The Weakest Link in Disaster Recovery

By Alex Bakman, CEO

Much of the focus of disaster recovery planning is on creating redundant data sites and backup tapes. Very often, a crucial component is overlooked: that of keeping current documentation for all IT configuration settings. Having such documentation and the original software discs can restore a network 40 percent faster than running backup tapes.

Access to the latest detailed configuration settings means faster disaster recovery. This paper demonstrates where having detailed configuration documentation fits in the disaster recovery process and how it aids in the rapid restoration of an IT infrastructure.

The characteristics of an "ideal" tool to solve the problem of collecting and documenting current configuration settings and how such a tool and the information it provides can solve some of the day-to-day challenges of managing an enterprise IT infrastructure, will also be outlined.

To be effective, most disaster recovery (sometimes called Business Continuity) plans require extensive testing, skilled personnel, access to vital records, and alternate recovery resources, including backup facilities. Plans are designed to restore an IT infrastructure—the business backbone of today's corporation—as quickly as possible.

For most organizations, information, and the technology that supports it, represents the organization's most valuable assets. Enterprise applications are deployed over multiple systems and databases—and across multiple locations—making documenting current configuration settings a very difficult task.

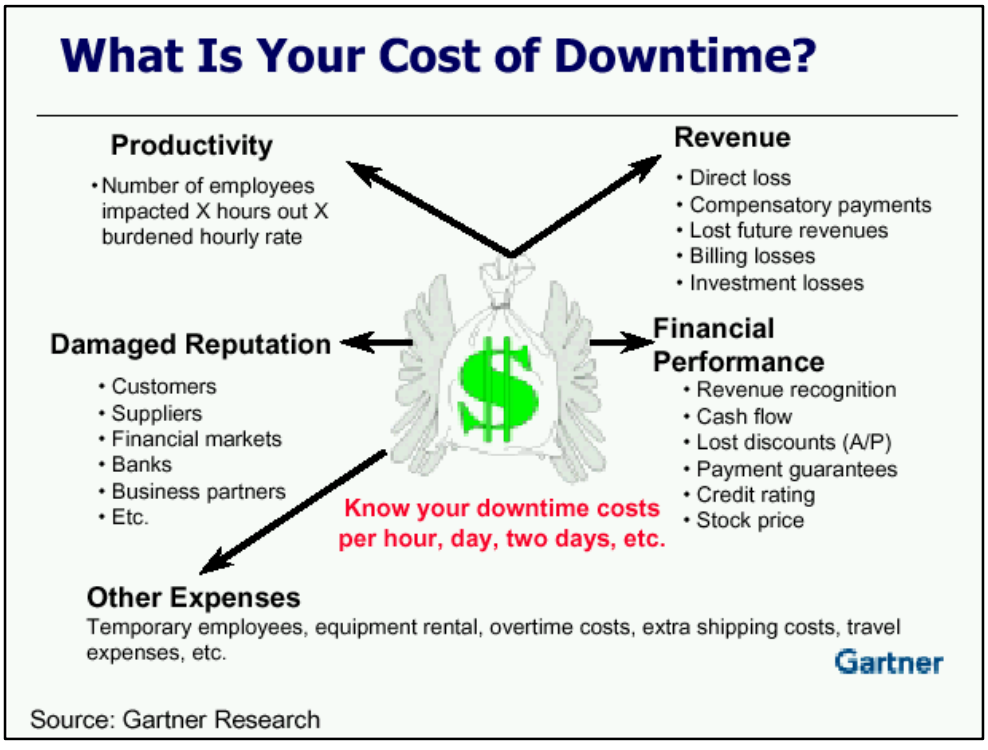

Time is Money

The high cost of business downtime goes beyond lost sales. Failure to perform can lead to contractual penalties. Customers who choose an alternate supplier may never come back—and even if they do, your cost of sales can increase due to a new competitive mix. If records such as invoices are lost, you lose revenue on delivered products and services.

While you are waiting to restore your IT infrastructure, you still have to pay salaries or suffer a possible public relations disaster. In the case of the September 11th tragedy, a company's reputation may not suffer, but stock prices, credit ratings, and cash flow may still be impacted.

Faced with the events of 9/11, enterprise IT departments are focusing more time and money on disaster recovery plans, equipment, and services.

Many enterprises have invested heavily to ensure the survival of the IT infrastructure (e.g. physical data center security and fire prevention equipment) as well as in redundant data centers (e.g. distributed corporate data processing facilities and/or outside services) to ensure business continuity. Most mission-critical applications have their data backed up to tape or other media and these archives (and other critical documents) are best stored in a safe site off the corporate premises.

Many disaster recovery plans include some level of IT infrastructure configuration data, collected as a "snapshot" at a given point in time. Typically, this a hardware and software asset catalogue which contains such information as vendor name, model number, serial number, location, etc. for hardware and vendor name, version number, service pack information, etc. for software.

Most enterprises feel they have all the bases covered with these products and services. However, the speed of business restoration is impeded by inadequate IT infrastructure documentation. Detailed knowledge of server, database, and router configurations is essential to re-establishing a working framework in which to restore corporate data and services.

The IT disaster recovery plan has, until recently, been viewed as a static document that sits in a three-ring binder on every IT mid-level manager's shelf, doing little more than provide comfort that the IT department is ready to do its part to ensure business continuity. Collecting this information from diverse platforms and "massaging" it into meaningful information (if that is attempted) takes a tremendous number of hours and most IT departments do not devote resources to keep the information current.

As a result, all configuration data collected in these documents rapidly becomes out of date due to the one constant in the IT world: change.

Until recently, most disaster recovery plans assumed the existing IT staff would be involved in the restoration process.

For a fire in a corporate data center, this might be true. However, in a natural disaster such as a tornado or flood where the area surrounding a data center may also be affected, IT staff will initially be more concerned with their families and homes than with their work responsibilities. Sadly, September 11th taught us that the unthinkable could occur. Even if IT staff is available to assist recovery, the multitude of IT platforms and the large

number of changes that occur on a daily basis limit their effectiveness to support a backup data center's restoration efforts.

Thus, the IT disaster recovery plan needs to be continuously updated with the latest configuration settings reported in a clear, consistent manner. All changes should be easily identifiable to preserve IT decisions.

1. There are three main reasons that detailed configuration data is not collected and kept current in many enterprises: Almost no company has enough IT staff. According to the Information Technology Association of America (ITAA), of the current US IT workforce requirement of 10 million, there are over 800,000 vacancies that cannot be filled due to the lack of trained talent. The workload increases, but hiring never keeps up.
2. The technical competence of individual IT talent varies with training and experience. Configuration documentation may seem an "entry-level" task that most professionals seek to quickly move beyond. Disparate IT staff members often collect different types of information and the quality of their reports varies greatly. The more senior IT people are assigned to "more critical" tasks, deployed by management where they provide the most perceived value for their salaries, which average $85,000 per year ($75 per hour). The hours needed to assemble, verify, and report configuration settings could amount to tens of thousands of dollars in a larger IT shop.
3. IT staff turnover ranges from 8-17 percent, depending on industry and geographic region. The costs of hiring and training new staff to replace lost employees is nearly triple the IT overhead cost (about $225 per hour). And when IT staff leaves, their knowledge of the corporate IT infrastructure leaves with them.

Are Backups Enough?

One of the most common reasons that detailed configuration information is not recorded is the belief that backups contain everything needed to restore systems into production.

The effectiveness of backup tapes depends upon the nature of the disaster. A system that experiences a simple power outage or hardware failure can easily be restored from backups, but restoring following a complete meltdown is another matter.

Critical information not contained in backups includes: hardware specification for each system, EEPROM settings, specific boot instructions, SCSI ID manipulation, BIOS versions, virtual memory swap space sizes, disk partition slices, space allocation considerations, recovery/re-installation prerequisite considerations, network services provided, network dependencies required for normal functioning, kernel parameters, initial system installation cluster, and configurations that affect storage devices. Typically, volume management software and RAID software is on the tape, but is useful in disk arrangement prior to reinstallation and restoration.

While preserving business data and transactions, most backup tapes contain no configuration data. If you have a hot backup site and don't even need to use tapes, you still need to build a system on which to restore the data.

Before you can restore the data, you need to reconfigure your IT infrastructure to support it.

Following are the five stages of a typical disaster recovery.

1. Immediate Response: This ensures the safety and evacuation of all employees, notification of appropriate management and continuity personnel, assessment, command center activation, and disaster declaration processes.

To demonstrate how immediate a response can be: On September 11th, the first disaster declaration to reach a backup service came at 9:02 am, 17 minutes after the first plane crashed into the World Trade Center.

2. Environment Restoration: This phase provides alternate space for the people and equipment. Generally people relocate to recovery work areas while computer equipment is sent to a hot site or data recovery center. Operating system software is restored for computer systems, while basic facility preparations are conducted for work areas.

3. Voice and Data Communications Restoration: This phase restores communication to the functional parts of the organization, customers, vendors, and, in some cases, the disabled facility.

4. Functional Restoration: This phase includes restoring systems for computer applications, which can only be done once the infrastructure is properly reconfigured. A critical element is the most recent security settings. You need to ensure that the restored applications do not have any security holes when they are returned to production.

5. Restoration and Synchronization: This final phase includes restoration of data from offsite locations through the use of electronic media and paper. Information that was not backed up may be lost forever. Most often data is protected at different times during the business cycle and synchronizing, validating, and reviewing data from different sources is a critical step in a successful recovery. Once reliable data is established, backlogged transactions that have accumulated during the recovery are processed.

Throughout the recovery process, detailed configuration documentation that contains change information allows the original IT staff and the restoration team to easily see, discuss, and alter any changes in configuration settings that occurred from the last safe settings.

Ideal solution for collecting and maintaining configuration documentation

The answer to the problem of collecting and maintaining detailed configuration documentation is quite simple: automation.

Where the process is being done manually (or with tools that provide parts of the required information), automated tools make this task easier by eliminating the work involved and increasing the speed that the information is collected. They also improve the accuracy of the information collected by superceding the human “error prone” involvement in sorting through and reporting mountains of input data. The information is presented and preserved in a consistent format.

Where the process is not being done, automated configuration reporting and change management tools make it possible to actually accomplish the task for the first time.

Automatic scheduling of configuration data collection eliminates the manual part of the process. Automation allows the data to be updated on a regular basis to ensure that the most current information is always available for disaster recovery.

Cost Savings Offset Purchase Price

Unlike an insurance policy where only a "disaster" realizes the benefits, detailed configuration information and documentation can be used on a daily basis to improve the operations of the IT infrastructure; in troubleshooting, security, auditing, and training.

Compliance Reporting is a subset of a larger IT management requirement that is driven by individual industry requirements for security—both of the data being managed and of the IT Infrastructure itself.

A critical component for being in compliance with these industry-specific mandates is possessing current and historical documentation that provides detailed configuration settings of the IT Infrastructure.

For example, the healthcare industry is working towards compliance with the Health Insurance Portability and Accountability Act of 1996—known as HIPAA. Within HIPAA is the requirement for the security and confidentiality protection of electronic health information. Ecora's products contribute toward compliance with the security requirements of HIPAA by providing current (and historical) detailed configuration reports to support auditing, security, and disaster recovery.

There are similar requirements in the financial industry, including Gramm-Leach-Bliley, mandated by the Federal Reserve System, which requires recording detailed configuration settings for security and backup. Firms that are ISO-9000-compliant or working towards that certification also require extensive documentation of IT processes and policies.

In summary, detailed configuration settings for all the major devices on your network should be kept accessible and up-to-date in order to speed recovery in the event of a disaster. Most companies do not maintain such documentation due to a lack of resources. Automated solutions exist that take the guesswork out of documentation and enable staff to schedule a variety of summary and change reports. Reports on settings are not only useful for restoring a network 40 percent faster in the event of a disaster, but also furnish crucial compliance information.

Alex Bakman (abakman@ecora.com) is founder & CEO of Ecora Software, makers of Configuration Auditor, a comprehensive solution for auditing, security, and disaster recovery.