BUSINESS CONTINUITY

by

Saeed Akbani

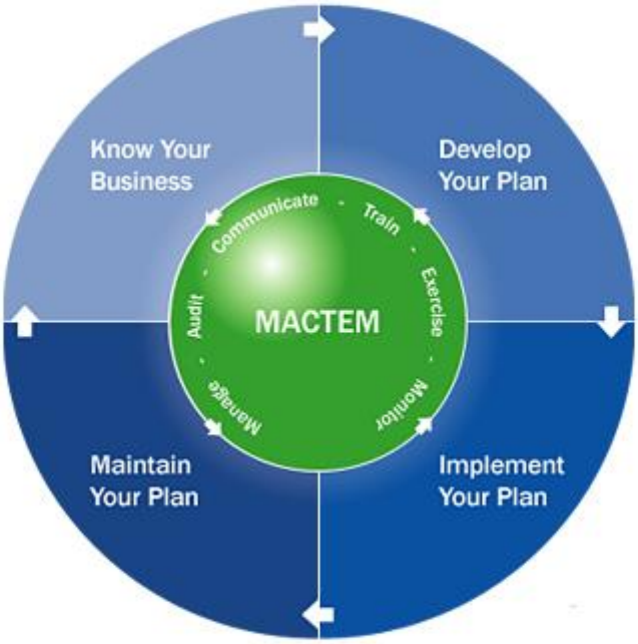

[1]

There is a saying: "Failing to prepare is tantamount to preparing to fail". This is more true in disaster planning than anywhere else. The world of information technology is definitely not immune to disasters. "What would happen to your business if it got hit by a devastating earthquake or flood?" The concept of disaster recovery initially encompassed building redundant data center(s) (backup facilities) away from the primary data centers, where the likelihood of the same disaster such as flood or earthquake destroying both facilities was negligible. Therefore, in case of a disaster, the business could recover most of its data within a reasonable period of time and continue its operation.

According to Gartner, the entire Business Continuity Process deals with a broader aspect of enterprise based recovery planning. This means that it covers an entire business's resources including workspaces, telephones, workstations, servers, applications, network connections and other crucial business assets. [2]

The concept of disaster recovery with initial focus on natural disasters, in recent times has evolved to a more holistic concept of Business Continuity. We define Business Continuity as:

> *"Measures that make a business more resilient and fault-tolerant so that it is able to withstand both large and small shocks while continuing its operations in a normal manner and keeping its digital assets protected"*.

A shock is a disturbance that disrupts normal operation of a business and causes loss of revenues, productivity, and customer goodwill. The shocks can come in the form of natural disasters like fire, flood, lightning and earthquake. These shocks can be caused by terrorist attacks, hackers, or can result from hardware or software failures. However, the shock can also come in the form of increased business activities (such as unexpected surge in website traffic) that digital assets are not designed to handle thus causing unexpected delays in transactions and worse, system crashes.

Today, business continuity encompasses the dimensions discussed below. These dimensions carry operational meanings on their own, but should not be considered independent of each other, as considerable overlap exists between them.

Disaster Recovery (DR)

As the name implies, disaster recovery simply means recovering from a loss of one or more digital assets that disrupts normal operation of a business. Some common disasters that businesses encounter and must be prepared to handle include:

- <u>Hardware Failure:</u> This typically occurs when a vital component of a server (or a PC) such as hard drive, memory, motherboard, etc. fails. Most servers come with a RAID array to make the server more tolerant of hard drive failures. There are different levels of RAID, all of which offer data redundancy, but more advanced levels allow machines to continue operation even after the failure, thus minimizing disruptions until the damaged hard drive is replaced. However, there is typically no redundancy for other components such as memory and motherboard.
- <u>Operating System Corruption:</u> This is a common problem with Windows where the operating system becomes corrupted. Until a few years ago, there was no other way, but to re-install the operating system. Now Windows creates checkpoints at specific intervals, which allows a network engineer to restore the OS to a previous checkpoint in minimal time in the event of corruption. The restoration typically takes significantly less time than re-imaging the system, but it is not always guaranteed to succeed. There are faster methods available which are discussed below.
- <u>Data Loss:</u> Data vital to a company's operation can become corrupted as a result of a hardware failure, OS corruption, or simply a glitch in the software used to create (or modify) the data. In such a case, data (contained in a database) must be restored from a mirroring or from another backup facility.

Failover

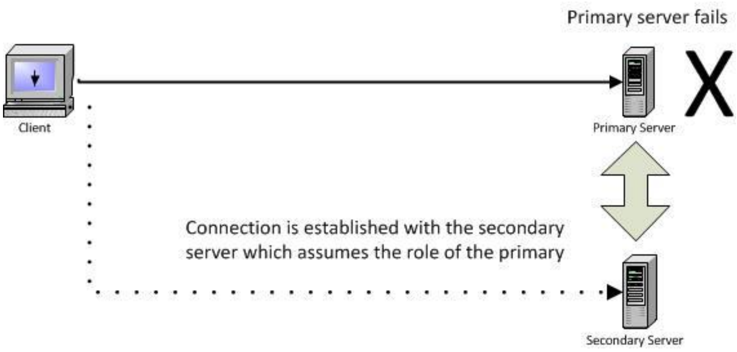
Primary server fails
Client
Primary Server
X
Connection is established with the secondary server which assumes the role of the primary
Secondary Server

Figure 1 – Concept of Failover

The principle of *Failover* calls for operations to be switched over (in most cases automatically). When a machine (usually a server) becomes inoperable, traffic is redirected to another (mirrored) machine. The other machine (usually a standby server) serves the purpose of redundancy. A simple configuration for failover includes a primary server and a secondary server where the data (along with the rest of the development) is mirrored. When the primary fails, the secondary server takes over the job of the primary, until the primary is fixed. Once the original primary server is revived, the roles may be reversed or it can continue its job as the secondary. The primary-secondary combination of failover is also helpful during maintenance of one of the machines.

The Failover concept guarantees minimal disruptions to a business through redundancy, but this comes at a substantial cost (at least double). Not only does a business require twice as many servers, but it may also have to upgrade its database systems. For example, the mirrored server concept in SQL Server is not available in the Express version, but rather in the Enterprise edition, which costs several thousand dollars more than the Express version. B2B and B2C firms that rely heavily on internet traffic for sales must take Failover into account or risk losing sales and customer goodwill.

Load Balancing

A sudden increase in business activity can become a curse for a business if its IT assets present bottlenecks during the surge. A business, that handles hundreds of thousands of transactions per day, will not have just one server, but a collection of servers or server farm working as a group. The concepts of load balancing and scalability (discussed later) are intertwined. The goal here is not to discuss these concepts in detail, but to raise readers' awareness. At its very core, the idea of load balancing is to distribute the computational load across servers so as to increase efficiency and minimize response time to users in a cost effective manner. One common way to achieve load balancing is to separate the web server from the database server, with the former dedicated to handling HTTP requests and performing computations, while leaving complex database transactions on one or more dedicated database servers. Another way to achieve load balancing is to place load balancers in front of servers (web or database) that accept requests (such as HTTP) and route them to appropriate servers (running in parallel) so as to balance the load evenly among servers. This all happens behind the scenes and appears completely transparent to the user.

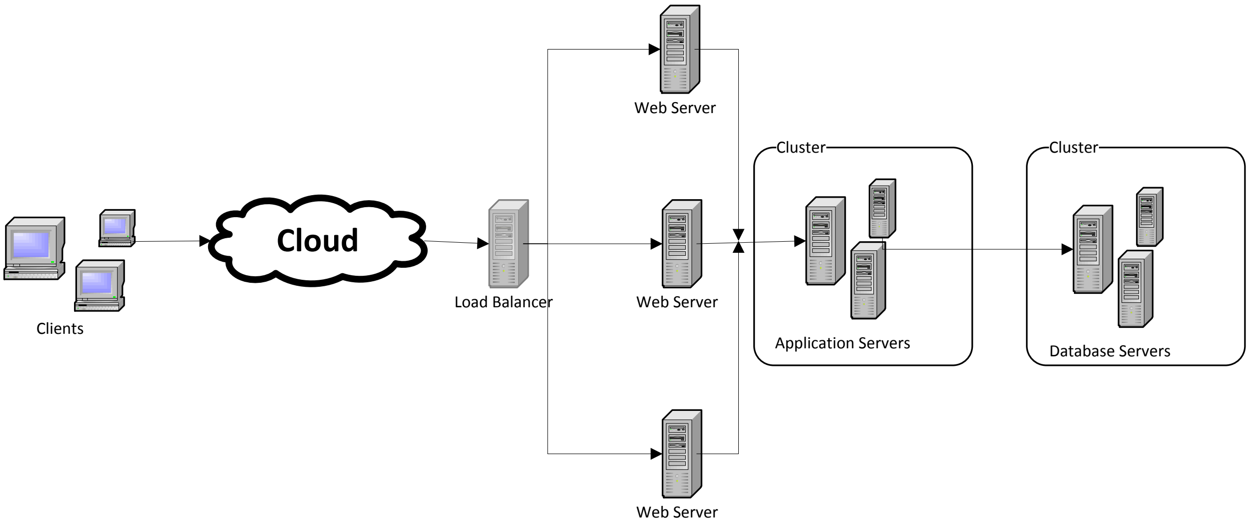

Figure 2 – Load Balancing

Scalability

Just because your business is experiencing 1,000 unique website visitors per day does not necessarily mean that it will stay that way next year. In fact, successful social media startups can experience exponential growth in website visitors from one year to the next. While this may not be true for most brick-and-mortar businesses, they must still think of investing in systems that are scalable when businesses experience growth.

Investing in scalability does not necessarily mean that you invest in infrastructure with enough capacity to handle several times the current work load, but that the design is flexible enough to increase capacity by incrementally adding to the infrastructure. A common example of scalability is when the data store (database) being loosely coupled with the application, can be moved to its own server hardware (often referred to as database server) for improved performance. An even simpler example of scalability can be found in server purchase. If the memory can be easily increased from 8 GBs (current) to say 32 GBs, when needed, by the purchase of 24 GBs of additional memory, you have addressed scalability in a small way.

The cost of incrementally adding to the system is a lot lower than to throw away and start again (24 GBs of additional memory compared to a new server). If the system is not scalable, load balancing may not be as effective, and the infrastructure is bound to experience bottlenecks resulting in outages. While it may be true that the cost of building a relatively scalable system is generally higher than of not, the payback in the medium to long term in most cases easily justifies the additional expenditure.

High Availability (HA)

As the name implies, HA is an approach to system design encompassing concepts discussed above (DR, redundancy, scalability, load balancing, and failover) to achieve uptime thus satisfying the requirements of a given SLA (Service Level Agreement). The availability (uptime), usually, expressed as a percentage (for example 99%), guarantees that the system will be running and catering to the needs of its users. When the system is unavailable it is considered down. Downtimes could be scheduled or unscheduled, and typically scheduled downtimes, for software, hardware, or network maintenance, are excluded from the availability calculations. Hence the expected unscheduled downtime probability can be computed as 1 – Availability. Therefore, one can expect 1%

unscheduled downtime for 99% availability. The table below (source: Wikipedia) provides details for two (2) cases of availability.

Availability	Downtime per Year	Downtime per Month	Downtime per Week
99%	3.65 days	7.20 hours	1.68 hours
99.8%	17.52 hours	86.23 minutes	20.16 minutes

Mirroring or Backup, Which one is better for redundancy?

At a fundamental level the difference between mirroring and backup is the frequency with which data (and maybe the image) is replicated across systems. In a backup situation, data may be replicated on an hourly, daily, or weekly basis, whereas, in mirroring the frequency of replication is more granular such as at the end of each transaction or event. Typically, as the granularity of level of replication increases, the delta between the primary and replicated system decreases. Therefore, the loss of data in case of an outage decreases. However, higher frequency comes with higher cost. The cost of mirroring is substantially higher than that of backup. As discussed above, mirroring is typically available to enterprise class users by software manufacturers like Microsoft, and small-medium sized businesses may find backup as the most cost effective solution.

Disaster Recovery Time

While it is always the desire to recover and resume business operations from a disaster as soon as possible (maybe even instantly), the price tag of instant recovery may be cost prohibitive. If the target recovery window from a major disaster is small, the price tag becomes large. Businesses, therefore, must be realistic and implement a disaster recovery plan that balances cost with the right recovery window, while at the same time investing in those areas that make the business resilient to small shocks, such as power failures, hacking, etc. which typically have lower price tags.

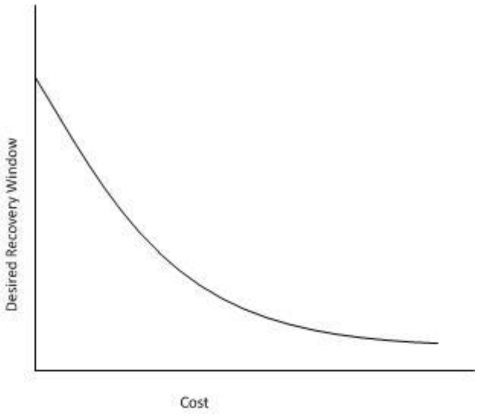

Figure 3 – DR Time vs. Cost

Cost Justification: How much should I be willing to pay for DR?

While companies selling DR implementations and BDR (backup and disaster recovery) appliances would like every business to buy the most sophisticated solution that guarantees resumption of business activity within seconds, the cost of such plans may be cost prohibitive for small-medium sized businesses. No matter how large or small, a business must go through some analysis to determine the maximum amount the business should be willing to pay for a DR plan. The DR plans may be different for different operations being supported by IT. For example, a server handling customer sales may require a more robust DR plan than a server handling accounting functions.

The following may provide a framework for determining the cost and targeting the right DR plan for a business operation.

	Outage Description:	**Amount**
1	Current Recovery Window (Hours)	
2	Probability of Occurrence (in a year)	
3	**Cost:**	
4	Indirect Costs	
5	Lost Revenues	
6	Loss of Customer Goodwill	
7	Loss of Productivity	
8	Total Opportunity Cost (add lines 5-7)	
9	Direct Costs	
10	Hardware Replacement	
11	Software Replacement	
12	Labor	
13	Total Direct Costs (add lines 10-12)	
14	**Total Cost (add lines 8 and 13)**	
15	**Expected Yearly Cost (multiple lines 2 and 14)**	

Table I – DR Cost Determination Framework

The above table provides a simple framework for determining the expected cost of an outage or a disaster for a business. Just like auto insurance for a motorist, a business must be willing to bear the cost for DR with a recovery window, whose price tag is equal to the expected cost of the disaster. A DR window much shorter than this with a significantly higher price tag means that the business may be overspending. Additionally, a business may lower the probability of an occurrence (line 2 in the table) by investing in other business continuity initiatives (failover for example), thus making IT more resilient to disasters.

Mission vs. Non-Mission Critical Systems

Obviously, the framework presented above would reveal much higher cost of outage for mission-critical systems as opposed to non-mission-critical systems. The loss of revenues and customer goodwill may be significant for mission-critical systems, and may be very low for those that are not. Hence a business may be willing to tolerate a longer recovery window for non-mission-critical system(s) than for mission-critical systems. It is for this reason we emphasize a 2-tiered approach to DR implementations, one that calls for a shorter recovery window (including failover) where both data and image are backed up and another one for non-mission-critical systems, where only data may be backed up.

Onsite vs. Cloud (Offsite) Backup

Most IT experts would agree that onsite backup is not enough, as the information remains too close to its origin and, therefore, remains susceptible to disasters. If the office building goes up in flames, onsite backup will most likely be of no use. This is where cloud backup becomes extremely valuable, as it provides businesses the opportunity to backup information to a remote location. With the price of cloud storage is decreasing rapidly, it would be negligent for a business to ignore this option. The old-fashioned method of copying server data to a flash drive or a tape and taking it home on a daily basis is inefficient, risky, and time consuming, especially given the low cost of cloud backup.

We, recommend adopting a 2-tiered approach consisting of onsite and offsite backup. The reason for this is data restoration time, which is much faster for data that has been backed up locally. For a local outage where the local backup is available, there is no need to rely on offsite backup. Local storage (in most cases a NAS or network attached storage) carries a one-time cost of only a few hundred dollars.

Backup Data or Data and Image Both?

At the very least, a business must backup data. However, whether a business should backup image (which includes OS, database, applications, registry settings, etc.) as well depends on the cost of disaster (see table above) and the recovery window desired. However, backing up the image may add significantly to the cost of both storage and internet bandwidth usage. If only data is backed up, then the recovery window is expected to be larger given the time it would take to restore the server (hardware and OS), and load the database and application(s) to their original state. Hence, for mission-critical systems where the business demands the shortest recovery window possible, it would be wise to back up both.

What is a BDR Appliance?

A backup and disaster recovery appliance is a device (in some cases as simple as a NAS) that uses VM (virtual machine) technology to combine disaster recovery with failover. A BDR appliance could come in various configurations offering local backup, cloud backup, or a combination of both. A BDR appliance backs up data as well as image (operating system) of the server. In case of a disaster, the BDR appliance could assume the role of a secondary server taking care of internal and/or external user requests, until the primary server is restored. This method ensures very little business downtime, if any. The cost of BDR appliances remains relatively high, and a business is advised to first go through proper analysis before selecting its BDR technology.

Business Continuity and Cloud Computing

While it is true that Cloud Computing models, such as Software as a Service (SaaS) and Infrastructure as a Service (IaaS) shifts the onus of disaster recovery to the service provider, a business client must ask the right questions to make sure it is adequately protected in case of a disaster. Some of the questions that could be asked include:

- Is my data backed up, if so where and how often?
- Is my image backed up, if so, where and how often?
- Am I getting failover?
- What is expected DR window?

A business must make sure that the right level of business continuity exists. Additionally, going with the cloud computing model may create vulnerabilities, such as the need for redundant internet connection, for a business which may not have existed before.

Other Considerations

Business continuity is more than just disaster recovery. Businesses must also look at failover, load balancing, as well as other disaster avoidance measures.

Selecting Server with Proper Configuration

While hardware manufacturers like to entice their customers with teaser prices, the price of a server class machine ends up being substantially more (3-5 times the teaser price). A business should make sure that its server purchase includes resilience and scalability.

Feature Type	BC Addressed	Description
RAID	Failover, Data Backup	A typical RAID consists of multiple drives where the data is replicated across drives. A more advanced RAID such as RAID 10 offers failover in that in case of a hard drive failure, the server can continue to operate until the damaged disk is replaced.
Redundant Power Supply	Failover	In case of failure of the main power supply, the secondary can take over
Ample Memory	Scalability	With the price of memory low, it is always a good idea to go with at least twice as much memory as you need. In addition, it is a good idea to make sure that memory can be increased incrementally in the future as opposed to scrapping original sticks that came with the machine.
UPS	Disaster Avoidance	Will protect server against power glitches and power outages. In case of a power outage, server will have time to shut-down in a normal manner thus avoiding disk, application, or OS corruption
Warranty	Disaster Recovery	Purchasing at least 3 years of warranty with NBD (next business day) or even the same business day will make sure the OEM backs its product for the certain period of time, thus preventing costly repairs.

Redundant Internet Connection

When a mission-critical operation relies on the internet connection, it is essential to have a redundant connection from a different supplier in case of an outage at the primary source. The framework presented in table 1 may be used to determine how much bandwidth in a redundant internet connection is necessary.

IT Security

A security breach can have more serious consequences for a business than a fire or flood, for example. The disruptive nature is, in most cases, no different. Hence insuring adequate security measures is part of business continuity. A security breach can occur from within or from the outside.

Security Aspect	Vulnerabilities	What can a business do
Network Security	Network and data could be compromised from outside	Make sure to have business class routers and firewalls in place with intrusion, virus, spyware, and malware, etc. detection being done at both the router level and also at the PC level. Make sure to have business sensitive data encrypted if possible. Passwords must be complex.
Cloud Applications	Business and customer data being stolen	Make sure the software application complies with minimum standards (HIPAA, PCI, etc.)
Policies and Procedures	Internal and External	Make sure to have SOP that is well communicated to all employees and vendors. A frequent audit may be necessary to ensure compliance

Succession Planning

A common mistake made by businesses is relying on their resources (employees and vendors) without any succession planning in place. Employees quit or are terminated, and relations with vendors are severed. No matter how great the relationship and trust, a business must plan for a smooth transition at some point in the future. It is essential to document employees' roles and responsibilities along with all the IT information (network topology, hardware/software landscape, email accounts, etc.). This way, the next vendor or employee can take over in case an employee or a vendor's services are no longer available.

DR Drill

A DR drill in principle is no different than fire drills in commercial buildings. Just because your data is being replicated somewhere, does not mean you can forget about it. Commercial backup software do provide notifications of completion of backup jobs, but it only means the process was run. In a changing IT environment, a backup job setup 12 months ago, may not be picking up recently created new folders. Frequent DR drills to simulate disaster conditions and recovery should be part of the process. At minimum, the goals of a DR drill should include the following:

- All the required data is getting replicated
- All the replicated data (along with environments) is recoverable. This mean it can be restored, and can also be accessed by the users
- Restoration can take place in the desired time frame (target DR window)

We recommend quarterly DR drills, and if this is not possible or practical, then at least once every 6 months.

Conclusion

In the year 2010, Forrester surveyed 2,803 IT decision makers. They discovered that the second highest priority of any business enterprise is improving their business continuity process and making their disaster recovery process more efficient. Moreover, the overall scope of Business Continuity and Disaster Recovery is also maturing day by day. Now companies have started to invest more in these processes. [3]

A business cannot continue its operations without adequate planning to withstand external and internal shocks to its digital assets. Investment in business continuity is akin to purchasing an insurance policy. Commercial policies may cover hard assets of a business, but the business needs to invest in protecting its data, loss of productivity, revenues and customer goodwill, and above all its reputation.

About the author:

Saeed Akbani is the president and CEO of Data Dynamics, Inc. He can be reached at (314) 438-5478 or saeed.akbani@datadynamics-inc.com.

How can Data Dynamics help?

Data Dynamics offers a whole array of online strategies to help our clients grow and dominate their industry. Having the necessary expertise, we can evaluate various technology options against your business requirements and select and/or develop solution to help your business gain a competitive advantage over your rivals in the most cost effective manner possible. Want to find out more? Contact us at (314) 438-5478 or email us at info@datadynamics-inc.com. You can also visit us at http://www.datadynamics-inc.com.

References

1. "Business Continuity Lifecycle." *Business Continuity and Disaster Recovery Planning Software*. N.p., n.d. Web. 12 June 2013. <http://www.erlogix.com/crisis_management.asp>.
2. "IT Glossary." Gartner. Web. 04 June 2013. <http://www.gartner.com/it-glossary/bcp-business-continuity-planning/>
3. RBZane Advisory Group - Business Continuity Consultants and Disaster Recovery Services."*RBZane Advisory Group - Business Continuity Consultants and Disaster Recovery Services*. Forrester, Web. 04 June 2013. <http://www.rbzaneadvisors.com/pdf/Forrester_bus_cont_disaster_recovery.pdf>

For more information, contact:
info@datadynamics-inc.com