Availability and Disaster Recovery: Basic Principles

by Chuck Petch, WVS Senior Technical Writer

At first glance availability and recovery may seem like opposites. Availability involves designing computer systems and networks to prevent down time from failures, while recovery deals with recovering from failures, including failures caused by natural disasters. These topics are really two sides of the same coin.

There are a few key principles in designing systems for high availability and rapid disaster recovery:

- Design redundancy into all hardware systems to minimize the effects of failures and disasters.
- Plan backups so that backup data is always readily and quickly available. Live mirroring to disk provides the highest level of backup safety and the fastest recovery.
- Maintaining redundant hardware and backups off site guards against disasters at your primary site.

Hardware and data protection strategies built around these key principles can enable you to return your system to operation quickly when failures or disasters occur.

Why You Need an Availability and Recovery Plan to Protect Your Data

Many businesses and agencies don't spend much time designing their computer systems to make recovering from a failure or a disaster easy. Consequently, when hardware failures or disasters happen, it takes far longer to get their computer network back into full operation than it should. They may lose business, or they may even go out of business if the disaster is severe enough and the company's resources are limited.

It's easy to dismiss failures and disasters as unlikely, but they happen every day. Electrical failures and power spikes, broken water pipes or fire sprinklers, earthquakes, floods, fires, and simple equipment failures are commonplace.

At any moment you could find yourself replacing cables, server components, disk drives, network routers, power supplies, and trying to recover data as well. Any one of those repairs could take your system down until you can get them completed. Even worse, uncontrollable factors could slow your recovery by days or even weeks: What if parts aren't available quickly? What if your entire computing system is destroyed? What if a local disaster affects not only you but also your employees, and they are too busy putting their homes back together to worry about your computer systems?

Could your business survive, and could you keep your job if it took weeks to get your computer systems working again?

If you design your computer systems to be fault tolerant, many failures will cause you no down time at all.

Here's the bottom line: If you design your computer systems to be fault tolerant, many failures will cause you no down time at all. And even those that do cause down time will not be nearly so catastrophic.

Wouldn't it feel great to great to discover a failure and in the next moment realize your fault tolerant system design ensures users will never even know a failure occurred? They keep right on working, you keep your job, and your business goes right on being successful.

It all comes down to planning and making wise choices in how you configure your systems. You have to design your system so that you control it, and it doesn't control you. The first step is planning for availability and recovery and then implementing your plan.

Planning for Availability

“Availability” or “high availability” has been a buzzword in IT for quite some time. Often it describes a particular product designed specifically for applications where little or no down time can be tolerated. But high availability is something that can be designed into every component of your computer system and not necessarily at great cost, especially compared to the cost of a major failure.

When you begin to look at availability, remember that your purpose is to create a system using redundant components and data paths so that if one fails, you can switch the load over to the remaining components with little or no down time and little performance impact.

Begin by listing each component in your network, and asking yourself these questions:

- What happens if this component fails? Is there a backup component? What about its internal components? Do they have backups?
- If you have to replace a component or do routine maintenance such as a software or firmware upgrade, can backup components carry the load while you do the work? Will they be overloaded?

Write down the component and its backup on your list, or better yet, create a full network schematic showing all components and subcomponents. If any component or subcomponent on your list has no backup, that’s a danger point. If that item goes down, potentially some users or even your whole network goes down. You need to figure out a backup plan for that item immediately.

If any component or subcomponent on your list has no backup, that’s a danger point.

Let’s consider the components of your computing network and look at strategies to improve their availability.

Power

Every component you have in your computing system requires power, so wiring your power sensibly can do a lot to protect the whole system. It goes without saying that you need surge suppression and filtering for every device to protect it from damage by power fluctuations. If you can afford it, you’ll also want a UPS or multiple small UPS’s to keep your system running during power outages.

Connect similar devices to separate branch circuits.

It is especially important to connect similar devices to separate branch circuits. If you have multiple network switches, power them from different breakers. If you have a server with redundant power supplies, connect one power supply to one circuit breaker, and the other power supply to a different circuit breaker. That way, if one breaker trips, your server still receives power through the other breaker. This kind of redundancy can be a lifesaver for the systems on your network.

Servers

You can buy multiple inexpensive self-contained servers, or you can buy rack-mounted server enclosures that accept multiple server cards. The racked server system may cost more, but it makes expansion easy. You just slide in a new processor card, and you've got another server. On the other hand, if a disaster happens, such as fire sprinklers turning on in your computer room, you have all your servers in one rack, and they will all go down together.

For high availability, the ideal server scenario is to have at least two units, physically separated by distance.

For high availability, the ideal server scenario, whether you go rack mounted or self-contained, is to have at least two units, physically separated by distance. They should be at least in separate rooms, better yet in separate buildings, and better still in different cities. We'll talk more about this later in the disaster recovery section.

Buy servers with redundant components installed, especially power supplies and disk drives, which are prone to failure.

Buy servers with redundant components installed, especially power supplies and disk drives, which are prone to failure. Preferably get at least two of everything in each server. Get redundant processors, redundant power supplies, redundant disk drives, and redundant networking cards. The extra expense is well worth it to reduce or eliminate down time from failures.

Talk with the manufacturer about the best way to configure the system so that these components back each other up. For example, some servers, operating systems, and software can detect failures and automatically switch over from the primary component(s) to the hot running spare components. If the system is well-designed, the switchover is invisible to users.

Even if you can't afford a high availability server that can automatically detect failures and switch over to secondary devices, buying servers with built-in redundant components is a necessity. Even a manual switchover when the spares are already installed and running will take far less time than taking cold spares off the shelf and installing them, or worse, having to go out and buy parts on an emergency basis and install them.

Clustering software allows servers to share processing.

It's also a good idea to cluster your servers. Clustering software, such as Microsoft Windows Cluster Service, allows servers to share processing so if one server goes down, the remaining processors can pick up the tasks of the failed unit and continue operating. This provides excellent failure protection for your users. Just be sure no single server operates at more than 50% of its capacity so its workload does not overload the other server(s) in the cluster if they have to assume its tasks.

Network

Redundant network switches and data paths improve availability of your network. By using multiple paths through multiple switches, you ensure that if a path breaks, you still have alternate paths available. This is called multipathing, and you can buy specific software packages to manage a multipath network.

Redundant network switches and data paths improve availability of your network.

For example, let's assume you have two servers sharing a disk array through a storage area network. You could simply run one network connection from each server through a single switch to the array, but if the switch breaks, your entire network goes down, and if a network card goes down in one of the servers, that server goes down.

For better availability, put two network cards in each server. Then buy two switches and create two networks as shown in the figure below. Run dual connections from both servers through both switches to your storage system. This dual redundant network ensures that if either network card in a server goes down, the server stays on the network, and if one switch goes down, you still have paths available through the other switch, and your network stays up.

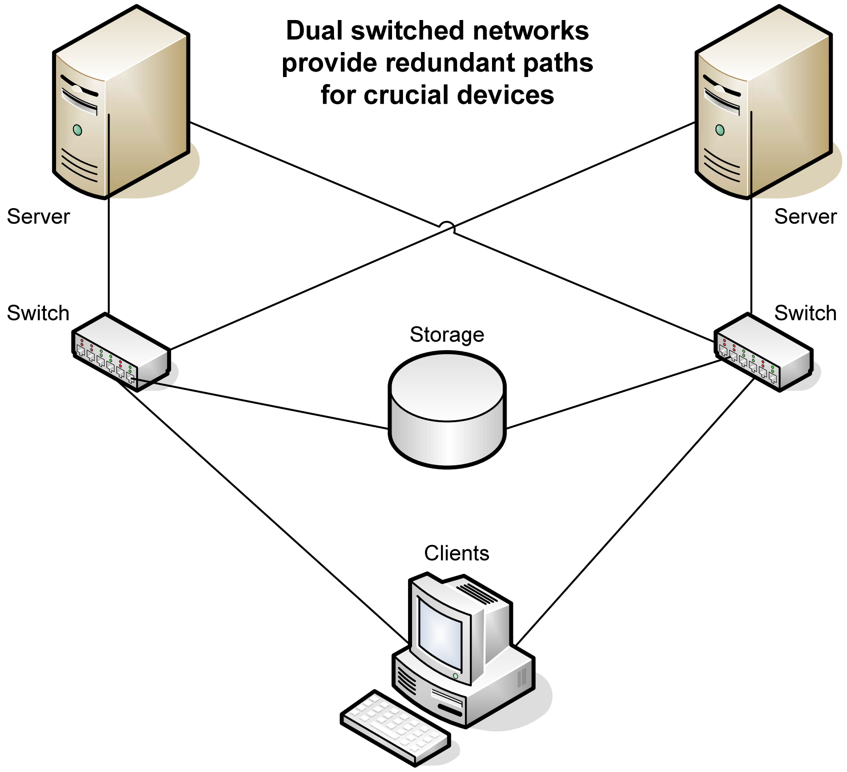

Dual switched networks provide redundant paths for crucial devices

A redundant network has further advantages for maintenance. If you want to take the switches down for firmware upgrades, you can take them down one at a time, leaving your network fully operational. Before doing this, though, be sure you don't have a workload heavier than 50% going through either switch so you don't overload the remaining switch when it handles all of the workload.

If economics prevent you from buying multiple switches, the next best thing is to buy a modular switch and route redundant paths through redundant line module cards in the switch. This protects you if a card goes down, but it won't protect you if the entire switch goes down.

On the other end of the network at the storage device, use multiple paths to separate ports on different network cards. As with the servers, this ensures that if one port or card breaks, you have another path available.

Planning for Recovery

What should you consider when planning for disaster recovery? Not surprisingly, many of the strategies you use to prevent down time from failures also enable you to recover quickly in the event of an unavoidable disaster.

- Data backups, local and off-site
- Redundant hardware systems, local and off-site

Data Backups

Snapshots of data taken throughout the day from the primary disks and copied to backup disks yield better protection because you back up data as users create it.

Any recovery plan begins with backups. You must have daily backups of all data at the very minimum, and you need to take the backups off-site to protect them.

You can do backups to tape, which is cheap but also slow and possibly unreliable because of frequent tape errors. Tape backups are easy to set up for automation during offline hours, but tape is notorious for errors as tape stretches and distorts with regular reuse. If you use tape, you'll want to test your backups frequently by doing a restoration from tape.

You can do offline backups more easily to disk than to tape, but buying extra disks is more costly than tape. Still, disk drives and disk arrays are getting so cheap, that it's hard to make a case for tape anymore. Again, you'll want nightly backups at the very least. More frequent snapshots of data taken throughout the day from the primary disks and copied to backup disks yield better protection because you back up data as users create it. You can buy snapshot software and schedule your servers to save periodic snapshots of their primary data to the backup disks.

Ideally, you can do live mirroring of your data to backup disks.

Ideally, if you have enough disks to give you adequate performance and space, you can do live mirroring of your data to backup disks. In this arrangement, you set up mirroring software to automatically copy every transaction to both the primary and backup disks. Alternatively, if you are using a disk array with RAID1 capability, you can set up the array to provide automatic live mirroring of data from primary disks to secondary disks within the array. Then if you lose data on a primary disk, you can quickly restore it from the mirror, or your server may be able to automatically switch over to the mirror if the primary disk or data fails.

The only disadvantage with snapshots and mirrors is the cost of disk space. You only get 50% efficiency from your disks because the same data occupies two separate and equal disk spaces. To overcome this, many disk arrays also offer other RAID levels that do not use as much disk space as mirroring but still offer good protection against lost data and disk drive failure. For example, a disk array configured for RAID5 writes data only once to a set of primary disks, but it writes parity information with the data. If a disk fails, you can use the disk array software to reconstruct the failed disk from the remaining data plus parity. In this way you get greater space efficiency than mirroring but you still have excellent protection for your data.

If you need the highest possible availability, use live mirroring with automatic switchover in case of primary data or disk failure.

If you need the highest possible availability, use live mirroring with automatic switchover in case of primary data or disk failure. With a backup constantly being made and constantly available, your data is well protected.

The difficulty with all of these backup arrangements is that you must get the backups off site to protect them from a potential disaster at your main site. This means you have to set up a data center separate from your main site and perform tape or disk backups at the remote location. If your company can't afford these separate facilities, a reasonably priced alternative may be to find a computer systems vendor who will provide backup services for you.

A computer systems and services vendor may be able to offer you a cost-effective set of remote backup alternatives:

- Accepting tapes for storage in their fireproof vault
- Performing nightly backups over the internet from your machines to their tape or disk drives
- Providing full live mirroring of data from your machines to theirs

Before you spend the money to set up your own remote data center, call your systems vendor to see if they can offer you a cheaper alternative.

Redundant Hardware

Previously in the availability section we talked about judiciously using redundant hardware to ensure you can recover quickly from a failure. But what if you were to locate your redundant hardware off site instead of in the same building? That's the essence of a hardware disaster recovery plan.

Just about everybody knows you have to back up your data, but backing up your hardware can be just as important.

Just about everybody knows you have to back up your data, but backing up your hardware can be just as important. It will do you no good to have your data backed up off site if your computer systems have been destroyed by a disaster. You would have to set up a new emergency computing center at some other location, essentially duplicating your entire computing infrastructure at a moment’s notice. Not easy, not fast, not cheap, and not a pretty picture.

If your vital servers and hardware systems are replicated across town, or preferably in another city, you are ready to handle any disaster

On the other hand, if your vital servers and hardware systems are replicated across town, or preferably, in another city, you are ready to handle any disaster. Your data is backed up or mirrored on the remote hardware, and the computing infrastructure is in place so that all you have to do is switch over to handling all of your computing at the remote site. This might typically take a few minutes to a few hours, depending on your setup, but it surely beats the days or weeks it would take you to throw together an emergency system from scratch after a disaster.

Again, you may not be able to afford to set up your own complete duplicate computer room in another location, but you may get the services you need from a computer systems house, for example:

- Pay a small annual fee for the option to temporarily use the vendor’s space and equipment if yours fails.
- Lease part time access to space and equipment from a vendor. For example, the vendor could do your nightly backups at his site using his equipment.
- Lease full time access to space and equipment from a vendor. For example, you do your remote snapshots or live mirroring to the vendor’s equipment.

Making a List and Checking it Twice

No matter what your hardware and data protection plans are, be sure you document them and test them.

No matter what your hardware and data protection plans are, be sure you document them and test them. In an emergency, you will need to know what equipment and data is stored where and who is responsible for each. Your plan should include at least these things:

- A written data backup procedure with task assignments
- A list of hardware and their locations or a network schematic showing each piece of hardware, its location, and the location of its backup
- Architectural drawings showing wiring and physical building layout
- A list of software licenses and license keys
- Data maps showing where all your primary and backup databases are stored
- Disaster recovery procedures with task assignments and a calling tree. Each person should know whom to call and what equipment and data are their responsibility to restore in what priority order.
- If you need help planning and writing backup and recovery procedures, WVS can help. Make sure all team members take document copies home, and store copies at your off-site location.

Periodically schedule a test of your recovery plan. If your documents are outdated or your team doesn't know what to do, your plan will fail. Even your data backups may not be trustworthy if you don't test them regularly to make sure you can use them to restore your data.

With good data backup and hardware restoration plans, your business will be able to tolerate routine equipment failures without incident. In the event of a disaster, your business could be one of the great survival and recovery stories.

Summary of Recommendations

- Design redundancy into all the components of your computer systems
- A UPS and filtering are standard requirements. Wire redundant power supplies to separate breakers.
- Set up redundant servers, preferably in different locations.
- Use clustering to enable servers to share workload and multipathing to ensure network resiliency
- Set up as much off-site data backup and hardware redundancy as you can afford. Vendors may offer cost-effective backup services.
- Create and regularly test a written recovery plan to restore hardware, software, and data in an emergency.

Call us when you need technical documentation

We're experts at technical writing of all types, including hardware and software documentation, marketing literature, white papers, procedures, and web content. We specialize in making complex information easy to understand and use.

Phone

Voice 1.530.265.4705
Toll-Free 1.800.529.9907
FAX 1.530.478.1387

Email and Internet

info@wvswrite.com
www.wvswrite.com

Address

WVS – The Technical Writing Company
552 Brock Road
Nevada City, CA 95959