An Enterprise Executive Top Five Article

The New Paradigm of Disaster Recovery

Business Continuity (BC) is an integrated, enterprisewide process that includes all activities—both internal and external to IT—that a business must perform to mitigate the impact of planned and unplanned downtime. This entails preparing for, responding to and recovering from a system outage that adversely impacts business operations. The goal of BC is to ensure the availability of information required to conduct essential business operations.

Information Availability (IA) refers to the ability of an IT infrastructure to function according to business expectations during its specified period of operation. When discussing IA, we need to make certain:

- Information is accessible at the right place to the right user (accessibility)
- Information is reliable and correct in all aspects (reliability)
- The information defines the exact moment during which information must be accessible (timeliness).

Various planned and unplanned incidents result in information unavailability. Planned outages may include installations, maintenance of hardware, software upgrades/patches, restores and facility upgrade operations. Unplanned outages include human error-induced failures, database corruption and failure of components. Other incidents that may cause information unavailability are natural and/or man-made disasters such as floods, hurricanes, fires, earthquakes and terrorist incidents. The majority of outages are planned; historically, statistics show the cause of information unavailability due to unforeseen disasters

# of Business Continuity & Continuous Availability

By Steve Guendert, Ph.D.

is less than 1 percent.

Information unavailability (downtime) results in loss of productivity and revenue, poor financial performance and damage to a business's reputation. The Business Impact (BI) of downtime is the sum of all losses sustained as a result of a given disruption. One common metric used to measure BI is the average cost of downtime per hour. This is often used as a key estimate in determining the appropriate BC solution for an enterprise. Figure 1 shows the average cost of downtime per hour for several key industries.

How Do We Measure IA?

IA relies on the availability of both physical and virtual components of a data center; failure of these components may disrupt IA. A failure is defined as the termination of a component's capability to perform a required function. The component's capability may be restored by performing some sort of manual, corrective action; for example, a reboot, repair or replacement of the failed component(s). By repair, we mean that a component is restored to a condition that enables it to perform its required function(s). Part of the BC planning process should include a proactive risk analysis that considers the component failure rate and average repair time:

- Mean Time Between Failure (MTBF) is the average time available for a system or component to perform its normal operations between failures. It's a measure of how reliable a hardware product, system or component is. For most components, the measure is typically in thousands or even tens of thousands of hours between failures.
- Mean Time To Repair (MTTR) is a basic measure of the maintainability of repairable items. It's the average time required to repair

a failed component. Calculations of MTTR assume that the fault responsible for the failure is correctly identified, and the required spare parts and personnel are available.

We can formally define IA as the period during which a system is in a condition to perform its intended function upon demand. IA can be expressed in terms of system uptime and system downtime, and measured as the amount or percentage of system uptime:

IA=system uptime / (system uptime + system downtime)

where system uptime is the period of time during which the system is in an accessible state. When it isn't accessible, it's termed system downtime. In terms of MTBF and MTTR, IA can be expressed as:

IA=MTBF / (MTBF + MTTR)

Uptime per year is based on requirements of the service for exact timeliness. This calculation leads to the well-known number of "9s" representation for availability metrics. Figure 2 lists the approximate amount of downtime allowed for a service to achieve the specified levels of 9s of availability. Figure 3 lists some typical Recovery Point Objective (RPO) and Recovery Time Objective (RTO) values for some common DR options.

Cost of Downtime by Industry

Industry Sector	Loss per Hour
Financial	$8,213,470
Telecommunications	$4,611,604
Information Technology	$3,316,058
Insurance	$2,582,382
Pharmaceuticals	$2,058,710
Energy	$1,468,798
Transportation	$1,463,128
Banking	$1,145,129
Chemicals	$1,071,404
Consumer Products	$989,795

Source: Robert Frances Group "Picking up the Value of PKI: Leveraging z/OS for Improving Manageability, Reliability and Total Cost of Ownership of PKI and Digital Certificates"

Figure 1: Average Cost Per Hour of Downtime by Industry

An Evolving Paradigm

The importance of Business Continuity and Disaster Recovery (BC/DR) for IT professionals and the corporations they work for has undergone considerable change in the past 20 years. This change has further increased exponentially in the years since Sept. 11, 2001, Hurricane Katrina in 2005, the Japanese Tsunami of 2011 and Superstorm Sandy. The events of 9/11 in the U.S. served as a wakeup call to those who viewed BC/DR as a mere afterthought or necessary "check in the block." The day's events underscored how critical it is for businesses to be ready for disasters on a larger scale than that previously thought. The watershed events of 9/11 and the resulting experiences served to diametrically change IT professionals' expectations, and these events now act as the benchmark for assessing the requirements of a corporation having a thorough, tested BC/DR plan.

Following 9/11, industry participants met with multiple government agencies, including the U.S. Securities and Exchange Commission (SEC), the Federal Reserve, the New York State Banking Department and the Office of the Comptroller of the Currency specifically to formulate and analyze the lessons learned from the events. These agencies released an interagency white paper, and the SEC released its own paper on best practices to strengthen

DEFINITIONS

Disaster Recovery (DR) is the coordinated process of restoring systems, data and the infrastructure to support ongoing business operations after a disaster occurs. DR concentrates solely on recovering from an unplanned event.

IT resilience is the ability to rapidly adapt and respond to any internal or external disruption, demand or threat, and continue business operations without significant impact. This concept of IT resilience is related to DR, but it's broader in scope in that DR concentrates on recovering only from an unplanned event.

Recovery Point Objective (RPO) is the point in time to which systems and data must be recovered after an outage. RPO refers to how much data your company is willing to re-create following a disaster. In other words, what's the acceptable time difference between the data in your production system and the data at the recovery site? How much data can be lost? What's the acceptable time difference between the data in your production system and the data at the recovery site; i.e., what's the actual point-in-time recovery point at which all data is current? If you have an RPO of less than 24 hours, expect to be able to do some form of onsite real-time mirroring. If your DR plan is dependent on daily full-volume dumps, you probably have an RPO of 24 hours or more. An organization can plan for an appropriate BC technology solution on the basis of the RPO it sets.

Recovery Time Objective (RTO) is the time within which systems and applications must be recovered after an outage. RTO refers to how long your business can afford to wait for IT services to be resumed following a disaster; i.e., the amount of downtime a business can endure and survive. RTO traditionally refers to how long your business can afford to wait for IT services to be resumed following a disaster and how much time is available to recover the applications and have all critical operations up and running again.

the IT resilience of the U.S. financial system.

One thing that wasn't well-appreciated prior to 9/11 was the concept of a regional disaster. A good working definition of a regional disaster is a disaster scenario, man-made or natural, that impacts multiple organizations and/or multiple users in a defined geographic area. In other words, it isn't a disaster impacting only one organization or one data center.

Organizations whose BC/DR plans focused on recovering from a local disaster, such as a fire or power failure within a data center, faced the realization on 9/11 that their plans were inadequate for coping with and recovering from a regional disaster. A regional disaster was precipitated by the World Trade Center attacks in New York City on 9/11. Hundreds of companies and an entire section of a large city, including the financial capital of the world, were affected. Power was cut off, travel restrictions were imposed and the major telephone (land line and wireless) switching centers were destroyed. Access to buildings and data centers was at a minimum restricted and at a maximum completely impossible for several days following 9/11. The paradigm for planning mainframe-centric BC/DR changed overnight. Organizations quickly realized they now need to plan for the possibility of a building being eliminated or rendered useless. Not to be Chicken Little and proclaim "the sky is falling," but what about the threat of terrorists using biological, chemical or nuclear weapons? What about another regional or super-regional natural disaster such as Hurricane Katrina or the 2011 Japanese earthquake/tsunami? Or, man-made disasters such as the 2003 North America power blackout? There are some examples of issues beyond the data center that organizations need to consider in their planning for a regional disaster. Two of these issues deserve the most attention.

The first primary issue beyond the data center is erosion in confidence in a business based on the company's reactions to the

UPTIME (%)	DOWNTIME (%)	DOWNTIME PER YEAR	DOWNTIME PER YEAR
98	2	7.3 days	3 hr, 22 minutes
99	1	3.65 days	1 hr, 41 minutes
99.8	0.2	17 hr, 31 minutes	20 minutes, 10 secs
99.9	0.1	8 hr, 45 minutes	10 minutes, 5 secs
99.99	.01	52.5 minutes	1 minute
99.999	.001	5.25 minutes	6 secs
99.9999	.0001	31.5 secs	.6 secs

Figure 2: Availability Percentage and Number of "9s"

DESCRIPTION	TYPICALLY ACHIEVABLE RPO	TYPICALLY ACHIEVABLE RTO
No DR plan	N/A-all data lost	N/A
Tape vaulting	Measured in days since last stored backup	Days
Electronic vaulting	Hours	Hours (hot remote location) to days
Active replication to remote site (w/o recovery automation	Seconds to minutes	Hours to days (depends on availability of recovery hardware)
Active storage replication to remote "in-house" site	Zero to minutes (dependent on replication technology and automation policy)	1 or more hours (dependent on automation)
Active software replication to remote active site	Seconds to minutes	Seconds to minutes (dependent on automation)

Figure 3: RPO and RTO Values for Some Common DR Options

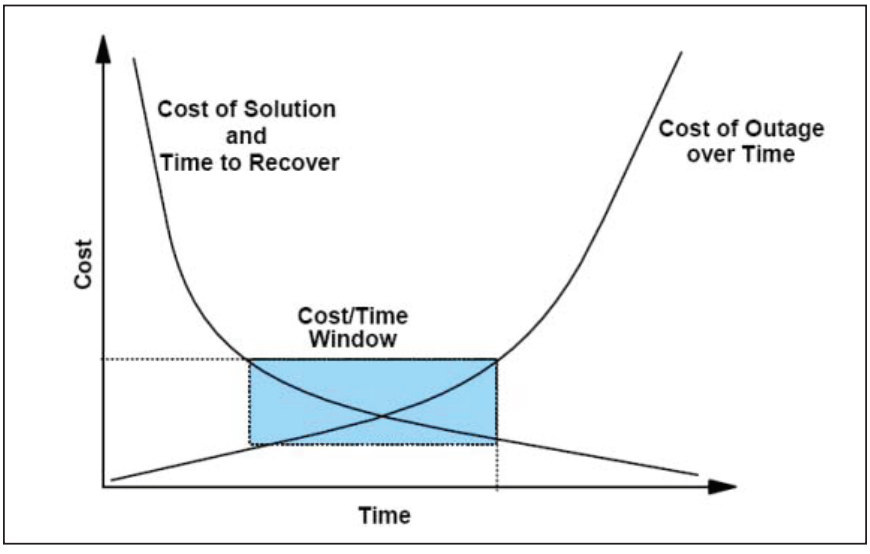

Figure 4: Cost of BC Solutions vs. Cost of Outage

DEFINITIONS - continued

Network Recovery Objective (NRO) is the time needed to recover or failover network operations. NRO includes such jobs as establishing alternate communications links, reconfiguring Internet servers, setting alternate TCP/IP addresses and everything else to make the recovery transparent to customers, remote users and others. NRO effectively is how long before you appear recovered to your customers.

Data vault is a repository at a remote site where data can be periodically or continuously copied so there's always a copy at another site.

Hot site is a site where an enterprise's operations can be moved in the event of a disaster. It's a site with the required hardware, operating system, network and application support to perform business operations where the equipment is available and running at all times.

Cold site is a site where an enterprise's operations can be moved in the event of disaster, with minimum IT infrastructure and environmental facilities in place, but not activated.

Continuous Availability (CA) is a new paradigm that encompasses not only recovering from disasters, but keeping your applications up and running throughout the far more common planned and unplanned outages that don't constitute an actual disaster.

regional disaster. This typically would be based on how the press and internal communications to employees reported on disaster-related events. In other words, what's the perception of the company's ability to react to the crisis, minimize the ensuing chaos, provide timely (and accurate) updates on the status of company operations and (internally) discuss how Human Resources (HR) issues are being addressed? In effect, organizations need to include in their BC/DR plans a communications plan as well as appoint a BC/DR communications director/coordinator. Perception is reality, and the more control an organization is perceived to have over its destiny, and the more upfront an organization is, the less likely there will be an erosion of confidence in the business.

The second primary issue is that an organization needs to be able to restart or execute restoration of business operations in a timely manner. Failure to do so can result in supplies not being able to reach the business to resupply necessary materials. Also, if a business isn't available to customers (i.e., they have no access to its products and services), they may look to competitors. This loss of revenue and customer loyalty has a direct relationship with the effects of a loss/erosion in business confidence.

The New Paradigm

BC is no longer simply IT DR. BC has evolved into a management process that relies on each component in the business chain to sustain operation at all times. Effective BC depends on the ability to accomplish five things:

1. The risk of business interruption must be reduced.
2. When an interruption does occur, a business must be able to stay in business.
3. Businesses that want to stay in business must be able to respond to customers.
4. Businesses need to maintain the confidence of the public.
5. Businesses must comply with requirements

such as audits, insurance, health/safety and regulatory/legislative requirements.

In some nations, government legislation and regulations lay down specific rules for how an organization must handle its business processes and data. Some examples are the Basel II rules for the European banking sector and the U.S. Sarbanes-Oxley Act. These both stipulate that banks must have a resilient back-office infrastructure. Another example is the Health Insurance Portability and Accountability Act (HIPAA). This legislation determines how the U.S. healthcare industry must account for and handle patient-related data.

This ever-increasing need for "365x24x7xforever" availability really means that many businesses are now looking for a greater level of availability covering a wider range of events and scenarios beyond the ability to recover from a disaster. This broader requirement is called IT resilience. As stated earlier, IBM has developed a definition for IT resilience: [It's] the ability to rapidly adapt and respond to any internal or external opportunity, threat, disruption or demand and continue business operations without significant impact."

Conclusion

There are several factors involved in determining your RTO and RPO requirements. Organizations need to consider the cost of some data loss while still maintaining cross-subsystem/cross-volume data consistency. Maintaining data consistency enables the ability to perform a data base restart that typically has a duration of seconds to minutes. This cost needs to be weighed vs. the cost of no data loss, which will either impact production on all operational errors in addition to DR failure situations or yield a data base recovery disaster (typically hours to days in situation), as cross-subsystem/cross-volume data consistency isn't maintained during the failing period. The real solution that will be chosen is based on a particular cost curve slope: If I spend a little more, how much faster is DR? If I spend a little less, how much slower is DR?

This ever-increasing need for "365x24x7xforever" availability really means that many businesses are now looking for a greater level of availability covering a wider range of events and scenarios beyond the ability to recover from a disaster.

In other words, the cost of your BC solution is realized by balancing the equation of how quickly you need to recover your organization's data vs. how much will it cost the company in terms of lost revenue due to being unable to continue business operations. The shorter the time period decided on to recover the data to continue business operations, the higher the costs. It should be obvious that the longer a company is down and unable to process transactions, the more expensive the outage is going to be for the company, and if the outage is long enough, survival of the company is doubtful. Figure 4 takes us back to some basic economics cost curves. Much like deciding on the price of widgets, and the optimal quantity to produce, the optimal solution is the intersection point of the two cost curves. **EE**

Dr. Steve Guendert is a principal engineer and global solutions architect for Brocade Communications, where he leads the mainframe-related business efforts. He is a senior member of the Institute of Electrical and Electronics Engineers (IEEE) and the Association for Computing Machinery (ACM), and a member of the Computer Measurement Group (CMG). He is a former member of both the SHARE and CMG boards of directors.
Email: stephen.guendert@brocade.com
Twitter: @BRCD_DrSteve

