Avoiding Disaster – "An ounce of Prevention"

White Paper

Published: May 2001 (with revisions)

Proactive Steps to Minimize Power Outages

Contents

Preface	1
Introduction	1
Availability – How much is enough?	2
An Ounce of Prevention is worth a Pound of Cure	3
The UPS	3
Redundant Power Supplies	5
Connecting the Redundant Power Supplies to Redundant UPS's	6
Failure Happens	6
Good Backups = Survival	7
Do you still need backups if you have a SAN?	7
Common Sense	7
Let Us Help You Succeed!	8

Preface

You are sitting at your desk, working intently on a critical project at your computer and the power goes out. There is no worse feeling than that – except of course if you are responsible for your company's enterprise computer systems and network, and you are not completely sure just how well protected you are from a power failure. This white paper will help give you some insight on how to avoid power failure anxiety.

Introduction

At the time of the writing of this white paper many areas of California are experiencing "rolling blackouts". A rolling blackout is an unexpected period of 30 minutes or longer when electricity is disconnected to targeted areas in order to reduce the total load on the state's power grid. So far there has been little or no advance warning regarding which areas will be affected next and therefore it is impossible to adequately plan or prepare for these outages on short notice.

The purposes of this white paper are to document proactive steps that can be taken to minimize the impact of unplanned service outages, and to discuss some elements of general systems availability. There is a certain element of overlapping inherent in these subjects, so the reader may find that some sections of the paper are useful in contexts other than that of a power outage.

While a power outage is not by definition a true disaster, it shares many of the same issues with a Disaster Recovery Plan. When it is not addressed, there can be a loss of services and data, and possible equipment damage, resulting in a longer unplanned services outage. Many companies incur serious loss of business, profits, and customer satisfaction with even short periods of unplanned services interruption.

Rolling blackouts were the inspiration for this document, but it should be noted that issues such as the loss of power are not limited to California. The impact is widespread and potentially devastating. It is good business practice to prepare for problems such as this in order to avert situations that could potentially lead to the execution of your Disaster Recovery Plan. As the saying goes, "an ounce of prevention is worth a pound of cure".

Availability – How much is enough?

The growth of the Internet has forced many companies to change the way they do business. In fact, most companies today are affected in some way by the Internet's rise in popularity as a medium for communications, sales, on-line transactions, and customer service. As a result, companies are finding that their old standards for systems availability are inadequate to meet their current needs.

As most of you are aware, the "nines" scale of availability has long been a standard for measuring not only the amount of time in which a system must be available (uptime), but also the amount of time in which it may be unavailable (downtime), for maintenance or unplanned outages. Many companies pay a premium for their systems and attendant maintenance in order to be guaranteed high availability. The more "nines" in the rating, the higher the level of availability and conversely the lower the allowable downtime. The actual numbers may surprise you. Per non-leap year:

- 90.0% uptime = 36 days and 12 hours of downtime per year.
- 99.0% uptime = 3 days, 15 hours and 36 minutes of downtime per year.
- 99.9% uptime = 8 hours and 46 minutes of downtime per year.
- 99.99% uptime = 52 minutes and 33 seconds of downtime per year.
- 99.999% uptime = 5 minutes and 15 seconds of downtime per year.
- 99.9999% uptime = 31.5 seconds of downtime per year.

It is important to differentiate between scheduled and unscheduled downtime, as well as service times. Often times it will be decided that a scheduled window of downtime on a periodic basis is acceptable and therefore does not factor into the "nines" equation. Also, many businesses are not 7 x 24 operations so there may be non-critical times identified during off-hours. Each business is different and it is important to determine what matters to your specific business, as opposed to jumping on the "five nines" bandwagon.

Obviously few systems will achieve 99.9999% uptime, although "five nines," or 99.999% is often achieved in telephone and banking systems. While your company may not require such a high level of availability, it also probably cannot survive with 90 or 99% uptime, either. Even 99.9% will be problematic for companies that rely heavily on their computer systems in the performance of their day-to-day business functions. It should also be noted that costs associated with a "high nines" system can grow almost exponentially as you extend beyond the 99% range.

An interesting note is that in this case, while availability refers to the uptime of the system, reliability means something completely different. Reliability refers to the likelihood of failure. In turn, a reliable system may ensure high availability, but when failures do occur, availability can also be affected by serviceability, which is a measure of how easily failures can be corrected. This can be more problematic in older and/or less popular equipment.

An Ounce of Prevention is worth a Pound of Cure

You can utilize many of the same techniques that companies requiring uptime levels of 99.9% or higher use. This helps to ensure that your business does not suffer needless loss of service or hardware failure in the event of power outages. Let's explore some of the areas where systems and practices can be reinforced to prevent unnecessary downtime:

The UPS

Preparation and redundancy are key to availability. While redundant servers are a commonly used preventative measure to ensure that processing is not completely halted by the failure of one machine and raid arrays help to eliminate the problems caused by single disk failure, a whole cluster of servers and their attached raid storage will be useless if their electrical supply is disrupted for any reason.

Most companies use some form of uninterruptible power supply, commonly referred to as a UPS, to support their computer systems. These range from small battery backup units that will power a single desktop computer long enough to allow work-in-progress to be saved, to large industrial units that support complex server farms and allow for an orderly shutdown of systems to prevent data loss and long restart periods. These devices may also provide another critical function, which is to condition the incoming mains current to prevent dangerous surges and spikes from reaching sensitive equipment. If a UPS is of a type that does not do this then it may also require a few milliseconds to begin powering the equipment connected to it. In that case it is imperative that the power supplies in the dependent machines maintain an equivalent number of milliseconds of power in the event that mains current is removed from them. It would also be highly advisable to provide some form of power stabilization to protect delicate equipment if the UPS does not provide this.

An often-overlooked point of failure is the UPS itself. If only one is in use and it fails the situation will be the same as having no UPS, possibly worse given a false sense of security. Redundant UPS’s are critical to systems that require high availability. Having redundant UPS’s are important, but so is having “cross wired” power supplies. The best configuration that we have seen is having two UPS’s running at less than 50% of capacity providing power to each and every piece of equipment. If one UPS fails then the other has both the capacity and the connections to sustain business operations (albeit for a reduced period of time). This is described in more detail below.

Even redundant UPS’s can fail so it is good practice to test every UPS on a regular basis. The most common failures are caused by the increasing inability of the UPS’s batteries to hold a charge as they age. Frequent testing will help to uncover this problem. In addition, many battery systems benefit from the occasional discharge and recharge cycles achieved during the testing process. Temperature can also affect battery life and capacity.

Size matters. A common error is to under-size the capacity of the UPS. Calculating the capacity of the UPS required by your installation is outside the scope of this paper, but manufacturers of these devices usually have clearly defined recommendations based on the power requirements of the equipment to be connected to them. It is far better to have surplus capacity than to have too little. Under sizing can result in damage to the dependent equipment or failure of the UPS to function at all, and the cost of replacing inadequate equipment is usually higher than the cost of purchasing oversized equipment to begin with. Also, be sure that the UPS you select automatically shuts off when it can no longer maintain adequate voltage and current levels, otherwise sensitive equipment may be severely damaged by decreased voltage and/or current as the UPS’s batteries weaken.

Depending on your installation, you may want to ensure that the UPS you purchase possesses an interface that allows it to communicate with your systems to instigate an orderly shutdown in the event of a power outage. These devices may connect directly to a server or to the network, and software installed on the dependent devices will begin a shutdown when signaled by the UPS that it is operating on battery power. Often this software will generate warning messages to the connected users to save their work and log off the system to avoid loss of data. An advisable additional option is paging support. With this option installed, pages will be generated and sent to designated personnel alerting them that an event is taking place. An example of when this might be a very desirable feature is if a power outage occurs during the night, personnel responsible for operations would be paged with the message that the system is shutting down. In this way they would be alerted that the system must be restarted when the power event is ended, thus eliminating otherwise costly unnecessary downtime once the power is restored.

Network appliances such as switches and routers are often overlooked when it comes to provision of a secure supply of power, especially those that are installed in an out-of- the-way place like a utility closet or on someone's desk. However, loss of power to these devices can wreak havoc. If a switch or router loses power, an entire segment of the network may fail. Heavily trafficked routers on large networks may require long periods of time to rebuild their routing tables and become functional again when power is restored, or worse still, may remember and propagate no longer current routing information causing wide spread network corruption and failure. It is important to remember these devices when designing your UPS installations.

Redundant Power Supplies

Another commonly used technique is that of redundant power supplies. A machine, such as a server or router, is critical to operations and may (or should) have multiple power supplies installed in it. These may be "hot" supplies that are always active and contributing to the total power needs of the machine, and the failure of one supply will not cause the machine to fail, or they may be "warm" supplies, meaning that they are connected and operating but only one of them is supplying the machine with power. So-called "cold" supplies are also used in compliment to either "hot" or "warm" supplies. These are installed, but not connected or supplied with mains current. They primarily shorten the maintenance time required to replace a failed supply, because they can be switched or connected into service, often without shutting down the machine they help to power.

Many enterprise level machines have internal monitoring systems that alert their operators to fault conditions or to impending faults. In the case of systems with redundant critical components, the operator or maintenance engineer may be able to take action to correct the fault without shutting down the machine. Some of these machines have separate power buses for the processor and disk array sections, allowing the processor section to be shut down for maintenance while the disk array is still accessible by other machines. These features may be important considerations in the purchase of a particular machine.

Connecting the Redundant Power Supplies to Redundant UPS’s

A further level of security can be provided by connecting a machine’s redundant “hot” power supplies to redundant UPS’s. In this way you can ensure that if one UPS fails, the other will automatically continue to power the machine, unless of course the power supply it is connected to fails as well. Critical enterprise machines often overcome this possibility by making the redundant power supplies themselves internally redundant, meaning that even if one of the redundant UPS’s fails, and the power supply that the remaining UPS is connected to suffers a fault, the internally redundant power supply will continue to function.

Is all of this complexity necessary? That depends on how critical uptime is to your organization. Some good questions to ask when assessing what level of uptime you need are:

1. How dependent is my organization on its computer systems to perform its business functions.

2. How much downtime can my organization withstand before it causes an unacceptable level of lost productivity, profitability, and customer satisfaction.

3. Are there well defined and documented procedures in place for manually performing our business functions in the event of prolonged downtime?

4. How does the cost of downtime compare to the cost of high availability?

Keep in mind that the period of time the electrical supply is actually off is not by any means the total time that your systems may be unavailable. A server, disk array, or database may require hours to come back on-line after an improper shutdown. Network appliances may require protracted periods of time to renegotiate and restore communications. Individual workers may have lost unsaved work. All of which may seriously impact overall productivity and have a tremendous associated cost.

Failure Happens

In spite of all good intentions, planning, and provisions for prevention of unplanned downtime, it does happen. Even the most redundant systems may fail: multiple power outages occurring before the UPS’s have a chance to complete their recharge cycles may cause them to fail to support their dependent systems long enough for an orderly shutdown sequence to be completed; mains events such as surges and spikes caused by the sudden loss of current may damage circuit-breakers and/or UPS’s beyond functionality, causing sudden loss of power to dependent machines.

Good Backups = Survival

In the event that a failure occurs, especially involving damage to equipment, data may be irretrievably lost. This could be catastrophic for a company. In this case the best cure again proves to be prevention, in the form of regular periodic backups performed on a schedule which is strictly adhered to. If data is lost due to a machine failure, it can be restored to the point at which the last backup was made, and with journaling techniques it may be restored to a point just before the failure occurred. Don't forget to verify your backups! Few things are more disheartening than to begin to restore a critical database only to discover that some or all of the backup tapes are corrupt or unreadable.

Do you still need backups if you have a SAN?

Yes! We were surprised to find that a common misconception is that mirroring data in a remote location is adequate protection against loss of data. Unfortunately this is not true. While it may certainly be useful in the event a disaster occurs on your main site, it must be remembered that most disasters do not happen in milliseconds. Instead they often happen over a period of minutes, and data that is corrupted on the main site may be mirrored on the remote site, in some circumstances rendering all of the mirrored data unreliable. Additionally, mirroring will not help the unintentional or deliberate loss of data. Just think "point in time recovery".

In the event of a sudden power loss there probably won't be much corruption of data to be mirrored on the remote site, as long as the whole main site fails at one time. However, if some machines fail and others do not a fascinating variety of interesting and unpredictable events may occur, many of which could corrupt data. Regular backups still remain the best defense against the unexpected.

Common Sense

An often overlook commodity when preparing for such events as rolling blackouts is good old-fashioned common sense. Think carefully about the events that will take place if the electrical supply to your place of business is interrupted, especially for a protracted period of time. Once your computer systems have been assessed, think about other aspects of your business. Will there still be telephone service so that customers can contact you to place telephone or fax orders or ask for information? Does your PBX require power to pass calls through? Will there be adequate light to see by - including in interior hallways and offices? Is there a possibility that someone may be trapped in an elevator? These are just a few of the many possible issues that should be addressed.

Take the time to plan for eventualities. Planning and prevention can eliminate or reduce the problems that your company may face during a power outage, as well as during many other types of unexpected events. An "ounce of prevention" can help ensure that such events don't become disasters. Please see our other White Papers on Disaster Recovery Planning and Project Management at http://www.comp-soln.com/whitepapers.

Let Us Help You Succeed!

Call today to discuss ways that Comprehensive Solutions can help your organization save money and achieve better results with your IT projects. We provide the ***confidence*** that you want and deliver the ***results*** that you need.

Download our “Disaster Recovery” Brochure

Back to White Papers
Back to Services

Comprehensive Solutions
4040 N. Calhoun Road
Suite 105
Brookfield, WI 53005
U.S.A.

Phone: (262) 544-9954
Fax: (262) 544-1236