware and Software, Engineered to Work Together