

A Review of Availability Technologies

a white paper from



**LAKEVIEW
TECHNOLOGY**



A Review of Availability Technologies

Introduction

When it comes to critical business data and applications, it is all about availability. You can build or buy the world's most sophisticated business systems, but if you cannot reliably access them and their related data when you need them, they are worthless.

Obviously, you don't want to leave your business exposed to downtime. But things occasionally go wrong. Every once in a while, a major catastrophe shuts down an entire data center at some company. And, sooner or later, every organization will experience a sudden hardware or software failure that could shut down its operations, unless it takes adequate precautions. Yet, much more often, your business will be stopped or curtailed not by something going wrong, but by perfectly normal and necessary planned maintenance. Furthermore, once your business operations scale up to take full advantage of your system's speed and advanced-processing capabilities, attempting to continue business operations once they've gone down may be utterly impossible, even during planned downtime.

Availability doesn't just happen. Specialized technologies and comprehensive strategies are required to ensure that your data and applications remain available even in the face of disasters or hardware and software failures, as well as during inevitable planned maintenance. This paper looks at some of these specialized technologies.

In this discussion, it is important to recognize the difference between *technology* downtime and *business* downtime. The former is unavoidable. Even with 100 percent reliable components — an impossible ideal — planned maintenance will still leave some of the elements of your IT infrastructure unavailable. There is no way to prevent it. Thus, the primary goal of most availability technologies is not to prevent technology downtime, but rather to mask its effects in order to prevent it from resulting in business downtime.

It is helpful to categorize availability technologies based on their scope. Some are designed to avoid the business downtime that results from the failure of a single component of IT infrastructure. When you use an availability product in this class, your business will still be exposed to some downtime if an unprotected component fails.

The other class of availability technologies offers a more holistic solution. This type of solution doesn't just prevent a single component from causing business downtime; it protects the broader system — hardware, software and data included.

Single-Component Protection

For a system to work, all of its component parts must work. This section looks at technologies that ensure component availability in four broad areas: power, communications, data and processing.

Power

Hot-Swappable, Redundant Power Supplies

A power supply is the component within a computer that distributes electricity, at various voltages, to all of the other electrically powered components. Power supplies occasionally burn out. Some systems support redundant power supplies that can overcome this problem. With a redundant supply, even if one fails, the computer can still draw power from the other.

In this redundant environment, power supplies may also be “hot swappable.” If so, you can replace a defective unit without shutting down the system.

Uninterruptible Power Supply (UPS)

In the sphere of availability, UPS is not a logistics company, but rather the abbreviation for Uninterruptible Power Supply. Here, “power supply” refers not to the internal unit described above, but to the external power source.

In the simplest terms, a UPS is a battery backup system for your main-building power. The computer always draws its power through the UPS. If the external power to the UPS is cut off, the UPS continues to supply power to the computer from its batteries, with no interruption. How long the UPS batteries can support the computer before they run down depends on the capacity of the batteries and the amount of power being drawn from them. High-end UPS units are typically rated at from one to a few hours, but some can last as long as eight to 12 hours and may have swappable batteries.

Most professional UPS units include some ability to automatically communicate with your computer system to let it know that power is coming from the battery rather than from the usual external power source. Once alerted to the situation, the computer system software can then perform some action, such as saving open files, notifying users of impending shutdown, and shutting down gracefully, all before the battery runs out of power.

Backup Generators

For organizations where “shutting down gracefully” is not an acceptable option, a backup generator should be installed. A fully integrated power-management system should include sufficient automation to start the generator well before the UPS runs out of power.

Multiple Power Grids

It is rare that a single facility will be situated so that it can economically draw power from two different power grids, but there may be some circumstances where this is possible. A much more common usage of this approach occurs when a company establishes redundant data centers. It can then strategically situate the two centers so that a different power grid serves each center. Then, if one power grid goes down, the other will likely still be available. This dual-data-center approach falls under the “comprehensive protection” classification that is discussed below.

Communications

Redundant Network Paths

The way to protect against communication failures is to ensure that there are multiple paths between your servers and your users. For the internal portion of your network (your Local Area Network or “LAN”), or for dedicated point-to-point communication lines (such as T1 lines between distant offices or data centers), your IT department will have to ensure this redundancy is present. This may require installing redundant equipment within your local operations or the leasing of communications lines from more than one vendor. This latter precaution is important because, as proven by recent history, entire regional or national telecommunications networks from a single provider can go down, leaving their single-sourced customers with no data or phone lines.

For any portion of your communications that travel over the public Internet, such as data for EDI or other supply-chain management systems, this is much less of a concern. Massive redundancy is inherent in the structure of the Internet. In fact, the military, which was one of its originators, designed it to survive a nuclear attack. Nonetheless, the redundancy exists only within the Internet itself. Your local connections to the Internet are just as vulnerable as your local phone service. To ensure the availability of communications over the Internet, you must establish redundant connections to it.

Data

In an availability context, data protection generally refers to safeguarding data stored on disk drives, sometimes called Direct Access Storage Devices (DASD). Under this definition, data refers to anything stored on a disk drive, which encompasses more than what most people consider as “data.” Here, data also includes program code along with the information (or “system objects”) that the system needs to carry out its tasks.

Before going on, it is also important to understand that some storage technologies, although often discussed in conjunction with availability issues, do not by themselves increase data availability. These are storage architectures known as **Direct Attached Storage (DAS)**, **Storage Area Networks (SAN)** and **Network Attached Storage (NAS)**. A detailed discussion of these architectures is beyond the scope of this paper, but from a high level, DAS includes the non-networked storage that you probably use every day in your PC. A DAS device is a disk drive that is directly attached to a computer, either inside the case or external to it, but connected directly via a cable.

As the name implies, SAN is a network that is designed specifically for and dedicated to storage. It is similar to a LAN, but rather than linking computers, it links storage devices.

A NAS device can be attached to a standard network. The data on the NAS device can be addressed as if it were a shared disk drive within a standalone computer. The difference is that the “standalone computer” in this case does nothing but host a storage device.

DAS, SAN and NAS are simply different storage architectures that can be used alone or in combination. None of them have any inherent availability features included as an element of the architecture definition, but all of them can have availability technologies built on top of or beneath them within the architecture. The remainder of this section examines some of those availability technologies.

Redundant Array of Independent Disks (RAID)

RAID spreads enough information redundantly across multiple disks to allow any missing information to be recalculated in case of a disk failure. While RAID protects the data, it does not protect against the failure of other disk-related hardware, such as a controller, an I/O processor or a data bus. So, while your data may be recoverable, it may not be immediately accessible if any of those components fail.

Disk Mirroring

Disk mirroring, when properly configured, can eliminate some of those single points of hardware failure that RAID leaves vulnerable. Often used in combination with RAID, this approach requires that data be concurrently written to each unit in a set of identical disks, incurring minimal increase in CPU overhead or system complexity. However, mirroring all data requires at least twice as many disks. And the storage requirements will likely be more than double that of unprotected data because of the additional information that the mirroring technology needs to do its job.

Data Replication

Depending on its configuration, data replication and its end effect may be similar to that of disk mirroring. The primary difference is that while mirroring works at the disk level to create an exact copy of a disk drive, replication works at the data level to copy data from one location to another.

Because of its nature, disk mirroring typically requires identical or highly compatible disk drives on either side of the mirror. In contrast, data replication can usually replicate data between different makes, models and sizes of drives.

Also, using data replication, you can usually select which data is replicated, possibly omitting non-critical data to improve performance and lower storage costs. Disk mirroring, on the other hand, typically mirrors the entire disk without exception.

Tape Backup

Saving data to tape is a last line of defense for protecting data. Tape is a very -low-cost medium that can easily be shipped off-site to protect it from any disaster that might strike a data center. Yet, since tape is a relatively slow-speed medium compared to random access disk, it should be considered as a last resort for data recovery. RAID, disk mirroring, data replication or one of the comprehensive approaches described below would be preferred methods for recovering quickly should online data be lost.

Data Vaulting

In addition to long recovery time, the tape-backup regimen suffers at least one serious drawback – the potential loss of data updates applied between saves.

Data is normally saved to tape and then transported off-site on some periodic schedule, such as nightly. If you need to recover data from tape and, for whatever reason, do not have access to a journal (data log) of recent updates, any updates that were applied since the last tape-save job will not be available.

One way around this problem is to employ “data vaulting.” Using this technique, nightly saves are still performed and the tapes are shipped off-site to a backup recovery site. In addition, a network link is established with the backup site. Any updates applied to the production database at the primary data center are automatically captured and transmitted in near real-time to the backup location. There, the updates sit in an online queue, but are not applied to a database.

If you need to recover data, your database can be reconstructed from the latest tape and then the changes in the queue can be applied. Each morning (or on whatever periodic schedule), when a new tape arrives at the backup site, the previous day’s changes can be deleted from the queue, since they will be included on the new tape.

Processing

Fault-Tolerant Systems

Some systems are built specifically for fault tolerance. They include redundant components within the system enclosure, as well as the smarts to automatically switch to the backup component when the primary fails.

Since the point of fault tolerance is to provide High Availability, without any stoppages, these components should also be “hot-swappable” so that a failed component can be replaced without shutting down the system. Without this capability, redundancy is lost after the first component fails, leaving the system vulnerable to a second failure.

Clustering

Depending on the hardware and operating-system environment, clustering can either be considered as a single-component availability technology or a comprehensive technology. The generic definition of clustering is a group of two or more systems that are so tightly coupled that they appear to the end user as a single system. In a Continuous Availability cluster, software automatically detects the unavailability of the primary cluster node (server) and automatically switches to a backup node.

Some hardware and operating systems implement clustering in such a way that some resources, notably disk drives, are shared among the nodes in the cluster. While this protects more than a single component (each node contains at least a processor and its internal power supply), it does not protect the shared resources. For example, consider a cluster that shares disk drives among the nodes. If a drive fails (or enough of them fail simultaneously to overcome RAID and/or disk mirroring, if implemented), all of the nodes in the cluster will fail because they will not be able to access the data and applications on the drive(s).

In other hardware and operating-system environments, including IBM® eServer iSeries™, clustering is implemented using a “shared nothing” model. This means that *all* resources, including data and applications, are duplicated on the backup nodes in the cluster. Using this model, clustering can be considered one of the comprehensive availability technologies, which we will discuss next.

Comprehensive Protection

Protecting components is important, but, as Poor Richard’s Almanac (1758) noted, “*A little neglect may breed mischief. . . for want of a nail, the shoe was lost; for want of a shoe, the horse was lost; and for want of a horse, the rider was lost.*” Sometimes, it is the failure of the smallest, least-expected component that can cause you to lose access to your data and applications. A more foolproof approach is to implement a comprehensive protection solution.

Redundant Systems

Fully redundant systems are fundamental to any comprehensive solution. Rather than just protecting individual components, this approach duplicates entire systems – hardware, software and data included.

Duplicating the hardware is the easy part. Vendors are happy to sell you all of the computers that you need.

Duplicating the software is almost as simple. If you buy packaged applications, you can simply buy multiple licenses. If you build your own applications, copying them to a second machine is easy. Yet, since you will typically upgrade your software more frequently than your hardware, keeping it up-to-date might be a little more difficult, unless you use some form of automation. To ensure a true Continuous Availability environment, software versions on both your primary and backup systems must be synchronized. If you upgrade one and forget to upgrade the other, users may find that their applications don’t function properly when they are switched to the backup system.

To adequately protect your enterprise, application and system data must be replicated from the primary system to the backup on a near real-time basis since, in any active business environment, it changes continuously and with a high frequency. MIMIX[®] DB2 Replicator[™], from Lakeview Technology, provides an excellent means of doing this. It recognizes modifications in a DB2/400[®] database and copies them to one or more backup databases.

But availability depends on more than just database files. Business-critical information also resides in non-database files such as PC (IFS) files, data areas, data queues, spool files and image files. And, if user profiles, device descriptions, application programs and other system objects are missing or wrong on a backup system, your applications may not run properly – or possibly not run at all. MIMIX Object Replicator[™] recognizes modification to non-database objects and replicates those changes to a backup system.

In high transaction-volume environments, performance of the replication solution is critical, both to protect the performance of the production applications and to minimize data latency between the two systems. The innovative MIMIX Adaptive Caching[™] feature, developed in conjunction with IBM, heuristically predicts upcoming data requests and notifies OS/400[®] so that it can minimize disk bottlenecks and streamline access by caching data in memory. You can optionally select Adaptive Caching to optimize the tradeoff between memory use and throughput.

In addition, MIMIX exploits the most advanced OS/400 journaling features to deliver high-speed, low-bandwidth replication. It maximizes flexibility and performance by supporting Remote Journaling, Minimal Data Journaling, IFS Byte Stream Journaling and Data Area and Data Queue Journaling.

Replicating data and objects is not sufficient to protect your business. You also need a way to switch users from one system to the other when the need arises. MIMIX Monitor[™], a component of MIMIX, makes this possible by continuously watching the MIMIX replication processes and enabling transparent, automated failovers to a backup system in the event of production-system failures. It also facilitates fast manual switchovers when the primary system must be taken offline for planned maintenance.

MIMIX Monitor can monitor most iSeries parameters on an interval, scheduled or continuous basis. It can also determine if your systems are running off the normal electrical grid or are receiving power from a UPS. On detecting that a UPS is the active power source, MIMIX Monitor can take any number of actions, such as issuing warnings to users or initiating a switchover to another system that is still on the grid. You can also use timed events to take further actions if the UPS remains the active source for more than some predetermined time.

MIMIX also protects availability during maintenance operations. MIMIX Active Server[™] technology allows you to perform critical maintenance tasks on your primary system without affecting the applications running on it. **Synch Check While Active** allows you to compare two systems to ensure they are exact replicas. **Resynch While Active** facilitates the resynchronization of your primary and backup systems, when necessary. **File Promotion While Active** lets you restructure your databases in the background and then promote them to production, without halting your applications. **Reorganize While Active** allows for Continuous Operations while you reorganize your databases to clean up and reallocate unused space and optimize performance.

With MIMIX in place, any single component of the primary system – the processor, the power supply, the disk drives, whatever – can go down and the backup system will still be available to take over production operations. The same is true if you need to perform maintenance on the primary system. In addition, if the primary and backup systems are in geographically separated locations, you gain a very robust Disaster Recovery solution. Even if your whole primary data center is taken offline, the backup in another location will be immediately available.

Clustering

A shared-nothing cluster, as described earlier, is a variation on the redundant solution. The difference is that with the previous solution, third-party software assumes full responsibility for the replication, monitoring and switching processes. In a clustered architecture, the operating system assumes responsibility for more of the monitoring and some of the switching functions. However, in the iSeries environment, third-party software, such as MIMIX, is still required to perform the replication and cluster-management functions.

MIMIX has been a pioneer in clustering support for iSeries. And, beginning with MIMIX V4R4, that support is built right into the core solution.

Summary

Your business depends on the Continuous Availability of its business data and applications. Myriad components combine to deliver those data and applications. The failure of even the smallest component can rob you of the availability you require. Therefore, while protecting individual components is important, the greatest availability gains are derived from the implementation of a comprehensive approach, such as that provided by solutions like MIMIX.

This paper has reviewed availability technologies, but complete solutions require more than just technology. They require a disciplined approach to managing availability that embraces a multi-element methodology focused on attaining Continuous Availability of computer systems. Achieving “Managed Availability” requires a full, end-to-end solution that combines powerful software with business-process and systems design, change management, technical support and long-term services that cover the entire computing enterprise, including applications, data, servers, operating systems and infrastructure. This Managed Availability approach facilitates consistent, predictable access to any data or applications whenever or wherever they are required.

For more information on availability technologies and the Managed Availability methodology in general, contact Lakeview Technology or one of its worldwide business partners.