



UNIVERSITY OF TASMANIA

---

Central IT

## **Departmental Server Deployment Strategies**

### **Discussion Paper**

Version 1.0.1 Draft

June 2002

Alistair Roberts

## **Introduction**

This paper describes some of the factors that should be taken into consideration by Departments, Faculties and Schools in planning the implementation of data storage systems at the University of Tasmania. The content applies to database, file and print servers, and should be the basis of any departmental data storage strategy. It should also be used as a basis for reviewing existing data storage strategies with a view to a program of service improvement.

## **Related Documents**

- Desktop Computing Policy
- Central IT Security Policy (Draft)
- ICT Usage Policy (Draft)
- CIT Corporate Systems Backup Business Case (May 2002)

## **Planning for Data Availability**

Availability management is a well understood and documented process, however the decision as to what level of availability management is appropriate should be driven by the business needs of the Faculty, School or Section.

The business and technical environment and requirements determine the economically appropriate level of availability. As the level of availability required is increased, so too does the attendant cost. However it is possible to raise the availability level across a Faculty for example by combining the separate resources of each school, within that Faculty in order to gain an economy of scale.

To give a practical example, a small server capable of storing 30GB of data may cost \$5000 to \$6000 to purchase. A much larger server with a capacity of 150GB of data may only cost \$15,000. The per-gigabyte cost is significantly lower for the larger installation. In addition, the attendant cost of maintaining and backing up a small number of larger servers is significantly lower than that for a large number of smaller servers. In short, rationalisation of services into larger server environments delivers economies of scale.

## **Strategies for Deployment and Backup of Server Equipment**

File servers are the accepted and recommended repositories for all corporate data. Concentrating and sharing data via servers dedicated to this purpose enhances security, reliability and flexibility, as these servers can be accessed by the appropriate clients from anywhere on the network.

It must be understood however that a server should be appropriately configured and resourced – it is not sensible to simply “build” a server that will contain critical data out of an uprated desktop computer, as the underlying hardware is unlikely to be adequate in terms of performance or reliability. Similarly, it is imperative that the real cost of adequate backup be factored into any server deployments. Central IT has made a proposal recently to install a corporate Backup System to address this requirement from Easter 2003.

### ***Strategy 1: Save and Restore***

(Recover within days)

The most basic level of protection is to backup your data to magnetic tape so that you can easily restore it if necessary. Magnetic tape is a highly reliable and relatively cheap media type, ideally suited to backup purposes. Backups are performed on a regular daily schedule, with tapes being archived offsite for a period of time.

Central IT recommends a 3-stage backup process, on a 2-year rotating cycle:

- Level 2 – Nightly incremental backups (backs up only that data which is different from the previous night). Three month's worth of tapes are retained.

- Level 1 – Weekly incremental backups (backs up only that data which is different from the previous week). Tapes are retained for 3 months.
- Level 0 – Monthly full backups. Tapes are retained for 2 years.

For capacity planning, the above schedule requires 36GB of tape capacity for every 1GB of data over a 2 year period.

However, it takes a long time to backup large amounts of data. Tape systems vary in throughput (ie speed) and the speed of the network and the network interface on the server can have an effect on backup performance. Modern tape systems are designed to span multiple taps and often utilise multi-loading “stacker” systems in order to automatically load a new tape when the old one is full. Modern tape formats such as DLT, SuperDLT and LTO are designed for high performance in terms of speed and capacity in order to store large amounts of data efficiently.

### ***Strategy 2: Effective Journaling and Commitment Control***

Your availability and recovery plan should also contain measures to capture any data changes since the last backup. All major database management systems offer journaling functions that meet this need.

### ***Strategy 3: Uninterruptible Power Supplies (UPS)***

Two dimensions of downtime impact system availability: frequency and duration. Power failures are the most common cause of abrupt system failures. Thus, Uninterruptible Power Supplies go a long way toward reducing outage frequency.

### ***Strategy 4: Disk Redundancy***

RAID5 (Redundant Arrays of Independent Disks) spreads enough information across multiple disks to allow the disk subsystem controller to recalculate any missing information in the event of a disk failure. Servers utilising RAID5 or RAID5+1 technology are recommended for larger systems and provide significant reliability and uptime benefits over non-RAID systems. In a medium sized server the additional cost of RAID functionality is only a small increase in overall cost.

RAID5 does not protect against the failure of other disk-related hardware, such as a controller, an I/O processor, or a bus.

Disk mirroring if implemented sensibly can protect against these more serious failures. Often used in combination with RAID5, this approach requires that data be concurrently written to each unit in a set of identical disks, incurring minimal CPU overhead or increase in system complexity. However, to mirror all data you must buy twice as many disks, and in order to achieve this level of redundancy, each set of disks should be housed in a separate enclosure driven off a separate controller.

In addition, systems with multiple “hot-swap” power supplies provide higher levels of redundancy against possible power supply failure. Quality systems will also have a large number of internal fans and be very efficient at removing internal heat. Of necessity, such systems are often very noisy.

### ***Strategy 5: Multiple Systems***

(Recovery within minutes)

The fifth level of availability management offers a significantly greater protection since it includes software to automate all of the system administration tasks introduced by journaling and to alleviate any CPU overhead penalty.

More importantly, even the best planned backup strategy and most complete mirroring scheme cannot eliminate all types of user processing interruptions. A multiple system approach can achieve this result.

Some service types lend themselves to multiple, relatively inexpensive systems easily. For example, Central IT utilise multiple medium sized systems for such tasks as NDS replicas, Windows Domain Controllers, and PUPS print servers. Larger systems in multiple include proxy servers and DNS controllers.

In order to provide redundancy for user data across multiple systems, generally some form of data replication is necessary. Central IT plans on implementing this with the Centra Administration file servers, where identical systems in Hobart and Launceston will replicate data via the WAN link in order to provide redundancy in the event of system failure or network failure. An added benefit is increased flexibility for users and higher performance due to the ability of the NDS system to connect users to the nearest server.

**Other issues to be considered:**

- Staff levels – point sensitivity, documentation, replacement arrangements whilst on leave
- Interoperability with Central IT standards/practices - “Look before you leap”. Departmental IT Administrators should contact Central IT before implementing new systems/projects in order to gather advice and to ensure that they are not unnecessarily duplicating a central service.
- Conformity with industry standards & common practice makes everyone’s job easier and is a useful mechanism for “future proofing” a system.
- Ease of upgrade/replacement should be considered – all systems are eventually replaced and a wise systems administrator plans for this at implementation.
- Avoid structural bottlenecks. For example, Central IT suggest the use of cnames (DNS aliases) to allow services to be easily relocated. Different services should be addressed through meaningful cnames rather than machine names in order to allow individual services to be independent of the machine hosting that service.