dantz®

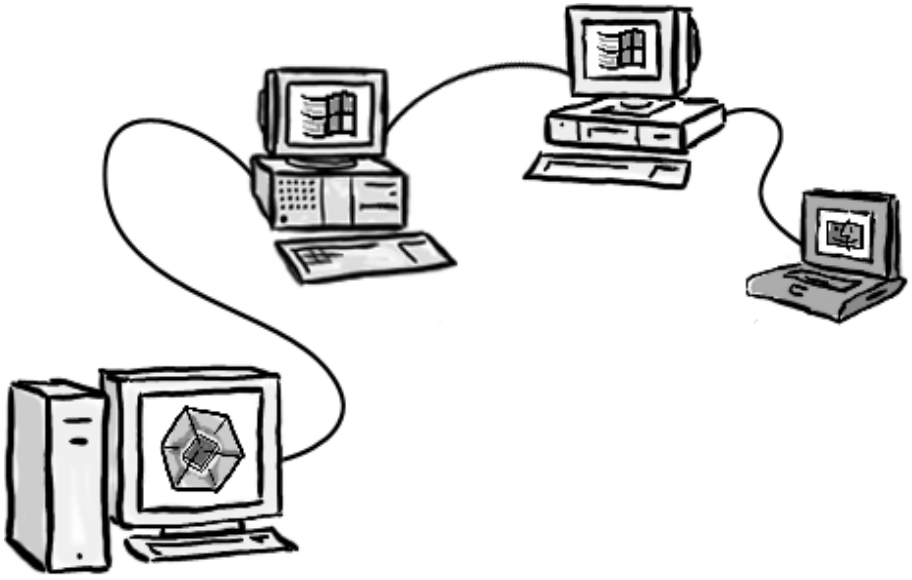

UNDERSTANDING BACKUP

GUIDELINES FOR A SUCCESSFUL BACKUP STRATEGY

WHITE PAPER

Contents

Overview

This document discusses the issues involved in setting up an effective backup strategy to prevent data loss on an individual or network of personal computers.

It first shows the value of your data and how it can be lost. Then it explains the distributed and centralized approaches to backing up networks. This white paper details the general hardware and software features necessary for a successful backup solution, then discusses the administration costs of implementing and maintaining the solution. It goes into more details about hardware features, benefits, and costs. Finally, it offers several software options for your consideration.

Why Back Up?—What's at Stake?

A large bomb explodes in the second level garage at the World Trade Center. The FBI shuts down the World Trade Center for two weeks to investigate the bombing. One year later 34% of the businesses in the World Trade Center without offsite backups failed.

In spite of these amazing statistics, the majority of organizations today fail to adequately protect their corporate data. Organizations around the world are storing vital data on personal computers—mainframes and minicomputers are no longer the sole repositories of the accumulated corporate knowledge. Despite this shift, most organizations still operate without a functional backup strategy to ensure the integrity of their critical, decentralized data.

Sporadic investments in earlier generations of backup software and hardware for personal computers have proven ineffective, mainly because they are cumbersome and go unused. However, the latest generation of personal computer backup and restore software, together with high capacity media and high speed storage devices, provides unobtrusive, reliable "electronic insurance" of a quality previously found only in mini- and mainframe computers.

Users who have implemented an organization-wide backup strategy have identified four major benefits:

1. *Significant cost savings.* As users store more and more information on personal computers, the value of the data soon exceeds the value of the hardware. In fact, the average personal computer used in business costs a fraction of what it would cost to recreate the data stored on the computer.

2. *Higher productivity.* When a computer or mass storage device fails, its user needs to be able to resume work immediately. Without backup, the user spends hours, days, even weeks rebuilding the hard disk and re-creating lost documents. As the user's productivity diminishes, so does the productivity of the user's organization. A dependable backup system protects information and maintains productivity.

3. *Simplicity for end users.* Users questioned about backup invariably respond that they know they should back up but tend not to because it is "too much trouble." Yet the technology exists today to remove that burden by using network-aware backup applications that centralize the backup of entire networks of personal computers.

4. *A more secure computing foundation.* As more and more tasks, such as productivity, presentation, and communication, are performed on desktop computers, a sound strategy of regular backup becomes essential. Without a solid foundation for computing, greater dependence on desktop computers results in greater vulnerability.

WHAT YOU CAN LOSE

Until you actually experience some sort of data loss, you may not realize to what extent loss is possible and the ultimate effects the loss may have.

Consider documents, for example. Because of the gradual nature of their creation, most users are not aware of the value of their total investment in electronic documents until they lose them.

And documents aren't the only form of data you can lose—consider all of the components of your complex computing environment that you can lose as well:

- *Operating System.* Windows® 95, 98, 2000, NT™, and Macintosh® operating systems are complex, with many files in many folders. You probably have an impressive collection of third-party drivers, fonts, control panels, and other system-related resources which would be tedious to reassemble after a disaster. Your ability to restore that precise system environment is a critical component of any backup strategy.
- *Applications.* Re-installation from disks, CDs, and the Internet is tedious at best and disastrous at worst.
- *Preferences.* Applications allow significant customizing during and after installation. You invest in a set of choices, gradually creating the configuration that works best. For example, a spreadsheet application may allow you to customize menus and tool palettes, and even automate data exchange with other applications. Should disaster strike, this customizing might be difficult and time-consuming to re-create.
- *Network.* Users are well aware of the complex and fragile nature of any network environment. This is an area where various hardware and software components interact with other original equipment and third-party products. Users whose network environments have evolved over time find they have dozens of interrelated software modules: libraries, extensions, control panels, data files, and even applications. An incompatible version of one of these files often results in network failure. Therefore, complete restoration of the network software environment has to be part of any backup strategy.

HOW YOU CAN LOSE IT

Data stored on computer is always in danger of being destroyed. All kinds of businesses—from the very large to the very small—lose critical data every day, usually for one of the following reasons:

Unintentional Computing

Most data loss occurs when users inadvertently delete files, folders, or even entire hard disks. Such "unintentional computing" may result from "mistaken identity" of data, disorganized hard disks, or inadequate training. Estimates suggest the vast majority of all lost data is accidentally deleted by users.

Hard Disk Failure

Because hard disk drives rely on precision tolerances, spinning disks, and floating magnetic heads, failure is a given. It is only a question of *when* a drive will fail, not *if*. With almost every personal computer employing a hard disk as its primary storage device, literally years of work may be stored there. Without backup, all of this data can be lost.

Virus Attacks

As a result of all the media attention, most users are now aware of the dangers of computer viruses. These renegade programs can "propagate" themselves from computer to computer, sometimes remaining benign but all too often destroying files or even reformatting hard disks.

Computer Theft

Personal computers are a favorite of thieves of high-tech equipment. Many organizations have lost almost every computer they own in a single night. Along with the hardware goes the data, which is often even more valuable.

Natural Disasters

The relatively unlikely occurrence of a fire, flood, or earthquake is off-set by the catastrophic loss such a disaster would inevitably bring.

Developing a Backup Strategy

Perhaps the most important consideration in developing a backup strategy for desktop computers is deciding whether to employ a *distributed* or a *centralized* strategy. For many organizations the default backup strategy is distributed—allowing the few users particularly concerned with data integrity to perform sporadic backups on their own. Many users and network administrators alike may not be familiar with the entire range of backup alternatives.

THE DISTRIBUTED APPROACH

You may want to keep the responsibility for backup in the hands of the person generating the data: the individual user. Users are often the first to appreciate the importance of the data they individually create, so they have a built-in incentive to back up that data. In addition, you don't have to designate individuals with backup responsibility for others.

A distributed strategy may work best for users such as graphic designers, who want to store data as both backups and repositories for large, infrequently accessed files they can retrieve at will. A distributed strategy may also be the only choice for users isolated in a satellite office or home office.

The key to success in implementing a distributed strategy is setting standards and following a consistent program of user education and training. The biggest single disadvantage of a distributed strategy is that individual users are notorious for avoiding backup. It is very difficult to achieve a mandated level of "data insurance" by relying on a user's initiative. Also, when a large number of users require a backup solution, buying separate storage devices becomes very expensive, and carting around storage devices from place to place is often impractical.

THE CENTRALIZED APPROACH

Local area networks (LANs) can centralize the management of backups. There are two types of centralized strategies.

Server Only

The desire to centralize backup administration and simplify the backup task for workgroups has resulted in many organizations deciding to back up only servers. Users are encouraged or even required to store important files on the server. While this strategy does

centralize the administration of backup on the network, it lacks complete data protection because users disregard regular backups. However, when centralized backup is implemented in phases, server backup is a good first step toward a comprehensive system.

Comprehensive

This strategy encompasses the backup needs of both servers and desktop computers. It is more reliable because responsibility for all administration tasks is removed from the user. It is more complete because all of the user's data is backed up, not just selected data files. This strategy requires little additional cost. Typically, all the data is collected and managed by a backup application running on a central computer with a high-capacity removable media device. While the central computer is often a file server with the backup application running in the background, it can just as easily be any computer on the network. See figure 1.

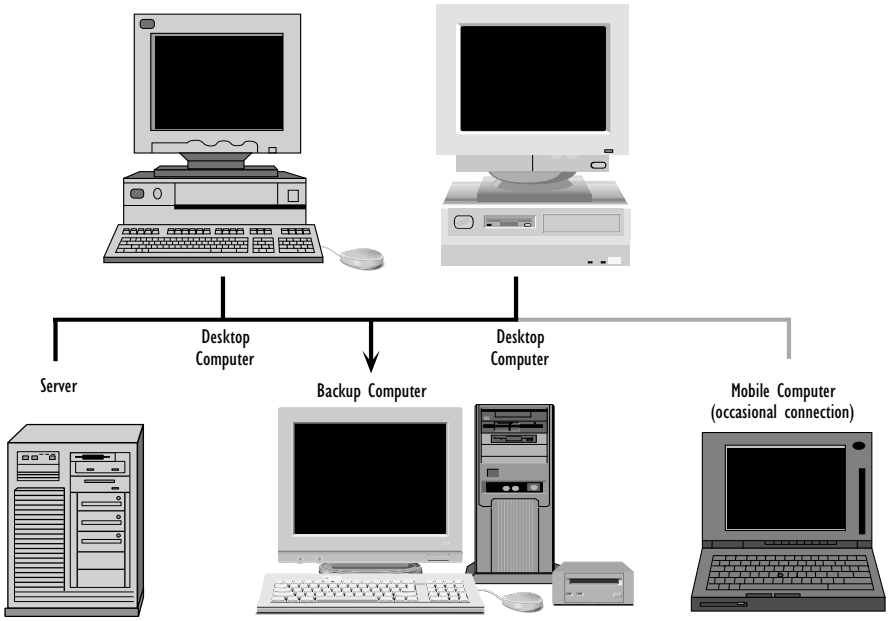

Figure 1: *A comprehensive strategy uses a designated backup computer (which can be almost any computer on the network) to back up not only file servers, but also desktop computers and even portable computers.*

The comprehensive strategy has the following advantages:

- Simplicity and transparency for users
- Lowest cost per user workstation
- Greater coordination and consistency of backup criteria
- Increased automation of the backup system
- Greater security for backup data

The single largest disadvantage to a centralized comprehensive backup strategy is the organization must designate individuals responsible for backup administration. Though the current generation of tools make this task relatively easy, the backup administrator must be willing to accept the responsibility. Fortunately, most organizations already employ people who provide this kind of service.

Evaluating Solutions and Requirements

The basic idea of backing up is simple: make extra copies of important files and put them in a safe place. However, determining the requirements of your backup strategy involves making decisions about such matters as performance, unattended operation, security, storage, and hardware. Much of the functionality of a backup system is determined by its software, so that is a logical place to start evaluating features.

RELIABILITY

Backups must be reliable to be effective. An unreliable backup is rarely better than none at all. Reliability typically concerns three areas: verification, reduction of user errors, and backup strategy.

Verification

There are many types and degrees of verification. First and foremost, the verification procedure should check each area of potential data loss. The highest level of security comes from a "double read back" verify. Once the backup operation is complete, all copied data is re-read from the source and compared byte for byte with the data on the backup media. Although this method incurs a performance penalty, it verifies that data was read correctly from the source, was transmitted via the storage device to the media without error, and may be read again from the media.

Reduction of User Errors

Features designed to prevent user errors radically increase reliability. Backup software should be easy to install, use, and understand. Otherwise, the backups may be unintentionally incomplete.

Backup software should request tapes by name and check each tape's name before writing to protect valuable data from being overwritten. The software should also record the precise state of documents as well as user and network environments to facilitate fast restores should disaster strike.

Backup Strategy

Reliability is greatly influenced by backup procedures. The protection provided by the backup solution not only depends on the frequency with which backups are performed, but also on the extent to which the backed-up files are themselves safe from loss or damage. In

essence, the specific strategy you implement determines the types of catastrophic events against which your users are insured. This is similar in principle to insuring against eventualities such as fire, flood, and theft.

You can increase the degree of backup protection by adopting the following schemes.

Rotating Among Backup Sets Backing up to tape, disk, or other media provides a second copy of vital data, but even tapes and disks can be damaged or destroyed. The solution is to rotate among multiple backup sets. For example, back up to the first set today, the second set tomorrow, and the third set the following day, then back to the first set. This reduces the risk of losing data.The key is to use software that properly handles incremental backups to rotating backup sets. See figure 2.

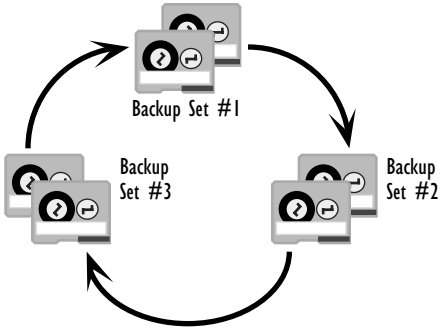

Figure 2: *Rotation among three backup sets.*

Off-Site Media Rotation Multiple backup sets protect against media damage, loss, and failure, but not against natural disasters such as fires, which can destroy all of an organization's computers as well as the very backup on which the organization depends for recovery. Consequently, backup systems provide the highest degree of reliability by regularly rotating copies of the backup off-site. In combination with a multiple backup set scheme, this creates as fail-safe a scheme of data protection as is possible.

PERFORMANCE

Performance is a key factor to evaluate because backup applications handle much more data than other software. Performance influences both capability and level of protection. Capability is affected because the greater the performance, the more you can back up in a given time. The level of protection is affected because the essence of backup is making multiple copies, and higher performance facilitates their creation and timely verification.

More important than the performance of any specific component is how the backup application, source drive, destination device, and "data pipe" work together as an entire system. For instance, the maximum throughput of a tape drive may be meaningless to the performance of the overall system if the software and data transmission components cannot

consistently deliver sufficient data to "feed" the tape drive. In this case, a faster tape drive could even decrease the overall performance of the system when "data starvation" causes the storage device to spend more time repositioning than copying data. See figure 3.

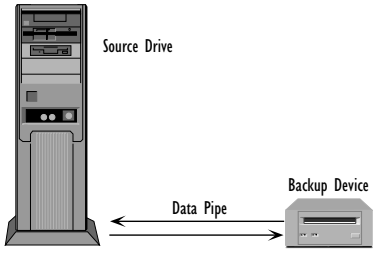

Figure 3: *Overall performance is limited by the elements of the "data pipe."*

Backup software should never sacrifice reliability for performance.

Full and Incremental Backups

The greatest improvements in performance result from backing up only new or unique data to each backup set. The first time a backup set is used, all desired files are copied. This is known as a full backup. Usually, backups subsequent to a full backup are incremental; only files created or modified since the last backup are copied. This saves considerable time and backup media space. A common strategy is to do regular incremental backups to save time and space, with periodic full backups to reduce the size of a complete backup set.

Backup administrators often shy away from incremental backups because most backup software does a poor job of restoring from incremental backups. Select backup software that allows you to execute a complete restore of just the most recent files from a combination of full and incremental backups in a single pass.

Optimized Backups

Cache and other temporary files do not need to be backed up because of their temporary or volatile nature. Because these files have no value, they should not be backed up. Excluding this and other temporary files from backup can reduce backup requirements significantly.

Compression Management

Compression of data, which can be handled by some backup devices and some backup software, can dramatically reduce the amount of storage media required for a given amount of data.

Many backup devices have compression abilities built in to their hardware so they can quickly compress data "on the fly" during backups. The backup application should recognize the compression ability of the backup device and use hardware compression when it can, but when hardware compression is unavailable the application itself should compress the data.

Data compression duties can be handled by some backup software. Although software compression may exact a performance penalty, it may be reduced because there is less (more compact) data to store on the media. A fast backup computer and backup software with a fast compression algorithm can make software compression so fast as to be transparent to the user.

Software data compression can be used to compress backups exchanged among computers and backup devices, as the hardware compression facility on a given backup device may not be compatible with the facility of another. In addition, software compression is used when saving to encrypted backup sets because encrypted data cannot be compressed effectively.

Performance for Network Backup

Performance is perhaps most important when servers or entire networks are being backed up, because of the sheer volume of data and the limited time in which to back it up. With a centralized strategy, increased performance can mean:

- Lower backup system costs because more servers and desktops can be backed up to a single central computer
- An increase in the frequency of backups, providing a greater level of data protection
- Backup of user and network environments in addition to documents

Because a centralized strategy often relies on the network to transport data, managing this part of the "data pipe" efficiently can radically improve performance of the backup.

Fortunately, backup software can use a number of techniques to enlarge the boundaries of network backup and improve performance. The most innovative technique dramatically reduces the amount of data to be transmitted during backup before the backup even begins.

Optimizing Network Transfers For most network operations, performance is determined more by responsiveness than by data throughput. This priority is reflected in network protocols that use relatively small block sizes and wait on acknowledgment. With backup, however, maximizing data throughput is much more critical than responsiveness. Optimum performance requires pipelined data flow with a separate acknowledgment "channel" that may trail the current data block by a significant degree.

Intelligent Data Reduction Identical files are often found on many computers on a network. Recognizing this redundant information and not recopying it unnecessarily saves both network transmission time and storage requirements. For many networks, intelligent data reduction decreases the data to be transferred during a full backup by one-third, increasing the amount of unique data that can be backed up. Intelligent data reduction's actual effect depends on the amount of "identical" files on the network and varies with each organization. Data reduction effectively increases both the amount of data that can be backed up and the effective capacity of a backup drive, providing room for growth. See figure 4.

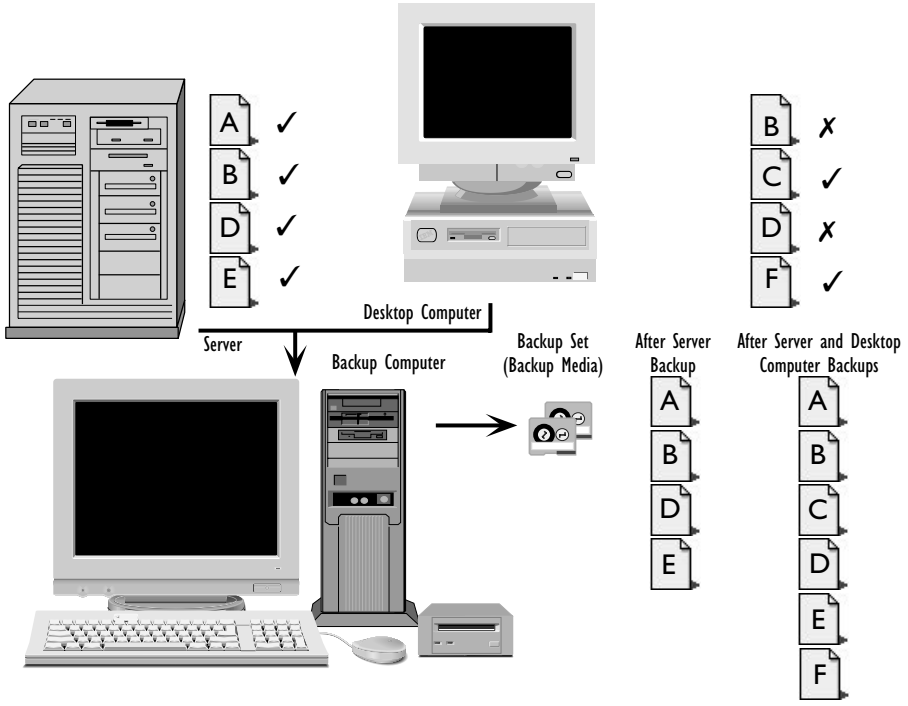

Figure 4: *The first backup to a new backup set copies all files. Subsequently, identical files are not copied.*

Performance for Individual Backup

Even with a distributed backup strategy, performance is often the most important criterion in selecting a solution. With a distributed strategy, increased performance can:

- Save time for users who perform backup tasks, increasing their productivity
- Increase the frequency of backups, providing a greater level of data security
- Allow the backup of user and network environments in addition to documents

COMPLETENESS AND FREQUENCY

Deciding how, how often, and which data to back up is similar in principle to deciding on insurance coverage. The frequency of backups largely determines your exposure to data loss; frequent backups reduce the amount of data lost should disaster strike. However, like lowering the deductible on insurance, increasing the frequency of backups has its costs: escalating the requirements for backup hardware and the amount of time needed for operation.

Most users realize that not all of their data is of equal value, nor is it modified by users at the same frequency. Users tend to want to back up some categories of files more frequently, others less frequently, and some not at all.

For example, it is common to find backup schemes in which documents are backed up most often, applications and system software less frequently. However, powerful features such as incremental backups and "intelligent data reduction" can supersede the need for these types of schemes.

Backing up all files instead of selected categories has significant advantages. When a restore is required, all documents, as well as the system, application, and network environments, can be recovered quickly and easily, usually in one step. Still, some special situations may require selective backup of documents, applications, system software, and so on. Some backup applications have the flexibility to support these criteria.

SECURITY

For many organizations, security is a primary consideration in evaluating backup solutions. Well-designed backup solutions can provide easy-to-use—yet powerful—options for securing data.

Most backup solutions duplicate an individual's data on removable media of one form or another. While this is usually the best way to store backup information, the copy may also unwittingly make an individual's or even an entire organization's most valuable data very easy to steal. Therefore, many organizations that place a premium on security should select backup products with security features that can be tailored to their requirements.

Security for Individuals

Security can be as simple and as inexpensive as storing the backup set under lock and key. There is a way to ensure that backups remain secure, even if they fall into unauthorized hands. Encryption is a common technique that prevents access to backed-up data by users without the encryption key, or password. From an administrator's perspective, encryption can be very easy to use, simply requiring the selection of a feature and the entry of a password when a new backup set is initiated.

When the encryption feature is selected, data is processed before being copied to tape or disk to encode it according to the encryption key. As a result, the data on tape is in a code that can only be decoded with the encryption key. The United States government has been one of the leading developers of encryption technology, and its Data Encryption Standard (DES) technique has become the de facto encryption method in the United States. With the rise of the Internet, other encryption methods have emerged as well.

Security for Networks

Centralized backup presents greater security challenges and opportunities. A centralized backup scheme has challenges in the form of more ways for an unauthorized individual to access data. On the other hand, centralization allows you the opportunities to implement a more consistent, unobtrusive, and comprehensive security scheme.

To provide additional layers of security with a centralized scheme, network administrators should protect backup sets under lock and key and/or encrypt the backup set itself.

Other potential points of entry should be protected in the following ways.

Network Encryption

Centralized network backup solutions depend on the network to deliver data. As a result, there is a risk of unauthorized "wire tapping" by a device that "peeks" at network data traffic. Therefore, network backup should provide safeguards to protect data on the network. The network administrator can prevent unauthorized wire tapping if the backup software ensures that all data sent from servers and users is encrypted before transmission.

Network Password Access

With the power to reach across the network and copy all of the network's most important information, security concerns require that the network backup application have rigorous safeguards. Control of access with passwords is perhaps the most important. The central backup application controls all backup access to other computers in the workgroup, so it should be protected by a password. To prevent users from gaining unauthorized access by using their own copy of the backup application, the software on the networked computers should likewise require a security code for any access. Along with password access to the application itself, this scheme ensures private, secured communication channels between backup nodes and the central application.

Individual and Organizational Security

In organizations with very rigorous security requirements, highly classified information may have to be strictly confined to the users themselves. In these situations a user with highly classified data can still benefit from the advantages of network backup that allows the exclusion of certain directories or allows the use of other desktop security software. This kind of flexibility lets organizations tailor the security of their backup solution to their exact requirements.

UNATTENDED OPERATION

The repetitive nature of backup makes it perfectly suited for automated operation, relieving the user of a burdensome task. To quickly and easily automate your backup strategy, backup applications must include two critical functions that allow unattended operation: scripting and scheduling.

Scripting

A script is a saved procedure involving several elements and a few steps, and its purpose is to automate the procedure of a given task. After you define a script, you needn't make the same repetitive sequence of choices for every backup operation. Scripting must be powerful enough to include all the operations and options you want, yet not so complex that you have to invest inordinate amounts of time setting up the scripts.

Invest in backup software that doesn't require all backup sources or destinations to be "on-line" at the moment the scripts are created because it is often impossible to ensure all sources, including notebooks, will be available at any given time.

Scheduling

Backing up is most effective and unobtrusive if it is done when the computers and the network are relatively inactive. Typically, the backup administrator is not present at those times, and scripts must be executed automatically. Required scripting features include multi-script execution and automatic repetitive scheduling.

Unattended operation results in dramatically lower administration costs and, in some cases, greater reliability. If you want to perform unattended backups, you need a device that can store an entire period of data on a single device. If you cannot perform a complete backup in this way, backup will be delayed as the software waits for new media. This is a particular problem when backups are scheduled overnight or on weekends when administrators are away from the backup computer.

You can never have too much storage capacity, as it directly translates into less frequent administrator intervention to swap media. With adequate capacity, both full and subsequent incremental backups can be automated, usually requiring no attention until the backup device is completely full.

Administration Costs

Whether your strategy is centralized or distributed, administration costs are likely to be a major consideration in evaluating the overall cost of the system. With a distributed strategy, spreading the responsibility for backup among the users hides the true cost of backup administration for the organization. However, in most cases these hidden costs reduce user productivity and decrease the reliability of the backup. Even though a centralized backup strategy requires a central administrator it usually results in lower overall administrative costs.

In either case, the design of the backup system in general and the backup software in particular can greatly lower administration costs. A backup system designed for unattended operation usually minimizes overall administration costs. In this case, the reliability of the system should be the main consideration. Enhanced reliability helps ensure that backups are uninterrupted by errors and that they occur continuously without the administrator's attention. Robust network support, error handling, and automated verification can increase overall reliability, reduce requirements for intervention, and minimize costs.

STORAGE

Determining backup storage requirements helps you select the right backup hardware. More storage capacity allows more unattended operations, making administration easier. Therefore, it is often useful to choose a backup device and media which have storage capacity greater than or equal to all the data to store.

Individual Backup

Even with the best floppy disk backup program, users outgrow floppy-based backup when they have about 40MB of data. With this much data or more to back up, floppy disks become unwieldy and the operation time becomes excessive. With most hard disks storing 2GB or more, most backup systems need a dedicated backup device, which avoids the tedious media swapping of floppy disk backup and allows unattended operation.

To determine minimum backup capacity requirements, simply calculate the total capacity of all the hard disks you want backed up. Selecting a backup device with capacity even greater than the minimum storage requirement brings several benefits:

- Unattended operation over longer periods of time; a full backup and subsequent incremental backups can be done without the need for the user to change media

- The backup device can support the inevitable increase of valuable data
- The extra capacity may allow you to implement a network backup system without any extra hardware requirements

Because of these benefits and low media cost, tape drives are a good choice for individual or network backups. However, if you want a more flexible device which has other uses in addition to backup, you should consider the multi-purpose drives discussed under "Selecting Hardware and Software."

NETWORK BACKUP

Conceptually, determining the storage requirements for network backup is little different from determining the capacity for individual backup. Simply add up the capacity of all the drives to be backed up on the network, including servers and individual computers. If unattended operation is important, select a drive with capacity greater than the aggregate capacity. Network backup almost always demands a high capacity backup storage device. This tends to narrow your hardware options to tape drives, as discussed in the following section.

This strategy yields benefits which are basically the same as those of individual backup. However, there are additional considerations.

Network Bandwidth

You can choose from several networking alternatives to provide an appropriate data pipe for any network backup requirement.

A slow network can be a bottleneck to a network backup system, and a fast network can transfer data as fast as some backup devices. A TCP/IP implementation over Ethernet cabling should be adequate for most installations.

If your backup system includes some implementation of intelligent data reduction (discussed on page 14), it minimizes the amount of data to be transferred over the network.

Workgroup Design

An appropriate strategy for multiple zone networks must take into account the possible performance penalty exacted by network routers, which can degrade network throughput. Administrators should carefully construct backup workgroups to avoid crossing routers unless they are modern, high-performance hardware routers. See figure 5.

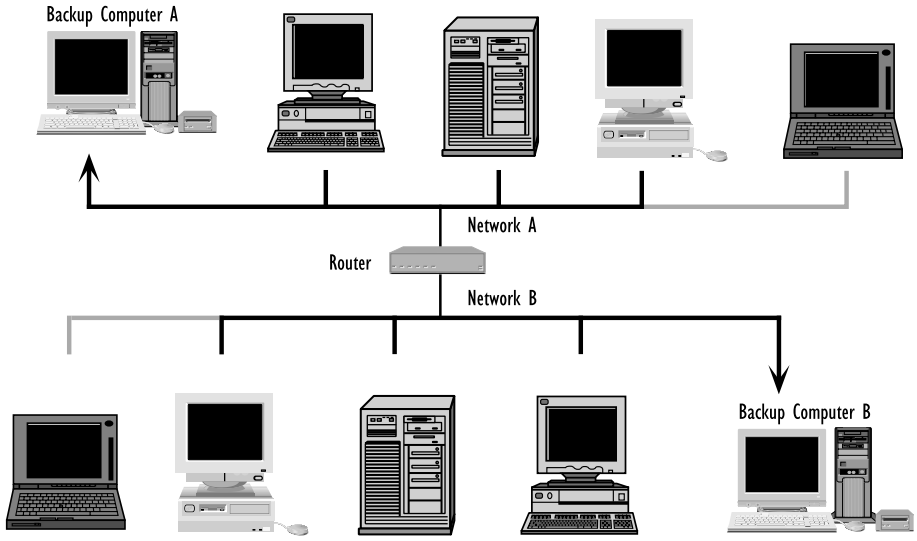

Figure 5: *Multiple backup workgroups connected by a network router.*

Dial-In or Remote Networks

Backing up computers over a modem or dial-in network connection is possible but not advisable because of the low throughput of the typical network phone bridge. The speed of your network bridge, the reliability of the connection, and the amount of data you wish to backup determine whether backup over a temporary connection is feasible.

FUTURE GROWTH

In addition to all the strategic issues discussed so far, also consider the potential growth of your personal computer or network. Larger hard disks will require more backup storage capacity. Backing up all the additional data in a limited amount of time may require faster hardware, including the backup computer, backup device, and the network over which the data travels. You never can have too much storage space nor too much speed for backups. It may even be necessary to implement more comprehensive backup strategies to handle future growth.

Backup Hardware Options

Hardware solutions fall into three categories: drives for network or server backup, drives for individual backup, and multi-purpose drives for a variety of uses. This section discusses the relative merits of all the options, and compares their advantages and disadvantages.

CAPACITY

In evaluating hardware—especially tape drives and particularly compression drives—the rated capacity should usually be viewed as the theoretical maximum. The capacities achieved in day to day use are often significantly less.

The actual amount of data that will fit on a given medium will vary due to many factors. For instance, a tape's capacity can be greatly influenced by the relative speeds of the backup computer and the tape drive. Another factor is compression, a feature of some drives and software which reduces the size the backup data occupies on the media. Theoretical maximum capacity representations refer to the amount of data before it gets compressed by a tape drive with hardware compression capability—and they often assume generous compression rates.

NETWORK OR SERVER BACKUP HARDWARE

Backing up servers or entire networks of computers requires hardware that can store extremely large quantities of data over a relatively short period of time—overnight or a weekend. Fortunately, many different reliable, high-capacity, high-performance backup devices are available.

DAT Drives

Four millimeter Digital Audio Tape drives have become very popular backup devices because of their low cost, quick speed, and low media cost relative to their high capacity. The drives use a helical scan method to store data on 4mm digital audio tape cartridges, though special data grade tapes must be used. Tapes written to by a given DAT drive can usually be read by DAT drives from another manufacturer, provided they use the same format and compression features, which is a useful compatibility consideration.

Many vendors offer several types of DAT drives varying in capacity from 1.3 to 20GB, while models with compression offer up to 40GB of storage. Many autoloaders and tape libraries

are available with one or more DAT drives. In recent years DAT drives have become an industry standard for backup devices by combining reliability, high capacity, and low cost.

8mm Tape Drives

Like DAT drives, 8mm tape drives are helical scan devices that can store large amounts of data on a single tape. Rather than digital audio tapes, 8mm drives use data versions of the tape cartridges used in 8mm video equipment. Entry level 8mm drives offer 2 to 7GB of uncompressed storage capacity. Higher-end 8mm drives are among the fastest and highest-capacity drives available, offering 20 to 50GB of uncompressed storage.

Travan Tape Drives

Unlike the DAT and 8mm tape mechanisms borrowed from other industries, the Travan mechanisms are a more traditional backup tape technology. Travan devices store data linearly on each tape. As a result, data is stored less densely than on DAT or 8mm drives. Travan devices range in capacity from 4 to 13GB, with compression drives potentially doubling capacity up to 26GB.

DLT™

DLT drives are among the fastest tape drives available. These mechanisms offer exceptional performance and capacity when used under optimal conditions. DLT drive models are available in different capacities up to 40GB, with compression drives potentially doubling capacity up to 80GB.

Autoloaders and Tape Libraries

Several manufacturers now offer autoloaders, which use one or more drives (typically tape drives discussed above) integrated with a tape changing mechanism. The device moves media between the drive and a magazine containing several media, multiplying the available storage capacity several times over.

Like multiple-disc audio CD changers which automatically play one CD after another extending your listening time, autoloaders extend the duration of unattended backup operation. When used with appropriate backup software, autoloaders allow multiple storage set backup strategies to be implemented in the absence of a person to change media.

Table 1: Network/Server Backup Hardware

Type of Drive	Advantages	Disadvantages
Digital Audio Tape (DAT)	• Affordable • High capacity • Low media cost	• More expensive than some lower capacity drives
8mm Tape	• High capacity • Low media cost	• More expensive than DAT and most other tape drives • Slow file restores
High-end 8mm Tape	• High speed can back up more data in a given time period • Very high tape capacity	• High price • High individual media price
Travan	• Affordable	• High individual media price • Typically lower capcity media
Digital Linear Tape (DLT)	• High speed can back up more data in a given time period • Very high tape capacity • Good media longevity, relative to other tapes	• High price • High individual media price
Loaders and Libraries	• Allow unattended backup even if storage demands are greater than a single medium • Ideal for use with backup servers	• High price

INDIVIDUAL BACKUP HARDWARE

Whereas network and server backup often require extremely high-capacity, high-performance backup hardware, individual backup places more importance on the cost of the hardware solution, hence the popularity of floppy disks.

Floppy Disk Drives

Because floppy disk drives come built-in with most computers sold, it is not surprising that floppy disks are popular for personal computer backups. However, their performance and capacity make them impractical for backing up the typical personal computer's 2 to 8GB hard disk. Even with highly optimized backup software backing up only your most critical 40MB of files, you have to spend three to four hours inserting and removing floppy disks. If you have more data than spare time, you should choose one of the other inexpensive storage options better suited for backup and archiving.

Travan Tape Drives

Travan drives are relatively inexpensive mechanisms with moderate capacity. They are a good choice for individuals with moderate backup requirements.

DAT Drives

Rapid decreases in the price of digital audio tape drives have made them popular for individual backup, as well as for network and server backup. They are a particularly good choice for users with large hard disks.

Removable Cartridge Drives

Many individuals do their backups with removable cartridge drives (such as Zip, Jaz, SuperDisk, and magneto-optical) which are not dedicated solely to backup. These drives are described next under "Multi-purpose Drives."

Table 2: Individual Backup Hardware

Type of Drive	Advantages	Disadvantages
Floppy Disk	• No additional hardware investment required • Low individual media cost	• Lowest performance and capacity of any backup device • Least reliable media
Travan tape	• Affordable medium-capacity drive • Moderate performance	• High individual media price
Digital Audio Tape (DAT)	• Affordable high-capacity drive • Low media cost	• More expensive than some lower capacity drives

MULTI-PURPOSE HARDWARE

The following drives are primarily used for archiving, data exchange, and long-term storage and are used secondarily for backup. They usually are more expensive or require much more expensive media than dedicated backup devices such as tape drives. However, if you need one of these drives for other purposes, they can do double duty as part of a backup system.

Removable Cartridge Drives

Zip, Jaz, and SuperDisk drives use removable disks or cartridges that approach hard disks in performance. Because they are random access devices, you can use them as either desktop-mountable volumes or as storage devices for backup and archiving.

Hard Disks and Servers

Using hard disks or hard disks installed as servers for backup is almost always the most expensive backup solution, albeit a potentially high-performance one. Keep in mind that

most users whose backup demands outgrow floppy capabilities need a backup device with capacity equal to or greater than their hard disk size. Therefore, a hard disk used for backup must be at least as large as the one used for normal use.

■ **NOTE:** Using hard disks as secondary storage devices usually makes sense only when special requirements to guard against hardware failures call for techniques like disk mirroring. With disk mirroring, instead of storing information on a single hard disk, you use a second hard disk to duplicate data. If one hard disk fails, computing can continue on the other one. Disk mirroring is not a replacement for backup because it only provides immediate protection against disk failures, not against accidental deletion, data corruption, virus, theft, or natural disaster.

Magneto-Optical Drives

Unlike disk drives which use a magnetic coating to store data, magneto-optical (MO) drives use light to store data on a removable optical disk. For most of these mechanisms, writing information onto the disk requires three passes over the disk, whereas reads require only one. This results in asymmetrical performance for applications like backup; backup is slow but retrieval of information (particularly of specific files) can be fast.

Their high capacity (relative to other removable cartridge drives) and flexible nature make MO drives a good choice for users with multiple tasks.

CD-R and CD-RW Drives

Recordable and rewritable compact disc drives use a laser to store and retrieve data on a removable, optical disc. Data stored on CD-R discs cannot be erased, so these drives tend toward archiving use rather than backup. CD-RW discs are rewritable and can be recorded over and over like floppy disks or removable cartridges; there is a limit to the number of rewrites, but you are not likely to encounter it with backups. CD-R and CD-RW discs, like CD-ROM discs and audio compact discs, are durable and long-lived, making them ideally suited to archiving.

CD-R and CD-RW drives require special software, as they are not directly supported by popular operating systems. They offer about 600MB of useful capacity, but their speed is slow compared to most other backup devices. Unlike MO and other removable cartridge drives, CD-R and CD-RW drives cannot be used like hard disks or floppy disks.

DVD-RAM Drives

Phase change drives are unique and flexible in that they work with both proprietary optical media and CD-ROM discs. The quad-speed CD-ROM capability, being "read only," is not useful for backup or archiving, but PD drives are capable of storing 650MB on each optical cartridge. PD optical cartridges work like removable hard disk cartridges, MO cartridges, and floppy disks.

Table 3: Multi-purpose Drives

Type of Drive	Advantages	Disadvantages
Removable Cartridge, including Zip, Jaz, and SuperDisk	• Flexible; can be used for transport, supplemental storage, backup, archiving • Fast file retrieval	• High media cost • Lower capacity media does not allow unattended backup of larger hard disks
Hard Disk, Including File Server	• High performance backup operation • Fast file retrieval	• Most expensive drive for backup • Capacity requirements usually preclude use as archive device • Does not protect against theft, disaster, or user error
Magneto-Optical (MO)	• High media longevity • Fast file retrieval	• Slower backup performance • Expensive media
Recordable and Rewritable Compact Disc (CD-R and CD-RW)	• High media longevity • Low media cost • Fast file retrieval	• Media space cannot be re-used to update obsolete files • Slow performance
DVD-RAM	• High media longevity and capacity • Reads and writes PD disks • Also functions as a CD-ROM drive	• High media cost • Slower performance

MEDIA COST

The following table lists several popular media from popular backup devices, the capacity and cost for each medium, and the cost per gigabyte and cost per megabyte of storage.

In the media cost table, given capacities are native. Compression can greatly affect the amount of original data squeezed onto a given medium and, therefore, the "cost per."

Media costs are the typical market costs at time of publication, though costs vary among vendors. Media costs also vary depending on quantity, and this table uses prices for media purchased in modest quantities. Figures are rounded.

Table 4: Media Costs of Popular Backup Devices

Type/Size	Capacity	Cost	Cost per GB	Cost Per MB
Hard disk	8GB	$400	$50	5¢
Floppy disk (high density)	1.4MB	56¢	$400	40¢
Zip	100MB	$13	$133	13¢
Jaz	1GB	$95	$95	10¢
SuperDisk	120MB	$13	$108	11¢
MO 5.25in	1.3GB	$43	$33	3¢
MO 5.25in	650MB	$30	$47	5¢
MO 3.5in	230MB	$11	$49	5¢
MO 3.5in	128MB	$10	$75	8¢
CD-R	600MB	$2	$3	less than 1¢
CD-RW	600MB	$9	$15	2¢
DAT 60m	1.2GB	$6	$5	less than 1¢
DAT 90m	1.9GB	$7	$3	less than 1¢
DAT 120m	4GB	$11	$3	less than 1¢
DAT 125m	12GB	$27	$2	less than 1¢
8mm tape 54m	1.2GB	$6	$5	less than 1¢
8mm tape 112m	2.5GB	$6	$2	less than 1¢
8mm tape 160m	5.0GB	$11	$2	less than 1¢
8mm tape AME 170m	20GB	$86	$4	less than 1¢
8mm tape SDX-T3N	25GB	$86	$3	less than 1¢
TR4	8GB	$32	$4	less than 1¢
Travan NS-20	20GB	$45	$2	less than 1¢
DVD-RAM 5.2GB	5.2GB	$47	$9	1¢
DVD-RAM 2.6GB	2.6GB	$29	$11	1¢
DLT CT-III	10GB	$35	$4	less than 1¢
DLT CT-IIIXT	15GB	$41	$3	less than 1¢
DLT CT-IV	20GB	$77	$4	less than 1¢

BACKUP SOFTWARE FROM DANTZ DEVELOPMENT

With headquarters in Orinda, California, European offices in Paris, and distribution throughout the world, Dantz has rapidly become one of the largest suppliers of backup software for standalone and networked computers. Dantz believes safeguarding and managing the vast quantity of data stored on personal computers is critical to maintaining a productive working environment.

Dantz publishes a wide range of backup solutions for individuals, workgroups, and networks. These products were designed to make data protection automatic, fast, and efficient. Dantz products are available through resellers worldwide.

RETROSPECT® EXPRESS BACKUP

Designed for individual users, Retrospect Express includes compression, fully automated backups, archiving, and one-step restores. It is optimized for use with removable media, including Zip, SuperDisk, DVD-RAM, and CD-R/CD-RW drives.

RETROSPECT DESKTOP BACKUP

Easy to setup and simple to use, Retrospect Desktop Backup gets the most out of every kind of storage device to provide reliable protection against data loss. Retrospect Desktop Backup includes all the features of Retrospect Express plus it supports tape drives and small tape libraries, includes software encryption for further data security, and is expandable to provide comprehensive and reliable network backups. Retrospect Desktop Backup is available for Windows or Macintosh computers.

RETROSPECT WORKGROUP BACKUP

Retrospect Workgroup Backup runs on any Windows-based computer or server (including Windows NT Server) plus 20 clients to effortlessly protect files in an entire workgroup. Its easy-to-use networking provides complete, low maintenance backup for any workgroup. Retrospect Workgroup Backup is available for Windows or Macintosh computers.

RETROSPECT SERVER BACKUP

Retrospect Server Backup provides the power to restore either a single file or an entire hard disk to any computer on the network. It includes 100 clients plus advanced networking features to provide comprehensive backup of all your networked data.

RETROSPECT CLIENT BACKUP

Add more networked Windows or Macintosh computers and servers to your Retrospect backups with Retrospect Clients, available in 5, 10, 50, and 100 packs.

Dantz Development Corporation
4 Orinda Way, Building C
Orinda, CA 94563 USA
Tel: 925.253.3000
Fax: 925.253.9099
info@dantz.com
www.dantz.com

Dantz Europe
50, rue des Archives
75004 Paris - France
Tel: (33) 1 40 29 11 00
Fax: (33) 1 40 29 11 09
europe@dantz.com