

Disaster Avoidance and Recovery for Oracle-based Businesses

DataMirror[®]
The experience of now.[™]

Disaster Avoidance & Recovery for Oracle-based Businesses

Executive Summary

Heightened security concerns and the all too real threat of domestic terrorism have focused attention on the need for high availability solutions and strategies for disaster recovery. High data availability is or should be top-of-mind for all organizations.

Every organization that relies on Oracle data to conduct business must be concerned about ensuring the availability and resiliency of their Oracle data assets. By resiliency, we mean the ability of a company to spring back or recover from any kind of unplanned or planned system downtime with minimal impact on the business. Globalization, business consolidation and the Internet have driven the need for 24/7 data availability. Companies now routinely have thousands of employees, customers and suppliers all over the world accessing their Oracle-based systems, 24 hours a day, 7 days a week. With costs of downtime ranging in the thousands of dollars per minute range, high data availability is a business imperative. If your database server is off-line or your business-critical Oracle data is unavailable for any period of time, your ability to transact business may be seriously impaired.

Disasters strike suddenly. Only prepared companies remain open for business as if nothing happened or with minimal inconvenience. Unplanned outages aside, every system must be taken off-line periodically for routine maintenance including hardware and software upgrades. For many companies, overnight and weekend maintenance windows to complete these upgrades no longer exist or have become too short to carry out the work required. Oracle-based companies across all industries must have a contingency plan in place in the event of downtime or disaster. Within this formal contingency plan, organizations require the following elements:

- Data and applications criticality analysis
- Data backup plan
- Disaster recovery plan
- Emergency mode operation plan

A high availability software solution and a proven disaster recovery methodology form the technology foundation of this contingency plan. This paper discusses high data availability planning for companies with Oracle-based computing environments. It outlines available options for data backup and disaster recovery and considers what to look for in a high availability solution. An overview of DataMirror's iReflect™ resiliency software for Oracle is also included.

Why back up Oracle database systems?

Although the availability of Oracle systems is reportedly high, things can still go wrong. How vulnerable is your “unbreakable” Oracle database to downtime or disaster? How well does your system cope with unplanned and, more importantly, planned downtime? With the growth of 24/7 e-Business and today’s global business hours spanning all time zones, many companies must be on-line all the time. This means that organizations require a contingency plan for ensuring continuous business operations.




Figure 1: Causes of unplanned and planned downtime.

Source: Donna Scott, GartnerGroup “Surviving in a 24 Hour World”

Downtime is at best an inconvenience to users. At worst, system failure can result in lost business, lost opportunities and diminished customer loyalty. If system outages prevent users from completing tasks and conducting business via corporate systems, the long-term effects may be costly and the damage to your business may be beyond repair. An outage can make it difficult or impossible to share order, inventory and other information with suppliers, trading partners or customers. Internally, if users are unable to access Oracle data stored on servers, decision making, business monitoring and other critical functions may grind to a halt.

Cost of Downtime

Cost of downtime is a major factor in deciding to invest in a high availability strategy. The elements required to make a reasonable estimate of cost of downtime include losses in labor, revenue and any other costs incurred due to the organization’s inability to access its Oracle data. Calculating the exact cost of downtime can be difficult given the range of possible variables. For example, lost access to an order or inventory database is much more costly during peak business hours than at 2:00 am. Lost customer goodwill and loyalty must also be factored into the equation. Exactly how much is a loyal customer worth over time?

Labor Costs

Part of the formula for calculating cost of downtime includes totaling the cost of normal operations for all staff affected by system outage. This involves estimating the extent to which employees are affected by downtime. Are some employees able to carry out other tasks during an outage? What percentage of their normal assignments are they able to complete? Order processing departments, for example, may be totally dependent on continuous access to inventory, shipping and cataloging data. Many companies find that even when they calculate a very conservative cost of labor during downtime, the figures can be astounding.

Lost Revenue

As with labor costs, arriving at an accurate figure for revenue loss due to downtime can be challenging. However, an assessment of percentage of revenue loss directly attributable to downtime can still be made. Lost revenue can vary greatly from organization to organization depending on industry, business model and other factors. An e-Business company or on-line retailer dependent on Internet-based order entry applications would be seriously impacted by downtime (essentially 100% lost revenue for the duration of the outage). Similarly, a call center in which staff enter orders directly into production systems would be seriously affected by system downtime.

Sample case study

There are often losses incurred due to downtime that are less tangible. Frequent or extended periods of downtime can adversely affect a company's reputation and stock price. Beyond lost revenue, downtime can also cause a close competitor to take the lead and make gains in the market while your organization loses competitive advantage. The sample case study below provides a guide and calculations that organizations can use to estimate the approximate cost of downtime. You can also visit <http://www.datamirror.com/solutions/highavailability/calc.asp> to access an on-line downtime calculator that can help you estimate the cost of downtime within your own organization.

A manufacturing company's primary Oracle database server goes down for 2 hours:

Annual gross revenue of the company	50,000,000
Number of employees	200 employees
Average value of an employee working hour (including salary, benefits and administrative cost)	\$ 55.00

Cost of data loss

Employee hours required to restore lost data	50 hours
Cost of lost data (hours needed to restore data) x (value of employee working hour)	\$ 2,750

Cost of employee downtime

Hours the system is down	2 hours
Cost of employee downtime (number of employees) x (hrs downtime) x (cost of employee working hour)	\$ 22,000

Cost of lost sales opportunities

Number of sales per year	10,000
Number of customers	1,500
Estimated number of sales lost due to system failure	2 lost sales
Cost of sales lost due to downtime (gross revenue) by (number of sales) = (cost per sale) x (number of lost sales)	\$ 10,000

Cost of lost customer goodwill

Estimated lost customers	2 lost customers
Cost of lost customers	\$ 67,000

(number of sales / year) by (number of customers) =
(number of sales / customer / year) x (cost / sale)
x (number of lost customers) = cost of lost customers

Estimated total cost of downtime

Cost of lost data	\$ 2,750
Cost of employee downtime	\$ 22,000
Cost of sales lost due to downtime	\$ 10,000
Cost of lost customers	\$ 67,000

Total cost \$ 101,750

The Bottom Line

As the sample case study calculations illustrate, the costs of downtime can be staggering. Given the potential for loss, having a reliable high availability solution in place becomes a business necessity. Tape backups provide a degree of availability but the incremental nature of these backup procedures may require considerable time and effort to restore data lost since the last backup. Paper trails may be difficult to trace and re-input into the system. In addition, tape backups must be scheduled for off peak hours and use up processing time that could be utilized for other tasks. For many types of organizations, a true high data availability solution is the only answer.

Implementing Oracle High Availability

High availability for Oracle is not so much a product as it is a project. It will involve input from almost everyone across your organization from the CEO (acting as a project champion or executive sponsor) and upper management to application programmers, network administrators and end users.

From a systems manager's point of view, a high availability implementation is never an "out-of-the-box" install. Customization, monitoring and a good amount of elapsed runtime are required before you can consider your high availability project to be fully functional. If a vendor or consultant tells you that high availability for Oracle works right out of the box, begin to question the vendor or consultant's definition of high availability.

When you invest in a high availability solution, you are purchasing a type of insurance. You are purchasing a product that will allow you to sleep at night knowing that your critical Oracle data is safely backed up on a recovery system. Should your production system fail, you can quickly switch users over to the recovery system and continue running your business with minimal interruption.

You are also building a long-term relationship with the high availability vendor. A typical installation may only take a day, but the planning, configuration, monitoring, customization, training and new operational procedures may take several months to implement. The good news is that high availability products on the market are sophisticated, robust and easy to install. High availability vendors are also experienced in the planning and configuration of disaster recovery.

Organizations considering a disaster recovery solution for Oracle must also clearly understand the methodology and approach employed by the software vendor or consultant to implement their disaster recovery solution. High Availability is a long-term technology partnership between the customer and the vendor. From the outset, both the customer and the high availability vendor need to agree on high availability needs, rough timelines for project completion and project risks and assumptions. Both parties must have a common understanding of disaster recovery requirements and the scope and business goals of the implementation.

What are the Options?

Some applications require more availability than others. The scale of a company's operations, the budget available for high availability projects and the feasibility of alternatives must be carefully considered before beginning a disaster recovery project. For example, straight media backup may be sufficient for a web-enabled Oracle database used for remote sales representatives to check daily or weekly updates to stock and pricing. The remote sales team may be able to function well enough without access to the database for several hours or even several days without seriously impacting business operations.

Media backups are seldom sufficient in the case of e-Business applications. In this case, the e-Business application provides on-line customers with real-time access to an inventory database. The need for continuous access to this inventory data, as well as customer profiles and other data-rich content, is critical.

There are several alternatives for Oracle high availability and disaster recovery. This section examines the options available to Oracle-based businesses including the benefits and potential drawbacks for each approach. The viability of any given solution depends on the organization's specific environment and individual needs for data and system access.

Media Backups

Media backups involve copying critical files directly to tape or optical disk for recovery purposes. This means that the informational database is manually restored once the system comes back on-line after a failure. Media backups do have a serious downside. They often require significant time and resources to restore the data, and critical data may be lost. They become impractical when the volume of data required for backup becomes large. The higher the volume of data that is backed up, the longer it will take to restore. Essentially, more downtime is required as a result of lengthy restoration processes.

RAID 5

RAID 5 (redundant arrays of independent disks) offers organizations the ability to maintain shadow copies of information across separate disks. It allows missing Oracle data to be restored from one disk to another in the event of a disk failure. While offering increased availability, RAID technology does not adequately protect against multiple disk failures or site failures.

Disk Mirroring

Disk mirroring involves duplicating data writes to a twin set of disks. In other words, if an organization's Oracle system has five disks, then the solution maintains a twin set of another five disks for backup and recovery. This protects against failures such as a power outage within one disk set, multiple disk failures and physical damage to disks. Disk mirroring is a more expensive alternative to RAID 5, however this approach does not require companies to purchase another Oracle system for backup and recovery and remains an affordable solution for many organizations. While offering increased availability, disk mirroring does not adequately protect against site failures.

File/Volume Level Replication

This approach is typically hardware-based with a software component. File/volume level replication solutions can create a duplicate of the production database and quickly apply it to a secondary database. File/Volume level replication uses disk-block replication and logical volume replication. In the event of a disaster this type of solution switches the machines to the secondary set of disks and brings the database back on-line.

File/volume level replication often requires costly dedicated hardware and high speed network resources. This solution requires that the hardware and software be exactly the same on both the primary and secondary which can make the upgrade process much more difficult.

Usually, file/volume level replication solutions are local or "same-room" solutions. These solutions are limited in their ability to be geographically dispersed because of the extremely high cost of implementing hardware at remote sites. For most companies where cost is a factor, file/volume replication remains a local solution that will create an exposure for the user if they lose their entire computer room.

Another issue when considering this type of solution is that since there is no secondary database (only a copy of the primary database files exist), the user will not have the ability to off-load any query and reporting activities or data warehouse replenishment to their secondary database and thus will not be able to realize any workload distribution benefits from the solution.

Additionally, this concept uses a synchronous transport method which means that the user might be constrained by their communications network. This means that any communications failure will result in the database stalling. File/volume level replication provides high levels of Oracle availability but at a high cost that may be beyond the means of many companies.

Inherent Methods

The Oracle database itself includes some inherent data backup and high availability capabilities. In the case of the Oracle 9i database, these capabilities include Real Application Clusters and Data Guard.

While these inherent capabilities may provide a sufficient level of performance scalability and disaster recovery for many companies, other organizations require a more advanced and comprehensive data backup and disaster recovery solution that complements and extends the inherent availability of Oracle.

Real Application Clusters (RAC)

While Oracle 9i Real application Clusters (RAC) are primarily a high performance scalability solution as opposed to a true high availability solution, RAC does provide a degree of availability by reducing recovery time after a failure. Real Application Clusters (RAC) allow an Oracle database to be accessed concurrently on two or more systems in a clustered environment.

The implementation of Oracle 9i Real Application Clusters results in a shared disk environment in which multiple systems share one disk subsystem. A single copy of the Oracle database is distributed across two or more servers in the cluster. In the event that one of the servers fails, the surviving servers can perform recovery on the failed Oracle database instance with minimal exposure to downtime. However, because multiple systems share the same disk subsystem, the RAC solution creates a single point of failure. Any problem in the disk subsystem can result in a major exposure to downtime.

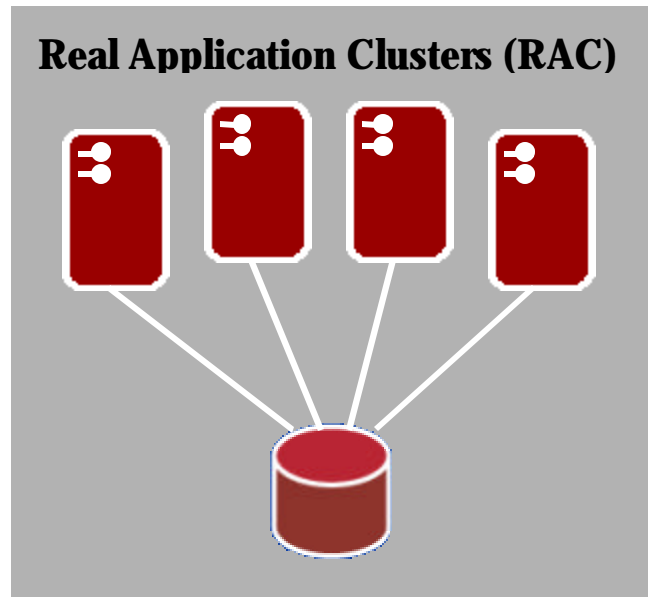


Figure 2: Real Application Clusters provide a solution for increasing the availability of Oracle 9i databases in a clustered server environment

Solutions like RAC often require costly hardware and may be difficult to setup and maintain. Due to high cost of deploying multiple servers and providing adequate bandwidth for remote replication, the Real Application Clusters configuration is usually a local solution only. This means that RAC does not support true disaster recovery in which data must be continuously backed up to a remote, geographically-dispersed recovery site. Solutions like RAC may also be very sensitive to changes to the hardware or software environment.

Oracle Data Guard

Oracle Data Guard for Oracle 9i is a point-in-time replication solution that involves copying Oracle Redo-logs to either disk media or a recovery system and then using a batch apply process. Data Guard takes a complete redo log, copies it to an archive log, then locks the secondary database while applying an exact replica.

Due to the batch apply of this process, the data on the recovery database will always be slightly behind or out of sync with the data on the primary database. This makes using the recovery system to run effective query and reporting or data warehousing very difficult. While solutions like Data Guard provide an effective solution for certain applications, log-based, batch-oriented backups do not provide true availability of the recovery system. The solution also exposes businesses to the risk of block corruption being propagated to the recovery system.

Data Guard can be configured to run in a number of different modes including synchronous or asynchronous transport mode. In asynchronous transport mode, Data Guard is a log-based replication software with a batch apply. It is important to emphasize that this process cannot provide a true real-time solution. Data Guard takes a complete redo log, copies it to an archive

log, then locks the secondary database while it applies an exact replica. Because of the latency of the apply, this process will result in two databases which are never fully synchronized. The secondary database will never be able to catch up to what has happened on the primary database.

Synchronous transport mode essentially guarantees data delivery. However, there are several considerations with using a synchronous transport in a real-world setting. Any data that is captured on the primary redo logs is sent to the secondary redo logs and only then can the data actually be applied to the primary database. This creates a co-dependency on the secondary system and the communications link for your primary database. If either fails, then your production system stalls. Additionally, latency increases dramatically if there is not enough bandwidth between nodes. This solution requires an extremely fast communications network which, due to cost constraints, generally limits users to a local or non-geographically dispersed solution.

Using this solution increases the points of failure for your primary database. There are several other issues with this approach: if there are any communications failures the primary database will become unavailable; the primary database is slowed down by the overall communications speed; and the primary database will be unavailable if the secondary system becomes unavailable. All of these situations will have an immediate impact on your business.

Advanced Availability

The trouble with media storage, RAID5, Real Application Clusters, Data Guard and other alternatives is that they all require some degree of system downtime. It takes time to copy critical data—time that could be better spent serving users. Additionally, the server may need to be taken off-line when it comes time to restore the database. This process is relatively straightforward but can take a significant amount of time to carry out depending on data volumes.

When inherent Oracle backup and recovery processes do not adequately meet the availability requirements of organizations, they must develop a more advanced strategy for ensuring continuous access to corporate applications and the data that powers their business. The ability to mirror Oracle data in real-time and provide a framework for high availability and disaster recovery is critical.

Components of a continuous Oracle resiliency solution

There are three major components involved in continuous 24/7 Oracle resiliency: mirroring software, role switching facilities and a process monitoring interface.

Mirroring

The most important element in an Oracle availability solution is the mirroring software that will synchronize the recovery system in real-time from the primary system. This software component must be fast, transaction-consistent and guarantee data delivery at the secondary site even in the event of a network or system outage. The mirroring software must also be

two-way to allow both the production and the secondary machines to act as a source and target to each other. Full two-way functionality gives the secondary system real-time capabilities to re-synchronize the production machine, while enabling users to continue working on the recovery system.

Ultimately, organizations look for a situation that involves switching users to the secondary site while recovery takes place on the primary production system. During operation at the backup site, Oracle database changes are being captured on that system. Once the primary system is recovered, the mirroring software can enable all the database changes captured on the recovery site to be replicated back to the primary system. In time, the primary Oracle system will be synchronized with the secondary system and users can be switched back to the primary system with minimal disruption.

Switching

A switching mechanism is a critical component of any high availability solution. The switching mechanism handles operational role switching between primary and recovery systems. In today's sophisticated network environments, a virtual switchover to a new system may be as simple as logging on to a second machine. Users and IT staff can be alerted to an outage at which time network traffic can be re-routed or users can click a different icon to establish a connection to the recovery system. Switching mechanisms automate user and role switching. Switching software can either notify administrators of an outage or initiate other conditional events after a user-defined delay. These types of high-speed switching facilities are designed to make system outages as seamless and transparent to users as possible.

Monitoring

The third component of any high availability solution is an interface for tracking the entire mirroring process. The ability to monitor historical and current activities is important as administrators need to know how the solution is performing. What is the time latency in moving data between primary and recovery systems? In any high availability environment, performance equals availability. Administrators must have the means to monitor mirroring performance and latency for system backup.

What to look for in an Oracle resiliency solution

Implementing a high availability solution for Oracle environments involves mirroring data in real-time between primary and recovery systems in order to secure operations from business interruption. The software used to enable such a solution should meet certain standards. It should be capable of refreshing the primary system while active. It should also be capable of moving data in a two-way configuration. That is, the primary server should be able to act as both the source or the target of data mirroring while it is active. If the software doesn't include this functionality, near continuous availability is impossible.

Transaction Consistency

The software should be transaction-consistent. This capability ensures that transactions are performed on the secondary machine in precisely the same order in which they occurred on the

primary server. Technologies use the Oracle redo log to track the order of transactions. Transaction consistency ensures data integrity.

Fault Tolerance

The software should include fault tolerance or guaranteed data delivery at the secondary server. If there is a network, system or communication outage during replication, the secondary server should reflect the last complete transaction. This eliminates the possibility of duplicate entries or partial transactions when it comes time to switch operations back to the primary server. This is absolutely critical for e-Business and financial transactions where the integrity of the transaction must be preserved at all costs.

Ease of Use and Administration

The high availability application should be easy to install and administer. The last thing companies need to worry about when setting up a high availability solution is custom programming and other time-consuming, costly and resource-draining activities. Data mirroring software should install quickly and be ready for operation with minimal effort. Additionally, users should be able to administer both the data mirroring, switching and monitoring functions from a single unified interface.

Selective Table Mirroring

The high availability software should give users the option to select the specific tables that they want to mirror from the primary system to the recovery system. From a performance perspective, it is not always desirable to mirror absolutely everything to the recovery system. An Oracle database may contain both critical and non-critical tables. Selective table mirroring allows users to exclude workfiles or any other non-critical data that is not deemed vital to driving the business.

In summary, companies requiring an Oracle high availability solution should ensure that they choose a package that integrates as many of the critical components as possible including mirroring, switching and monitoring. Piecemeal solutions—separate components for operational switchover, data mirroring and performance monitoring—can involve separate installation, configuration and administration for each element. Taking all of these factors into consideration, organizations require a fully integrated solution for Oracle high availability and disaster recovery.

About DataMirror iReflect

DataMirror iReflect is designed specifically to virtually eliminate system downtime for Oracle environments. The software provides a powerful solution for disaster avoidance and recovery. iReflect uses intelligent database-level mirroring to provide a cost-effective solution to help businesses running Oracle achieve superior uptime and availability for their precious data assets. It mirrors Oracle database transactions from the primary system to the recovery system in real-time. The software also has the ability to detect primary system failure and invoke operational switching to enable highly available business operations in Oracle environments. iReflect is easy to configure and supports log-based replication of Oracle data-types.

DataMirror iReflect is a tightly integrated, fully functional product that will enable you to get the most benefit from your Oracle systems in a high availability environment. It is designed specifically for easy configuration, testing and monitoring. The product is a single code based

package that handles data mirroring, operational switching and monitoring. There is only one product to install and maintain and one product interface to the end user. iReflect is installed on both the primary and the recovery Oracle system within a matter of minutes.

What happens behind the scenes? iReflect's high performance capture, transmit and flow (CTF) architecture leverages the Oracle Redo Log changes via a process called scraping. iReflect constantly monitors redo-logs to ensure that all changes including inserts, updates and deletes are captured and flowed continuously across a communications line to the recovery Oracle system. Once these changes are received on the secondary system, they are applied to the recovery Oracle database tables. While production users are updating the primary Oracle database, transactions are being sent as quickly as possible in real-time mode to the recovery system where they are staged and then applied to the target database. iReflect also gives users the flexibility to select the data that is mirrored to the recovery system so that administrators can choose to exclude workfiles etc.

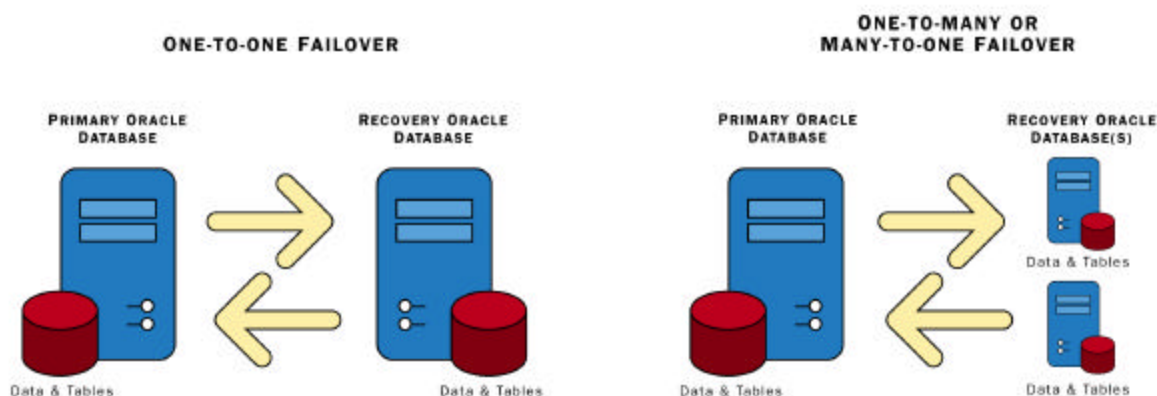


Figure 3: DataMirror iReflect supports both one-to-many and many-to-one failover for Oracle databases. The solution forms the foundation of a disaster avoidance and recovery plan while enabling operational efficiencies including workload distribution, query off-loading and the ability to conduct business 24/7.

Built on the Oracle 9i architecture for maximum performance, iReflect enhances the inherent availability features of your Oracle databases to provide thorough protection against all failures. Whereas tape backups may require considerable time and effort to restore data lost since the last backup, iReflect provides a comprehensive and easy-to-use resiliency solution in the event of planned or unplanned downtime. iReflect's flexible architecture supports one-to-many, many-to-one and two-way replication scenarios to ensure that the high availability solution can easily accommodate your current systems and configuration as well as scale and evolve as your environment changes to future-proof your investment.

Key Features and Benefits of iReflect

Key Product Features:	Business Benefits
<p>High speed operational switching Should an outage occur, iReflect can be used to rapidly and seamlessly switch physical operations to a recovery system so that normal activities can resume. iReflect automatically detects when the primary system has failed and invokes automatic failover or other corrective measures specified by the user.</p>	<p>Human error has been cited as the chief cause of system downtime. The ability to automatically invoke operational switching with minimal human intervention helps to ensure maximum uptime and a smooth, seamless return to normal business operations.</p>
<p>Powerful log-based replication iReflect utilizes the Oracle Redo-log to provide intelligent, real-time mirroring at the database level. By leveraging the Oracle Redo-log to enable real-time data-level mirroring, iReflect provides high performance while ensuring data integrity and accuracy.</p>	<p>In an Oracle high availability environment, performance equals availability. iReflect's real-time, log-based replication ensures high performance mirroring between primary and recovery systems to help achieve near zero latency for backup and recovery so that the data you need is where it needs to be for recovery in the event of an outage.</p>
<p>Real-time database-level mirroring Changes to the source Oracle database are mirrored in real-time with near zero latency to ensure maximum uptime and availability with minimal human intervention.</p>	<p>iReflect provides data availability when you need it most with maximum uptime, speed and performance. Data is replicated at the database level, ensuring that block corruption is not propagated. This ensures data integrity and minimizes exposure to system downtime.</p>
<p>Real-time Auto-registration Generic specifiers can be used so that new tables are replicated automatically to the recovery database. Auto-registration saves time and resources by automating tasks that would otherwise need to be performed manually.</p>	<p>This feature reduces administration costs while maintaining the integrity and accuracy of data transactions. Auto-registration also maximizes applications resiliency by ensuring that any changes, adds or deletes made to an application's Oracle data store are automatically mirrored to the recovery system, making businesses much less vulnerable to data loss or corruption.</p>
<p>Event-driven status monitor Interactive, real-time inquiry screens snapshot current activities while historical replication statistics are tracked for trending analysis. Multiple views allow you to view the relationships between object specifiers, groups and target systems as well as monitor replication on both the source and recovery system.</p>	<p>The ability to monitor Oracle systems and ensure the performance of continuous mirroring operations delivers enhanced peace-of-mind for administrators and business leaders. Comprehensive monitoring along with alarms and alerts can result in faster implementation and lower administration costs.</p>
<p>Intuitive Java Graphical User Interface 100% Java-based GUI offers point-and-click flexibility for configuring and administering the high availability environment.</p>	<p>The portability and platform independence of Java helps iReflect maximize flexibility and ease-of-use while dramatically reducing maintenance costs. Companies can realize a rapid return on their software investment</p>

What added value does iReflect deliver?

With DataMirror iReflect, companies not only gain peace of mind knowing that they have a reliable, high performance resiliency solution in place in the event of an outage, but they also gain improved overall system performance. Through a workload distribution strategy, organizations can off-load user query activity to the recovery Oracle system to improve overall efficiency of the environment. Because iReflect mirrors data in real-time, companies can actually create a shadow Oracle database on the secondary machine and use it for batch jobs, queries, reports, building data warehouses and other resource-intensive tasks that may ordinarily impact performance on the operational system. Data integrity and consistency is ensured and organizations can extend their technology investment by maximizing system resources.

Ease of Use

iReflect installs without the need for programming changes to existing applications. Intuitive graphical screens simplify the setup process, allowing users to map selected data to recovery systems without scripting or coding. iReflect is a single-setup product with no need to run separate installs for each component. Users are not required to run separate tasks from different applications. Data mirroring, switching and monitoring processes are all centrally administered. iReflect also features additional capabilities to simplify initial synchronization between primary and recovery systems.

Full Range of Products and Services

DataMirror in-house professionals are available for any implementation assistance that organizations may require. Seasoned Oracle high availability experts can help outline requirements, install software, test the implementation and help with any other project needs. DataMirror brings extensive experience and expertise in high availability and clustering solutions. DataMirror also offers global 24/7 support via telephone, e-mail and the web.

DataMirror's broad product family provides Oracle customers a single-vendor, end-to-end solution for Oracle resiliency, multi-platform data integration and data monitoring. When iReflect is implemented alongside DataMirror Transformation Server[™] for Oracle and LiveAudit[™], Oracle customers have all the tools required for a comprehensive enterprise data management strategy.

DataMirror Transformation Server captures, transforms and flows data in real-time between Oracle, Microsoft SQL Server, Sybase, DB2 UDB and XML across a range of computing platforms including UNIX, Linux, Windows NT/2000/XP, IBM OS/400, OS/390 and z/OS. LiveAudit for Transformation Server allows users to capture all data that is added, changed or deleted in an application database to create real-time audit trails that preserve historical information which would otherwise be overwritten. LiveAudit allows companies to monitor, track and report on critical database information for enhanced security and added protection against user error, tampering or other activities that may impact business operations.

The complete DataMirror package including comprehensive software solutions, methodology and support results in a Oracle resiliency solution with a difference.

Summary

As stated earlier, high availability is a long-term project and expectations must be managed. Don't expect to get your Oracle high availability solution fully up and running to the point where you can switch over users to a recovery Oracle system in a matter of hours or days. You need to plan, configure, monitor and give your Oracle-based applications time to run through their most critical cycles to ensure that the high availability environment meets all of your requirements—especially during critical month-end or billing processes as peak load is more likely to occur during these periods.

Organizations also need to ensure that you have the resources required in terms of hardware and communications. There will be some additional disk requirements on both the primary Oracle system and the recovery system. Communications will also have to be sufficiently large enough to flow all of the transactions that are occurring during peak production runs across to the recovery system.

DataMirror consultants or certified business partners can help you to determine these specific requirements. Testing and monitoring are also very important and will require some elapsed run time as well as time for possible problem solving and optimization.

In closing, it is important to remember that as an organization grows more dependant on Oracle data for e-Business, decision support, order processing, inventory control, accounting and other operations, the need for resiliency and high availability becomes imperative. The many tangible and intangible business benefits of investing in an Oracle resiliency solution include:

- 24/7 access to key Oracle databases
- Disaster avoidance and recovery for Oracle-based businesses
- The ability to avoid the costs of both planned or unplanned downtime
- The ability to enhance service levels and promote customer retention and loyalty
- Highly efficient distributed data processing and workload balancing
- The ability to ensure the continuity of supply chains and e-Business environments

While there are several alternatives for Oracle high availability, none offer the integrated design, high performance architecture, ease of use and overall security of DataMirror iReflect.

Copyright © 2001 DataMirror Corporation. All rights reserved. DataMirror, iReflect and *The Experience of Now* are trademarks or registered trademarks of DataMirror Corporation. All other brand or product names are trademarks or registered trademarks of their respective companies.

About DataMirror

DataMirror (Nasdaq: DMCX; TSE: DMC) delivers solutions that let customers integrate data across their enterprises. DataMirror's comprehensive family of products includes advanced real-time capture, transform and flow (CTF) technology that gives customers the instant access, integration and availability they demand today across all computers in their business.

Over 1,500 companies use DataMirror to integrate their data. Real-time data drives all business. DataMirror is headquartered in Toronto, Canada, and has offices worldwide. DataMirror has been ranked in the Deloitte and Touche Fast 500 as one of the fastest growing technology companies in North America.

DataMirror[®]

www.datamirror.com

1 800 362-5955