



# How to minimise your computer network problems

Most people rate *'reliability'* as one of the top three characteristics they want in their computer network. In the modern business world, a trouble-free computer network is a "must have", not a luxury.

Each computer network is comprised of a number of components. To consistently deliver a smooth running network, all components need to be working in harmony. In this guide we'll focus our attention on the heart and brains of your network - the server (sometimes referred to as the 'file server').

The server is the centre of the computer network. Typically it is the storage area for all management data files and other important information. The server is also frequently used to store the organisation's e-mail messages.

An unhealthy server can affect the stability of the computer network by causing "crashes", which can disrupt the efficient operation of the entire office and put at risk your valuable data.

## Server location

Ideally the server should be kept in a secure location (a locked office) to minimise the likelihood of unauthorised access via the server console. The console should remain locked at all times when not in use, and only staff with a "need to know" should have the passwords of administrator-level functions. (Generally this should be limited to one or two staff plus the network service-provider).

An old saying that is very applicable to the computer environment is that "An ounce of prevention is better than a pound of cure". It is important that the server be located in a well-ventilated area, where the ambient temperature should be kept between 23 - 25 degrees Celsius. This is an aspect of server housekeeping that is often overlooked, but time and again we have observed that those sites that keep their server "comfortable", experience far fewer hardware problems. As a general rule of thumb, if the computer room is uncomfortably warm for a human, then it is too hot for the server.



## Disk sub-system

The disk sub-system or “hard drive” is one of the most important components within a server and is responsible for file storage. An uncommon - but not unheard of issue - is a lack of free disk space (i.e. capacity). This can bring a file server to a sudden and complete halt. Therefore, Capacity Planning is an important network management task that should be reviewed periodically.

Remember that disk drives are mechanical devices, in contrast to some of the other components like memory, motherboard and processor. This is why, over time, disk drives typically have a higher failure rate than their “non-mechanical” siblings.

Disk drives are also more complex to repair. For example, if the memory or motherboard were to fail (even though there would be some server down-time), once the faulty component has been identified and replaced, the server will be ready to resume operation as before. In contrast, once a failed disk drive has been replaced, all the information that was previously stored on the original disk drive must be restored.

For this reason a single-disk server, is referred to as not being “fault-tolerant” or “redundant”. A failure of the disk drive in a single-disk server will result in a catastrophic “crash”. Time to recover from such a failure will depend on the amount of data to be restored and the type of backups held. The best-case scenario is ½ day and if an “offline backup” is not held, 2-3 days is not uncommon.

Given the amount of disruption that can be caused by the failure of a single-disk server, a fault-tolerant disk subsystem is highly recommended.

A number of options are available depending upon capacity, budget and performance requirements. Regardless of the option chosen, there is a uniform benefit - the failure of a single-disk within the server does not disrupt server operation. When one disk fails the server continues to operate normally without any interruption.

## Power supply

It is important that the power supply has the capacity to do the job properly. With workgroup servers’ sometimes there is a tendency for the power supply to be insufficient (called “under-spec”). This means it doesn’t supply “clean” power to all the components in your server.



The ramifications of this short-supply can be quite insidious - intermittent hardware “lock-ups”, which are difficult to troubleshoot because you have to replace different components within the server until you isolate the culprit (or in the worst case, culprits).

Fortunately this problem can be avoided by taking two steps. First check that the power supply is rated for at least 300W, and up to 450W if you’re running more than two disks or other power-hungry devices like an internal DLT tape drive.

Second, install a UPS (Un-interrupted Power Supply). A UPS is essentially a battery, which continues to provide power to the server in the event of loss of power. This has the benefit of allowing the server to be shutdown in the regular manner during an extended power outage, rather than an immediate “dirty” shutdown (like when the plug is knocked out of the wall).

A good quality UPS has an added benefit in that it provides “filtered” power to the server’s power supply, which is good for the power supply, and thus all of the server’s components. In some areas, particularly where there are nearby manufacturers using “heavy” equipment you can experience surges and/or brownouts in the power being supplied. Even some air-conditioners can have this same effect. Over time these irregularities can seriously impact on the reliability of the server’s power supply if a UPS does not protect it.

## Backups

Everyone knows the importance of backups. Yet on a number of occasions when we have visited a new site for the first time, we are surprised to learn that although the staff were changing the tape on a daily basis, nothing was actually being backed up! You can imagine the horror felt by management when they realised they had no backup at all!

A whole book could be devoted to the topic of backups, but as a general guide the following key items ought to be considered:

1. Make sure you understand which directories on the server contain the essential user data – and ensure that those directories are included in the nightly backup.
2. Ensure that your e-mail messages are included in the nightly backup.

3. Configure the backup job to eject the tape after completion of the backup. This has two benefits: first if the tape isn't ejected the following morning – that's a good indication that there was a problem with the previous nights' backup that warrants further investigation. Second, by ejecting the tape you reduce the likelihood of inadvertently overwriting the previous days' backup because somebody forgot to change the tape.
4. Configure the backup job to print a copy of the summary report after each backup. No report means that there is likely to have been a problem that requires investigation. Ensure that the reports are filed, and reviewed at least weekly for any files that are consistently being missed.
5. Tape media should be rotated such that certain backups are held off-site. In a worst-case scenario where there is a fire at the office, the backups won't be much use if they are destroyed along with the server.
6. It is important to periodically verify that the backup is working properly by performing a test restore of selected data files (make sure that the test restore doesn't overwrite the production data – i.e. you restore to a different destination). Then open the file to make sure the information is intact.

## Disaster recovery

Fortunately these situations are rare, but it's the scenario we all dread. In the context of the server, a "disaster" is an event that results in the server being unable to start properly or an event that significantly compromises the server's ability to perform its tasks. Examples of such events are:

1. Catastrophic hardware failure of the disk sub-system or the entire server (e.g. fire).
2. Virus strike that cannot be repaired.
3. Software update that has significant unexpected consequences that cannot be easily undone.

In each of the above situations, the solution is to restore the server from a backup. Although, in this case it is the server operating system and your server's unique configuration (user accounts, passwords, printer configurations, applications etc.) that need to be restored; plus possibly all the user data depending on the exact nature of the problem.



In this scenario the server is “off-line”. That is, the server is not available for use. So the disruption in the office is considerable, and thus the speed of recovery is important.

The type of backup that is recommended to recover from this scenario is different from what we discussed in the previous section. This is a fairly technical area, but in short, one way that backups can be classified is whether they’re “on-line” or “off-line”. An “on-line” backup is recommended for regular backup of user data and other important information stored on the server. While an “off-line” backup is recommended for backup of the server operating system and associated system files (e.g. registry, directory etc.).

## Virus protection

With the prevalence of e-mail “worm” viruses in recent times there is a heightened awareness of the need to implement measures to protect against viruses disrupting operations.

There are a number of anti-virus products on the market and we have our preference. But regardless of the product chosen, when evaluating a site’s virus defence strategy, we look at the following:

1. Does the vendor release virus definition updates on a regular (at least daily) basis? Virus screening is only as good as the definitions that are used to recognise hostile payloads.
2. Can the vendor’s program be configured to update the definitions on an automatic basis? Refer point 1.
3. Can the vendor’s program be configured to update all the workstations from a central point? Unfortunately new viruses are being released all the time, so you need to have a simple process to keep the entire network up-to-date.
4. What approach is being used to detect e-mail viruses? Are they being screened at the workstation? (Good). Or are they being screened at the e-mail server? (Better).
5. Is one product being used for all virus screening, or is a separate product being used for incoming e-mail? A separate program is preferred, as it means that you have two sets of virus definitions from two different vendors scanning for hostile payloads.

6. Have measures been taken to prevent e-mail worms from propagating in the event that a virus does manage to infiltrate the defences? Network access controls can be put in place to stop an e-mail worm that breaches your defences from being able to propagate outside of your network and reach any of your clients or suppliers.

## Internet connection

One mistake we see from time-to-time (whether it be a dial-up connection or ADSL/cable) is the modem for Internet access is connected directly to the server. This represents a genuine security risk as this means that the server is then directly connected to the Internet. This is an issue for two reasons:

1. Exploits can be launched directly against the server and if the attacker is successful they will have control of, or be able to damage, your most valuable computer. The consequences may be devastating.
2. The server is configured to “share” resources (files, printers, services). So if the server is directly connected to the Internet it will be ‘visible’ to any port-scanning trawl by an attacker. Again this makes the server very vulnerable, as you’re relying on your ISP to protect you from outside attack. Generally ISP’s run fairly “open” policies so as not to “break” client’s applications and thus generate un-wanted technical support calls. Your network needs better protection than your ISP provides.

The recommended approach is to attach the modem to a dedicated Internet appliance (e.g. firewall/router) – or in certain cases a PC based firewall – that does not share any resources and has limited function.

Such a host does not represent an attractive target to an attacker. Moreover even if the host were to be compromised (which we have not experienced to date) as the function of the system is restricted – at worst all that could be accomplished is a disruption to Internet access – the internal network (server and workstations) would be unaffected.



## Conclusion

A healthy server is a critical component of a healthy network. However, it's important to understand that a reliable computer network requires more than just a reliable server.

This guide has discussed some of the more common issues we encounter with servers in a networked environment. The information provided is general in nature and is not intended to replace specific diagnosis of your particular computer network problems. A qualified computer network service provider will be able to provide the most accurate assessment of your needs.

An experienced network service provider will be an asset to your organisation by assisting you to manage all elements of your network. A well-structured computer network will free up your management time, reduce interruptions and create a more productive work environment.

Zero Effort Networking (ZEN) is experienced in accurately identifying computer network problems and acting swiftly to fix them. Since 1996 we have helped clients implement stable, reliable computer networks that take the worry out of network management. For further information you are welcome to contact us.

Zero Effort Networking Pty Ltd  
ABN: 38 082 434 446  
Telephone: 02 9676 3541  
Facsimile: 02 8569 2012  
Email: [info@zeroeffortnetworking.com.au](mailto:info@zeroeffortnetworking.com.au)  
Web: [www.zeroeffortnetworking.com.au](http://www.zeroeffortnetworking.com.au)  
Office: Suite 15, 18 Third Avenue, Blacktown NSW 2148

No part of this publication may be reproduced without written permission of the authors. Copying and distribution by any means is strictly forbidden. Additional copies may be obtained for free by contacting Zero Effort Networking.