

Creating A Highly Available Database Solution

Advantage Database Server and High Availability

TABLE OF CONTENTS

- 1 Introduction
- 1 High Availability
- 2 High Availability Hardware Requirements
- 2 High Availability Personnel Requirements
- 3 Advantage Features for High Availability
 - 3 Server-Side Aliases
 - 3 Online Backup
 - 3 Replication
 - 4 Failover Clustering
- 4 Conclusion

INTRODUCTION

Information makes today's businesses run. Often, businesses need access to important information 24x7. Applications that manipulate this data have become mission critical and thus storing and management of this data is critical.

In order to provide this mission critical data to employees and customers you must focus on three elements; Hardware, Software and Personnel. Each of these pieces is critical to achieving the level of availability that you need. More uptime requires more personnel and hardware resources and more robust software. Careful consideration must be used when determining how high a level of availability you must achieve.

HIGH AVAILABILITY

For our purposes High Availability (HA) is defined as how much the resource is available within a given time. This is generally expressed in a percentage of up-time which is how the expression 5 nines (99.999%) came about. Achieving this high standard requires the right combination of hardware, software and personnel.

The percentage availability can be calculated using the following formula: $\text{availability} = (\text{total time} - \text{total downtime}) / \text{total time}$. Therefore if your system is down for one shift (8 hours) per month you will achieve 98.9% uptime. ($98.9\% = (8760\text{hr} - 96\text{hr}) / 8760$). To achieve 99.999% up-time the total downtime in a year is less than 5 minutes.

Number of 9s	Percentage Availability	Downtime per Year
1	98.9%	4 days
2	99.0%	3 days, 15 hours, 36 minutes
3	99.9%	8 hours, 46 minutes
4	99.99%	53 minutes
5	99.999%	5 minutes

As you can see it gets more and more difficult to get to a higher level of availability. Keep in mind that downtime includes the time it takes to realize that the system has failed as well as the time it takes to get the system back up and running. This requires constant monitoring to ensure that failures are detected and repaired in the shortest time possible.

Good hardware and personnel are critical to providing up-time for your critical business systems. Hot swappable RAM and Hard Drives along with redundant systems assist in reducing maintenance downtime. IT personnel can ensure that secondary systems are available when primary servers require maintenance, service packs or other upgrades.

HIGH AVAILABILITY HARDWARE REQUIREMENTS

A major step toward high availability can be achieved by adding a robust storage solution into your environment. The use of a solid Network Attached Storage (NAS) device or Storage Area Network (SAN) is an excellent way to increase availability. Although many servers can be configured with a Redundant Array of Inexpensive Disks (RAID) the data becomes unavailable if the server fails. By moving the storage to a separate system the data can be accessed by many different servers. Keep in mind that the use of a NAS or SAN does not replace the need for regular backups.

State of the art server hardware can also minimize downtime and provide higher availability. Servers can be configured with redundant components such as power supplies and network cards to provide a more robust environment. Hot-swappable components such as Hard Drives and RAM allow for upgrades and replacement of failed components without shutting down the system.

To attain higher levels of availability redundant servers must also be used. These additional servers are used in case the primary server experiences a catastrophic failure. These standby servers are usually classified in three categories; cold standby, warm standby and hot standby.

Cold standby consists of a server configured similar or identical to the primary server. If the primary server fails a copy of the most recent backup from the production server is applied. When this process is completed the server is brought up on the network and users can once again use the system.

Warm standby consists of a server which has a copy of the database and is receiving updates via log shipping, synchronization or replication. In the event of primary server failure the warm standby server can be promoted to the primary server quickly. This method takes much less intervention than a cold standby server.

Hot standby is the most robust solution and requires the least amount of intervention. This is accomplished by using a failover cluster which automatically starts the database server on another node within the cluster if a failure occurs. The database is stored on a shared drive array or storage area network (SAN). Since the data is stored on an independent storage system the data is already up-to-date and database activity can resume as soon as the database server service has been started on another cluster node.

HIGH AVAILABILITY PERSONNEL REQUIREMENTS

Although hardware is a major part of a high availability solution someone needs to configure and maintain it. Good personnel are a critical piece in providing a high availability solution. A knowledgeable IT staff can identify hardware requirements, configure the network, monitor the solution and provide training to the end-users.

IT staff need to have a thorough knowledge of the hardware and software components of the high availability solution. The staff needs to be able to quickly recognize and remediate problems to ensure the downtime, remember that the time to identify that the system is down can be a significant portion of the actual downtime. Monitoring of the system can be automated; however, all failures will require human intervention to ensure the systems stability.

Documentation of all failure recovery procedures is also vital to the success of your high availability solution. Having thorough documentation on how the system is configured as well as procedures when components fail can greatly minimize downtime. Changes to the environment must also be well documented to determine their impact on the catastrophic failure recovery plan.

In some cases having a test environment identical to the production environment may be necessary. This allows the IT staff to test the impact of upgrades and patches to the system prior to implementing them in the production environment. By testing patches and upgrades the most efficient and quickest methods for upgrades can be determined.

ADVANTAGE FEATURES FOR HIGH AVAILABILITY

Advantage Database Server has many features that provide high availability to end-users all while preserving our core ideals of low administration and ease of use. These features include Server-Side Aliases, Online Backup and Replication. Each of these features can be used to preserve data and allow for access to this data in the event of a catastrophic failure. Additionally, Advantage can use the latest advancements in hardware and operating system features such as multi-processors and multi-core processors as well as 64-bit operating systems and various clustering solutions.

Server-Side Aliases

Traditionally, Advantage supported client-side aliases which allowed developers to specify data paths on each client. With the release of Advantage version 8, server-side aliases were introduced. Server-side aliases allow the developer or network administrator to define an alias for the data path on the server itself. Clients connect to this alias name instead of a data path.

This provides two major benefits. First, it allows for greater security since the data path is hidden from the end-users and it no longer needs to be shared on the network. Second, it allows for data to be stored on a NAS device. The server-side alias can be pointed to a NAS using a UNC path.

Online Backup

Backups have been a standard in the Information Technology world for many years. Most IT departments have a nightly, or more frequent, backup plan for all critical company data. These systems are very reliable and provide for recovery of data if a critical failure occurs on a server. However, most backup software can cause problems when working with data that is in use. Applications that must be used 24 hours a day can make backups problematic since the data is always in use. Additionally databases can be left open inadvertently, for example an application left running overnight.

In order to resolve these issues Online Backup was added in Advantage Database Server 8. The Online Backup works with both data dictionaries and free tables that can be accessed by the Advantage server. Online Backup creates a snapshot of the data and sends this snapshot to a specified location. This backup will be logically consistent and preserve all referential integrity whether the data is in use or closed.

When using data dictionaries differential backups can be performed improving the efficiency of the backup process. Differential backups identify changes since the last backup and apply these changes to an existing backup. This is different than an incremental backup which puts changes into a different location from the original (full) backup.

Differential backups are useful for very large databases where only a portion of the records are frequently changed. This reduces the amount of records that are backed up reducing the time required to perform the backup. In the event of a catastrophic failure the differential backup can be restored to the original location providing data current as of the last backup. In contrast, each incremental backup would have to be applied to the original backup to bring the data up-to-date.

Replication

Replication can be used to provide a real-time copy of data to another server which can be used for reporting or as a warm standby server which could be used in the event of primary server failure. By using replication to keep the standby server up-to-date with current changes if the primary server fails, the standby server becomes the primary server and clients connect to the new server and continue working.

For this solution to work you would start with identical databases on both servers. The primary server would be configured as a publisher and the stand-by server as its subscriber. Changes to the primary server would be pushed in real time to the stand-by server as long as a connection is available. Your application would have to contain an option for connecting to the stand-by server in case of primary server failure.

It is also important that the stand-by server only be used when the primary server is not available. The data would have to be copied back to the primary server after it has been repaired. No replication would occur while the primary server is unavailable.

Replication also provides a mechanism for providing updated data between locations. This ensures that all branch offices contain the same data and are kept up-to-date with all of the other offices within the company.

Failover Clustering

With the need for high availability to mission critical applications and data many companies are using cluster technology. Clustering allows multiple servers to work together to provide services on the network. Clustering also allows services to be moved between servers, or nodes, within the cluster in the event of failure.

Advantage Database Server can be used within a clustered environment to provide fault tolerance. Advantage can run on one node of the cluster at a time providing data to its clients. If the node fails the Advantage service will be transferred to another node in the cluster making it available for clients again providing a hot standby environment.

This structure provides many benefits to users as well as IT staff. Using a cluster provides a quick solution for a hardware failure since the Advantage service gets moved to another active node, allowing application users to restart their application and continue working. The Advantage service is moved automatically and requires no IT intervention.

Additionally, scheduled maintenance can be done without interrupting service to the end-users. Administrators can move the Advantage service from one node of the cluster to another. This allows IT staff to apply patches, upgrade hardware and other maintenance of the servers themselves without interrupting the end-users.

CONCLUSION

The need for a high availability solution is usually driven by your business needs. The ability to provide information 24x7x365 can be a daunting task. However, it can provide great rewards to your customers enhancing your business.

High availability is not an automatic process; it requires the right combination of hardware, software and personnel to make it work. The Advantage Database Server is capable of working with many of the latest hardware advances and provides many high availability features. Proper implementation and monitoring of these features is critical to providing the highest level of availability to end-users.

The need for high availability is driven by your customer's business needs. Keep in mind that not all situations call for maximum availability. However, if there is a need Advantage Database Server can provide the availability to keep your data safe and ready for your application and users.

iANYWHERE SOLUTIONS, INC.
WORLDWIDE HEADQUARTERS
ONE SYBASE DRIVE
DUBLIN, CA 94568-7902
U.S.A.

North America
Advantageinfo@iAnywhere.com
1 800 801 2069

Germany
ADS-team@iAnywhere.com
+49 (0) 7032 / 798 - 200

United Kingdom
AdvantageUK@iAnywhere.com
+44 (0)117 333 9000