Linotype-Hell

Technical Information

Fault Tolerance

As prepress becomes fully electronic, how we handle our digital data becomes extremely critical. Loss of data, by whatever means, is a stunning setback to any project. To avoid this kind of catastrophe, many computer users build in safety nets of various kinds. These techniques range from the simple precautionary techniques to extremely expensive hardware solutions.

While data may be lost through a computer virus or an unintentional deletion, the intent in this document is to cover losses that occur due to equipment failure of one kind or another.

Non-stop computing

Non-stop computing (also called fail-safe computing) describes the ability of a system to continue to function even when any hardware or software component within the system fails. Fault tolerance may also be used in this sense, although in general it provides something short of non-stop computing.[1] For example, fault tolerance may only protect a single component of a system like the disk drive. If the disk drive fails, a backup will take over. However, if something else fails (for which fault tolerance is not built in), then the system will fail.

Non-stop computing systems, as you might imagine, can be incredibly expensive. Thus users often settle for redundant, parallel or duplexed systems that build in an acceptable level of fault tolerance.

Redundant, parallel, & duplex

The meanings of the terms redundant, parallel, and duplex are overlapping. To confuse the matter, these terms may mean one thing when referring to one concept and something entirely different when referring to another. (See box below.)

In the purest sense, a duplexed component in a computer system, is one that waits on standby while its partner performs the function. If the partner fails then the duplexed component takes over. In a parallel system, both components operate concurrently. In the event of a failure, the component that didn't fail can continue system operation and provide access to critical data. In a redundant system, either the data or a component is duplicated. For example, the data may be written out twice in two different locations, or, there may be two pieces of hardware such as the disk controller.

***Duplex** – One waits, the other takes over when necessary. Requires a manual intervention in the event of a failure.*

***Parallel** – Both work at the same time. In the event of a failure, the working component takes over the work of the other. This happens automatically.*

***Redundant** – Function is duplicated, two components do the same thing, but not at the same time. Usually requires a manual intervention in the event of a failure. (Exception: Some RAID implementations.)*

[1] Some authors define fault tolerance very strictly: "the data is always available and always correct." However, in a Macintosh and PC (personal computer) environment, the term is used much more loosely.

Both duplexed and redundant systems usually require manual intervention. Manual intervention requires users to do something (reboot for example) to get back up and running while automatic systems require nothing from the user. A parallel system would be automatic. A duplexed system could be up faster than a redundant system. Price is a factor in the decision to choose redundant, parallel or duplex, but these questions are also important:

- How long can the user afford not to have access to their programs or data?
- How much data can you afford to lose entirely?
- How much downtime can you afford waiting for a replacement part?

Typically, these redundant systems have two of everything but still require some type of failover[2] from the failed component to the backup system. This failover may or may not require manual intervention and may or may not allow for data loss depending on the level of redundancy and the integrity mechanisms in the software.

UPS

One specific form of fault tolerance is an uninterruptible power supply (UPS). A UPS is a battery backup system that protects against loss of power. When the normal electrical power goes out, the battery backup contained in the UPS supplies power. Most UPSs also supply power-line conditioning and surge protection.[3]

A UPS can protect anything that is plugged into it, as long as it can meet the power needs of the device or devices. For this reason, many UPS manufacturers discourage attaching a laser printer to a UPS because of the amount of power the laser printer draws on start-up.

Since a power outage normally knocks out all the power in a building, many users choose to use a UPS only on the file server. This allows the file server to shut down gracefully (i.e., close out the file system properly and also notify any users who are still up that the server is shutting down). Most UPS manufacturers offer software that automatically performs this function by communicating with the file server through a serial connection.

UPS support ranges from minimal (enough power for a graceful shutdown of the file server) to extensive (where all critical workstations are equipped with a UPS with enough battery backup to run for a long period of time). A standby power supply (SPS) provides slightly less service than a UPS. While a UPS runs continuously running from the power supply's batteries, an SPS doesn't kick in until the power goes off.

Disk storage

Another area where fault tolerance is commonly built into a system is in hard disk storage. Three methods for disk fault tolerance include disk mirroring, duplexing, and disk arrays (See below.) With any type of fault tolerant disk storage, the actual data redundancy or backup can be performed in hardware or software. Hardware solutions offer considerable performance improvements over software, but they are also more expensive.

Disk mirroring – Disk mirroring is a common form of data redundancy, whereby data is written out to two identical drives simultaneously. Data is constantly checked to ensure its accuracy. Then, if one of the physical drives fails, the system keeps running with all data intact. The bad drive can be removed and replaced at the user's convenience. This type of protection is normally found only on file servers primarily because of the cost involved in providing this level of fault tolerance.

Disk duplexing – Duplexing is an extended form of disk mirroring where each of the two mirrored physical disk drives is connected to a separate disk controller.[4] This offers additional protection over a single controller

[2] The term failover is a combination of the words failure and switchover. It describes the process of picking up on the function of the failed component. Warm-failover is failover that does not require human intervention. Cold-failover does require human intervention. Here is a non-computing example: warm-failover is like a generator that starts automatically when the power goes off, cold-failover is like a generator that you must start on your own.

[3] Power line conditioning and surge protection are topics that go beyond the scope of this document. A good source for information on this topic is *PC Power Protection* by Mark Waller. This book was published in 1989 by Howard W. Sams & Company, Indianapolis Indiana.

[4] Disk controllers are either integrated into the motherboard or are separate cards. You typically see them on motherboards on Macintosh computers and some PC workstations. File servers almost always have separate cards because they use large capacity drives which benefit from high performance controllers.

environment in that should the controller fail, you still have a separate controller providing access to your mirrored data.

Disk arrays – RAID (Redundant Arrays of Inexpensive Disks) is the newest form of data redundancy. Data is distributed through a disk array containing two or more drives. The data from one disk is replicated on the others. When a drive in the array fails, that drive can be physically replaced and the data rebuilt because enough of the data is duplicated mathematically so that a single drive failure does not cause loss of data. Some systems allow the drive to be replaced and the data rebuilt while the system is operational. Depending on the particular product, it may offer redundant fans and/or power supplies. The power supplies are each capable of keeping the system running. This protects the system against failure of the power supply, but not from power failure, that would require a UPS.

There are many levels of RAID. Some of the most commonly mentioned ones are RAID 0, RAID 1, RAID 3, and RAID 5. RAID 3 is particularly useful for applications that use large images.

Fault tolerant processors

Some technologies such as RISC[5] computers offer fault tolerance through additional processors. These processors are not necessarily redundant but instead are symmetric or parallel. Still, if one goes down, another processor can take over (but not necessarily automatically). Some network operating systems (NOS) such as Novell NetWare® SFT (System Fault Tolerance) provide this capability through the use of two physical PCs. If one file server goes down, the other is still functional.

Fault tolerant networks

In a Thinnet (also called Cheapernet) network, there are typically two areas problem areas: cable breaks and network interface cards (NICs). Fault tolerance can be built into a Thinnet network via multiple NICs. Then if a single network segment goes down, the other NIC provides a way to get those down nodes up and running quickly.

Recently, 10-Base T networks have been developed which have intelligent hubs connecting all nodes in a star configuration. If the cable for a particular node is damaged or disconnected, it does not affect the other nodes on the network.[6] Many 10-Base T manufacturers support SNMP (Simple Network Management Protocol) which enables the system administrator to monitor the network from a single station. SNMP management software enables the system administrator to be warned of a failure as it happens. It also identifies the problem as a network problem.

The concept of a hub is important here. A hub is the center of a star network. All nodes plug into the hub. Hubs can be passive, active or intelligent. *Passive* hubs provide interconnection between all devices. *Active* hubs

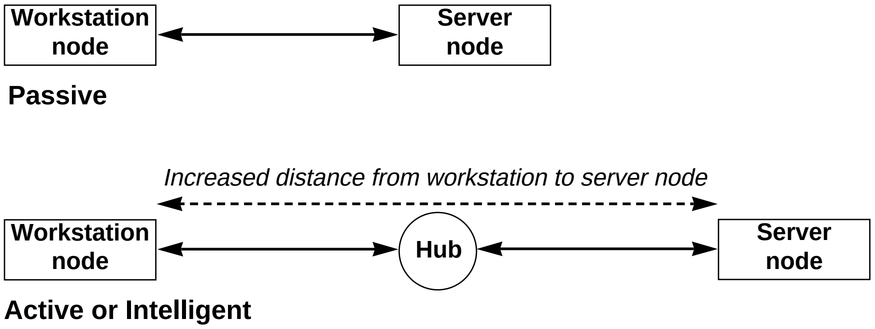

Figure 1 – An active or intelligent hub can effectively increase the distance possible between a workstation node and a server node. The actual distances depend on a number of different factors including the network software and the type of wiring.

[5] RISC stands for Reduced Instruction Set Computing.

[6] While hubs can help prevent problems due to damaged cables or nodes, if a hub goes down, all the network legs leading into the hub will go down.

require power and increase the signal in a similar manner to a repeater. This means that nodes may be separated by a greater distance. (The maximum distance is from the hub to the node instead of from one node to the other, see Figure 1 on the previous page.) *Intelligent* hubs require power and not only increase the signal but also know where the packet is going so that it can broadcast only to that node. All 10-Base T networks require hubs. A Thinnet network may or may not have a hub.

LinoServer

Regarding the Linotype-Hell product line, UPS and HADA are offered as options on LinoServer (HADA stands for High Availability Disk Array and it supports RAID). Linotype-Hell also supports multiple Ethernet adapters and additional central processing units (CPUs) for LinoServer.

LinoServer with HADA also offers a technique called Automatic Data Reconstruction (ADR), which provides fault tolerance in case of a hard disk failure. With HADA, if a hard disk fails, the system does not have to be shut down to replace it. In addition, the data that was on that disk will be automatically reconstructed. With HADA, the ADR feature automatically reconstructs the data onto a new disk. This happens transparently to the users of the system. There is no downtime due to disk failures. Within LinoServer, the HADA feature is referred to as the CLARiiON™ Disk Array.

Another feature of LinoServer is Error Correcting Code (ECC). ECC is a method of storing more than a single bit of parity. (A single parity bit is usually attached to each byte to detect errors in transmission.) ECC tells you if you have an error and how to repair it.

Conclusion

These are the key techniques used in achieving a fault tolerant installation. Of course price and production realities also play a critical in determining the appropriate level of fault tolerance for a computer system. Ultimately, the question comes down to the level of importance you place on keeping your system operational and your data secure, intact, and available.

In a future article, a related topic will be considered: data protection. This soon-to-be released article will cover backup procedures, computer viruses, and data recovery.

Acknowledgements

The material for this technical information article was submitted by Eugene O'Brien of Linotype-Hell Company. Many thanks also to Mark Astmann, John Flaherty, Marco Mularoni, and Bruce Roth of Linotype-Hell for their help in producing this article.

Please direct any questions or comments to: Jim Hamilton, Marketing Department, Linotype-Hell Company, 425 Oser Avenue, Hauppauge, NY 11788
(For subscription information on the Linotype-Hell technical information series, please call 1-800-842-9721.)

November 1993, Part Number 7023