

Business Continuity & Disaster Recovery Planning

By Dr. Jim Metzler,
Ashton, Metzler & Associates

Sponsored By:



Produced By:



Table of Contents

1. Introduction	3
2. The Movement to Implement Contingency Planning	3
3. A Characterization of Disaster Recovery and Business Continuity	5
4. The Continuum of Disaster Recovery Options	5
5. The Key Components of a Contingency Plan	6
6. High Availability Network Design	8
6.1 Robust LAN Design	8
6.2 Robust WAN Design	9
7. Technologies that Enable Contingency Solutions	10
7.1 Load Balancing for Servers and Contact Centers	10
7.2 Wireless Solutions	10
7.3 Expert Systems	10
7.4 MPLS Enabled Frame Relay WAN Services	10
7.5 Content Distribution Networks	10
7.6 Back-Up Power	11
7.7 Re-direction of Phone Calls	11
7.8 VoIP	11
7.9 SRDF	11
8. Other Techniques that Enable Contingency Solutions	11
8.1 Network Simplification	12
8.2 Geographical Diversity	12
8.3 Hosted Solutions	12
9. Processes that Enable Contingency Solutions	12
9.1 Link to Other Key Processes	12
9.2 Perform a Regular Inventory Analysis	12
9.3 Have Arrangements for Back-up Equipment	12
9.4 Create an IT Crisis Management Team	13
9.5 Identify and Document Methods of Efficient Communications ..	13
9.6 Establish a Succession Plan	13
9.7 Educate the Employees	13
9.8 Testing the Plan	13
9.9 Recovery	13
10. The Funding of Contingency Planning	13
10.1 IT Governance	14
10.2 Business Case Analysis	14
11. Summary	16
Business Continuity at Coca-Cola Enterprises	16
Business Continuity at EMC	17



Dr. Jim Metzler

In addition to running two network organizations, Jim Metzler has worked in software development, network engineering, marketing, product management, consulting, and market research.

Jim's current interests include Business Continuity, Outsourcing, Voice over IP, and IP-based VPNs. He is currently beginning the deployment of The Business of IT Initiative. The goal of this initiative is to focus the industry less on technology and more on the ability of IT to add business value.

Jim can be reached at jim@ashtonmetzler.com.

1. Introduction

Disaster recovery and business continuity planning are not new topics. Both have been around for a number of years. However, while the concepts have been around for a long time, they were not that widely implemented. In addition, the catastrophes that were envisioned in most disaster recovery or business continuity plans were somewhat narrow in scope. As one financial firm in New York City explained: “We had business continuity planning prior to September 11. One part of that plan was to fly our traders to London where they would be as functional as they had been in New York City. Who knew that they would shut down all of the airports?”

For the purpose of this document, the phrase “contingency plan” will be used to refer to a plan that contains some combination of disaster recovery and business continuity planning. From the perspective of the network executive, implementing an effective contingency plan is one very visible way that the network function can demonstrate value to the business. To assist these network executives, this paper will create a framework for the successful creation of a contingency plan. This framework will focus on the Information Technology (IT) component of the plan in general, and on the network component of this plan in particular. It is intended that this framework will be modified to fit the needs of virtually any enterprise, whether they are looking to create a contingency plan that is somewhat simple and narrowly focused, or one that is more complex and broadly focused.

Part of the framework that will be developed in this document is a characterization of what is generally meant by the phrases “Disaster Recovery” and “Business Continuity.” Included in this framework will be a discussion of some of the key technologies, techniques and processes that comprise the various solutions to a company’s contingency requirements. Note that no attempt will be made to create anything close to an exhaustive list of relevant technologies, techniques and processes. The goal of this discussion is merely to exemplify the vast array of technologies, techniques and processes that can be part of a contingency solution.

The framework that will be developed in this document will also discuss funding. In particular, the framework will include some analysis as to how the IT organization can engage in discussion with the business decision makers in order to ensure that IT receives sufficient funding to implement the appropriate level of contingency planning. As will be demonstrated in this paper, it is not possible to have a meaningful discussion about contingency planning without also discussing funding.

This document is not designed to be yet another abstract dissertation on how to do contingency planning. To avoid that fate, this document will incorporate insight from a number of leading

edge companies who either already have deployed an effective contingency plan, or who are currently in the process of doing so. Note that the internal policies of a few of these companies preclude their being specifically identified in this document.

2. The Movement to Implement Contingency Planning

In March of 2002, Ashton, Metzler & Associates, in conjunction with Key3Media, asked over 600 network professionals to indicate the primary new technology that they would spend time on over the next 12 months. Their answers are depicted in Figure 1 (see page 4).

What Figure 1 indicates is that for many companies, implementing a Disaster Recovery/Business Continuity plan has become a priority. One company for whom Disaster Recovery and Business Continuity is critical is Coca-Cola Enterprises (CCE). According to John Ridley, Sr. Enterprise Network Architect at CCE, CCE is currently working on a Disaster Recovery/ Business Continuity project that is anticipated to cost more than 400 million dollars, the largest IT project in the history of CCE (See Story, “Business Continuity at Coca-Cola Enterprises,” page 16).

Peter Brown, Director of Architecture and Cross Border Services for PricewaterhouseCoopers’ internal IT organization, points out the critical leadership role that IT needs to play in the development of a company’s Disaster Recovery/Business Continuity plan.

“Many times the IT function needs to facilitate the process that results in the Disaster Recovery/Business Continuity plan,” Brown says, because “it is only within the IT function where all of

NSI SOFTWARE **Protect Your Data without Boundaries**

Visit the **NSI® Software** Web site to learn more, and to receive the top five disaster recovery tips **EVERY IT manager must know.**

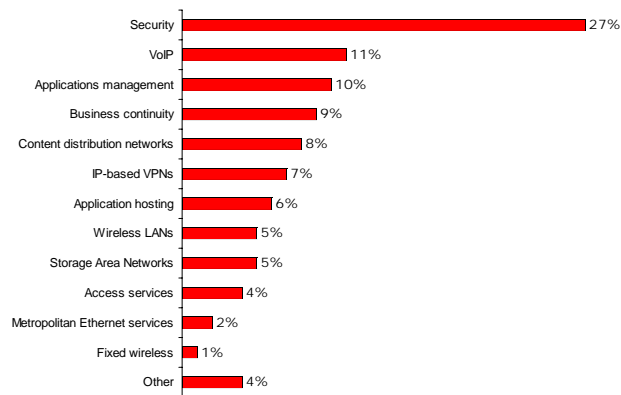
Double-Take® is the replication engine used by **HP, Dell and IBM** to protect data on select **NAS and server products.**

Double-Take®
True, real-time data replication for Windows

the various elements of the business come together in a manner such that the company's vulnerabilities are visible."

Suppose, Brown says, "A plant within the meatpacking division of a company determines that they don't need and can't pay for Disaster Recovery/Business Continuity. However, what is visible

Figure 1: Technologies on Which the Most Time Will be Spent in 2002



Source: Ashton, Metzler & Associates and Key3Media

to the IT organization is that if the gateway to transmit and receive orders is down, then no one knows what to ship, where to ship it, or who to bill. Based on this visibility, IT is in the best position to facilitate the dialogue so that the business leaders understand the criticalness of this 'simple gateway' and the need for a Disaster Recovery/Business Continuity plan to keep the facility, and perhaps the division, operational."

While companies of all sizes and in all industry segments know that they need to implement more effective contingency planning, most companies are struggling to answer some key questions relative to the planning that they know that they must do. These questions include:

- What exactly is business continuity?
- Is business continuity just disaster recovery with a new fancy name?
- How does one achieve business continuity?

In the abstract, disaster recovery and business continuity represent extreme viewpoints in terms of how to plan for a catastrophic situation. For example, a strict interpretation of what is meant by disaster recovery is that it is concerned with recovering from a disaster once that disaster has occurred. At the other extreme, a strict interpretation of what is meant by business continuity is that it is concerned with keeping some number of key business processes functioning in spite of a catastrophe.

Mark Endry, Senior Vice President and CIO of J.D. Edwards & Company, has a somewhat broad interpretation of the phrase Business Continuity.

"When I hear the term Disaster Recovery (DR) I generally think about the process of putting Information Technology resources back into production after a catastrophe. When I hear Business Continuity Planning (BCP) I think about the much larger challenge of getting the entire business back into production. To some extent, the Business Continuity Plan ensures the results that come from putting an excellent Disaster Recovery Plan into action can be put to good use. Having all of the IT resources up and running is not much use if the business is not prepared to make use of it," Endry says.

The importance of tightly integrating IT's contingency planning into what is being done by the Business Unit Managers is echoed by Kirk Corkery, Corporate Chief of iSERV, a branch of the Canadian Government that provides all of the IT to the province of Toronto. "In my operation there are two applications for which I can completely recover the system after a catastrophe. However, because the user controls the output and has not implemented a business continuity plan of their own, the application will be useless because there is no output."

It should also be noted that there are tight linkages between developing a contingency plan and developing a security plan. For example, Terry Dymek, Director of Global Telecommunications for EMC's internal IT organization, says: "It is through a careful assessment of your company's operation that a viable business continuity plan is developed. Each business process should be examined in detail to determine its sensitivity to the loss of current data and the ability of people to operate from various locations. Once that analysis has been performed, the appropriate business continuity

Enjoy uninterrupted nights with

BACKUP EXPRESS.

The perfect backup and restore partner.

Get a FREE Trial and FREE Nightshirt.

solution can then be implemented” (See Story, “Business Continuity at EMC,” page 17).

The development of a security plan requires performing a very similar risk analysis to the one described by Dymek. In addition, a security breach such as a virus attack can certainly impact an organization’s ability to function. However, to manage the size of this document, the framework that will be developed in this document for contingency planning will not directly address security issues.

3. A Characterization of Disaster Recovery and Business Continuity

Disaster Recovery in this document will be viewed as the subset of Business Continuity that deals solely with computer systems. By that is meant that the typical disaster recovery plan does not take into account other resources, such as a company’s employees and the daily interactions that they have with customers, suppliers, and distributors.

In contrast, the scope of a business continuity plan that focuses on IT includes everything that is covered by disaster recovery, and also typically includes other functions, including:

1. Implementing wireless voice and data connectivity to provide connectivity in case wireline services are not available;
2. Providing for “hot stand-by” office space as well as voice and data functionality for key personnel in case a given building becomes inoperative;
3. Implementing backup call center capabilities in case a call center becomes inoperative;
4. Establishing, communicating and testing alternative processes.

Serge Minassian, Vice President of IT Strategy Planning and Relationship Management for Avaya’s internal IT organization, distinguishes between the relative importance of Disaster Recovery and Business Continuity: “As companies become more global and operations become more centralized, business continuity is becoming an increasingly more important option over simple disaster recover. Imagine not being able to enter sales orders in Asia because of a data center fire in the U.S.”

The next section of this document will discuss a variety of ways to implement a Disaster Recovery solution. One of these solutions involves a hot stand-by data center. If a hot stand-by data center has been implemented, then the data at the primary data center is being continually replicated at the stand-by data center. If there is a failure at the primary data center, the stand-by data center becomes the production data center in a matter of milliseconds.

As previously mentioned, a business continuity solution could well include hot stand-by office space. However, even in those situations in which back-up office space has been arranged, the switchover from a production office building to the back-up office space will most likely take several hours. One of the major reasons that it takes several hours to shift over to the back-up office space is that it will typically take several hours for the affected personnel to get from their homes or normal office building to the back-up office space.

The issue of where to have the back-up office space exemplifies some of the key trade-offs that have to be made in designing a contingency plan. It is clearly desirable to have the back-up office site close to the production office space, because this reduces the time and expense that it will take to get the employees to the back-up office site. However, having the back-up office space close to the production office space increases the probability that the catastrophe that made the production space unusable will also make the back-up space unusable.

4. The Continuum of Disaster Recovery Options

This section of the document has two goals: First, to demonstrate the continuum of alternatives that can be used to implement a Disaster Recovery solution. The second goal is to use for this demonstration a topic that is a key component of any contingency plan — data protection.

Note that it is not the goal of this section to develop a complete description of data protection. What this section will do is to describe some of the more popular approaches to providing data protection.

YOUR FIRST LINE OF DEFENSE

PROTECT YOUR ASSETS & SPACES

MONITOR & RECORD ACCESS

GET EARLY WARNING

PREEMPT PROBLEMS

NETBOTZ™

INTELLIGENT MONITORING OF CRITICAL ASSETS AND SPACES

The advantages listed for each of the following scenarios assume the contingency plan is appropriately designed. For example, the primary advantage of having a hot stand-by data center is that the company will continue to have uninterrupted access to its applications even if a catastrophe has occurred. However, if the contingency plan is not appropriately designed and the same catastrophe takes both the primary and the back-up data center offline, then this advantage will not be realized.

Scenario #1: Tape Backup OffSite

In this scenario, a company writes the contents of its key databases onto tape backup on a regular basis. Daily backup is common. Those back-up tapes are then sent via truck to some remote site for storage. If the company does encounter a disaster at a data center, the company contacts the site that is storing their back-up tapes and has the tapes sent via truck back to the company.

The primary advantages of this approach are that it is easy and low cost. One disadvantage of this approach is that any data that was created subsequent to the last tape backup is lost. Another disadvantage is that, depending on the type of disaster that occurred, the affected data center may be inoperative for days.

Scenario #2: Electronic Vaulting

This scenario is similar to scenario #1, except that it eliminates the need for trucks. In this scenario, the company's key databases are transmitted over a Wide Area Network (WAN) to a remote site. Organizations that implement this scenario typically back up their data more frequently than they would if they had implemented scenario #1. Hourly updates are common. As such, one clear advantage of this scenario vs. scenario #1 is that there is less lost data. However, this scenario typically costs more than scenario #1 and does not reduce the amount of time that the affected data center is inoperative.

Scenario #3: Remote Disk Mirroring

This scenario is similar to scenario #2. The primary difference between the scenarios is that an organization that implements scenario #3 performs the updates on a continuous basis. As such, this scenario has the advantage of eliminating the possibility of lost data. This scenario, however, also tends to cost more than scenario #2 because it tends to require higher capacity WAN links. This scenario also does not reduce the amount of time that the affected data center is inoperative.

Scenario #4: Implement a Cold Site

In each of the preceding scenarios, the data center that was affected by the disaster could be inoperative for an indefinite amount of time. In this scenario, a cold site is established as a backup to the data center. A cold site is a site that will have the

equipment installed and configured only after a contingency plan has been invoked.

This scenario can be implemented in conjunction with any of the preceding scenarios. The advantage of this scenario is that it limits the amount of time that the company cannot access its key applications. The disadvantage of this scenario is that it adds additional cost and complexity.

J.D. Edwards' Endry notes the complexity associated with this approach. "Companies opting for the cold site option need to plan carefully for the equipment at the cold site. For example, if they make their tape backups with older tape drives they might have a difficult time obtaining the same kind of drive on short notice to reload the tapes at a cold site."

Scenario #5: Completely Duplicated Hot-site

This scenario exemplifies what is thought of as a Business Continuity solution. It differs from scenario #4 in that at the remote site there are also mainframes and/or servers that are capable of running all of the appropriate applications. The advantage of this scenario is that if a catastrophe does make a data center inoperative, the remote site can take over operations with little or no impact. This, of course, assumes that the catastrophe did not also affect the remote data site. The disadvantage of this scenario is that it is expensive.

5. The Key Components of a Contingency Plan

In order for network professionals to choose the appropriate contingency plan, they need to be able to answer the following questions:

Multihoming to the Internet is a key component to any business continuity strategy - improving the reliability and redundancy of IP networks. However, multihoming is only the first step. Join Sockeye Networks, live - via the web, to explore technologies that improve network performance, stability and visibility in a multihomed environment.

June 27, 2002
2 pm EST (11 am PST)
To register visit
www.sockeye.com

Question #1: What functions are so business critical that it is worth it to the business to provide protection against a catastrophe for these functions?

PricewaterhouseCoopers' Brown points out the difficulty of generating simple answers to this question. "The key applications that require continuity vary greatly by industry. In the professional services industry, remote access and e-mail are mission critical applications, whereas in a manufacturing environment your ERP systems may be most critical. It is absolutely necessary to get the information on application criticalness from business owners, not their IT counterparts, because when value/cost is associated with a potential non-availability of an application, business management often times sees different priorities than do IT professionals."

Avaya's Minassian advises that just knowing what functions are business critical is not enough. "It's important not only to know what is business critical but also if the parameters are expected to change during a disaster. For example, are more or less users expected? Are reduced response times acceptable? Knowing the answers to these questions can change your approach to BC/DR, your design, and your cost."

Question #2: How much protection is needed for each component of the business; i.e., is it acceptable to be up and running the next day?

ISERV's Corkery provides some insight into the issues that surround answering question #2. "In 1984 I had the opportunity to be involved in the recovery of a major data center for a large telco following a fire. Though the building was not physically damaged, the electrical systems had to be rewired, and this took about four weeks. Although we had classified all applications into either "Vital" or "Non-vital" and only the Vital applications were scheduled to be recovered, over the four weeks every application became vital to the business user. As a result we re-developed the list into Maximum Tolerable Outages. This provided the needed prioritization of what had to be up when but recognized that everything eventually did."

Question #3: Against what types of catastrophes is the solution supposed to protect against? These types of catastrophes will be referred to below (see Table 1) as being Type 1, Type 2, and Type 3.

Question #4: What solutions are available? What is the cost of the proposed solution(s)?

Question #5: Who is the decision maker for the contingency plan? What is their willingness to pay for the identified levels of protection?

Endry believes it is not possible to have a meaningful discussion on contingency planning without discussing funding. "Many

companies identify electronic mail as a critical business function. They need to then determine if 100% of the electronic mail capacity needs to be duplicated, or if they can operate for a while with reduced capability. The answer to this question can significantly impact costs."

Typically, the cost of a contingency plan is directly related to the answers to the questions that are listed above. In particular, the cost of a contingency plan increases with an increase in the set of functions for which protection is needed, the amount of protection that is required, and the types of catastrophes that the solution is supposed to protect against.

As mentioned previously, the events of September 11 have turned contingency planning into a hot topic. However, there is currently a tendency to think of contingency planning only relative to protecting against monumental events, such as the terrorist attacks that occurred on September 11. Such an approach is shortsighted. A better approach, when thinking about question #3, is to delineate three types of catastrophes, as depicted in Table 1.

Table 1: Classes of Catastrophes

Type of Catastrophe	Characterization of Catastrophe	Expected Length of Outage
Type 1	Normal outages of a system or a component of a system, due to factors such as a hardware or software failure, or bad change management	24 Hours or Less
Type 2	Force Majeure that impacts a building or a campus; i.e., fires, loss of electricity, sabotage	24 hours to 7 Days
Type 3	Force Majeure that impacts a region; i.e., widespread power outage, floods, hurricane, terrorism	8 Days or Greater

Clearly the contents of Table 1 need to be customized for the somewhat unique situation facing each organization. However, the value of using a tool such as Table 1 is that it begins to appropriately structure the conversation between the IT organization and business decision makers. However, for this conversation between the IT organization and the business decision makers to truly be meaningful, the cost of the various options needs to be included in this conversation. The topic of the funding of an IT contingency plan will be discussed in Section 10 of this document.

6. High Availability Network Design

The purpose of this section of the document is to present some general guidelines relative to the design of highly available networks. Section 7 will then detail some particular techniques and technologies that can be used to further enhance network availability.

It should be noted that designing a cost-effective, highly available network requires a mix of technologies. Device-level redundancy can eliminate single points of failure (SPOFs) within devices but cannot protect a distributed network from link failures or “soft” failures, such as operator error, mis-configuration, or software errors. The best approach to protecting a distributed network from failures of all types is through a resilient network design that provides multiple parallel paths between key end systems, thus eliminating SPOFs at the network level.

A failure event in a resilient network typically results in traffic failing-over to the surviving path(s), temporarily increasing the load on remaining devices and links. Device-level redundancy and link redundancy can be exploited to limit the frequency of these fail-overs. When fail-overs do occur, QoS mechanisms can maintain a high level of service by automatically allocating a higher percentage of network resources to critical application traffic.

It should also be noted that designing a cost-effective, highly available network typically increases the complexity of that network. Avaya’s Minassian advises network professionals to “balance the complexity of your network design with maintainability. You can engineer some very good complex self-healing networks. Just make sure that you have the ongoing talent to support and troubleshoot them in times of crisis.”

6.1 Robust LAN Design

As of a few years ago, it became common to implement large campus switched LANs in a hierarchical three-tiered design. This design corresponds to the physical topology of wiring closet, site backbone and campus backbone, as shown in Figure 2.

Each tier of the network is comprised of a number of replicated, fault tolerant modules that are optimized for the required functions in that tier. Note that a small enterprise network might be implemented as a single-tier LAN corresponding to one or more wiring closet modules, plus the addition of a WAN router. A medium-sized enterprise network might be implemented as the combination of the wiring closet and site backbone modules.

The most common switched LAN topology has been based on Layer 2 switching in the wiring closet tier, with Layer 3 switching in the site and campus tiers. This approach, sometimes called an L2/L3/L3 approach, has the advantage of preserving the subnet structure from the earlier hub/router model of LAN design. While


the campus tier can also be successfully implemented with Layer 2 switching, Layer 3 campus backbones are considerably more scaleable and robust. Layer 3 switching allows fast fail-over times, which directly translates into higher availability. Layer 3 campus backbones also provide for full link utilization and load-sharing over parallel, equal cost paths through a meshed Open Shortest Path First (OSPF) backbone.

One drawback of the L2/L3/L3 model is that the uplinks from the closet level are switched at Layer 2. In particular, the use of the spanning tree algorithm can prevent configuration of parallel load-sharing connections from each wiring closet switch to redundant Layer 3 site switches. Not being able to configure these parallel connections reduces the availability of this particular LAN design.

An effective way to increase availability without violating the spanning tree algorithm is to use Link Aggregation (LAG) to create point-to-point links between the wiring closet switch and first Layer 3 switch at the boundary of the core as shown in Figure 2. The spanning tree algorithm views each LAG as a single logical link, and within the LAG, load is balanced among the members based on MAC address and/or port denomination.

One increasingly common modification of the L2/L3/L3 model is to extend the Layer 3 boundary to include the uplink ports on the wiring closet switch, which allows OSPF Equal Cost Multi-Path (ECMP) load sharing in the risers as well as in the backbone. With ECMP in the risers, it is possible to design an end-to-end redundant network without the cost/performance overhead of stand-by secondary uplinks that carry traffic only when primary paths are interrupted by failures.

Figure 2 can also be used to show how some network-level features can be employed in a three-tier, resilient campus network design with full link and device redundancy from the wiring



closet to the server farm. With Layer 2 switching in wiring closet switch “A”, link and router redundancy are provided with LAG and a protocol such as Virtual Router Redundancy Protocol (VRRP). If the primary LAG or primary site Layer 3 switch “C” fails, traffic is re-directed to the secondary site switch “D” via the standby LAG. Alternatively, OSPF Layer 3 switching on the uplinks of wiring closet switch “B” allows load balancing across the equal cost paths to switches “C” and “D”.

In the site and backbone tiers, Equal Cost Multi-Path OSPF supports full link and device redundancy. In the server farm, OpenTrunk 802.1Q VLAN trunking between the L2/L3 switches “E” and “F” enables the server to be dual homed on a VLAN that spans the two server switches.

6.2 Robust WAN Design

To implement network design redundancy in the WAN, all major network sites should consider having dual points of access, with dual routers providing parallel, load-balanced connections to the WAN. Note that this will typically involve working with the Local Exchange Carrier (LEC) to determine if truly diverse access is available from these sites. Also note that where this is available, it will be expensive.

Another technique, commonly implemented in smaller locations, is to use ISDN as a back-up media. Typically, the ISDN circuit is automatically activated when the primary WAN service (i.e., private line or Frame Relay) is not working. It is common to use two B channels, and to use both inverse multiplexing and compression. The result is a circuit that can be expected to deliver on the order of 256 Kbps. Note, however, that this technique will be of no use if the cause of the outage is a cut in the LEC access circuits. If that is an issue and diversely routed access across terrestrial circuits is not readily or economically available, alternative strategies, such as microwave access to the service provider, or the use of satellite links, should be evaluated.

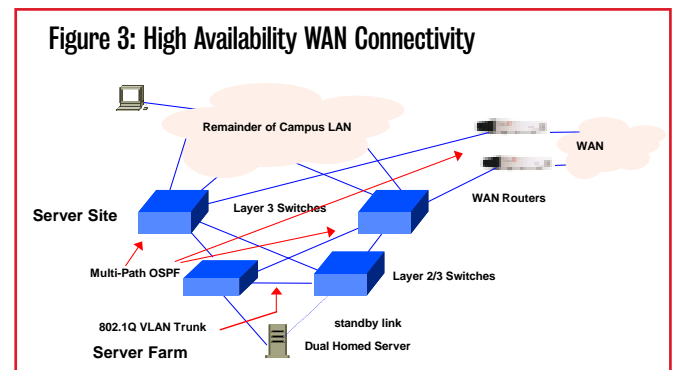
It is common to get service-level or up-time guarantees from service providers. In theory, this moves the issue of providing high availability from being a design issue to one of negotiation and network management. It is an issue of negotiation, because while it is common for service providers to supply service level guarantees, in most cases the penalties for not performing are relatively trivial. Hence, network executives looking to deploy a highly available WAN need to be able to negotiate more meaningful service level guarantees.

While negotiating more meaningful service level guarantees sounds like an important component of a contingency plan, Avaya’s Minassian advises: “Don’t bank on an SLA (Service Level Agreement) helping you in times of an outage. Avaya had an experience where the last leg of a circuit was sublet to a wireless provider in the Netherlands. There was a network outage, but

labor laws prevented the technician from being dispatched for 12 hours. Because of these legal limitations, the SLA provided no support from a Disaster Recovery perspective.” Minassian advises that if the WAN is critical to your business, you need to have route diversity and redundancy.

Providing high availability is also an issue of network management because the monitoring of the service level guarantees requires the implementation of relatively sophisticated service management tools and processes. Where organizations cannot afford to rely purely on the SLA offered by service providers, the dual network connections can be made to redundant network services from competing WAN and Internet service providers. One example of this is that many corporations split their call center (i.e., 800) traffic over multiple service providers.

Figure 3 depicts one WAN access architecture that is highly available. Note that the dual access routers are connected to redundant Layer 3 switches in the server site tier to eliminate any SPoFs between the campus LAN and the WAN.



ISERV’s Corkery says that “in planning for a high availability network it is highly desirable to ensure that you have route diversity. Depending upon the areas that you serve, however, some additional investigation may be required. We have had several situations where, though using different carriers, the physical transport was the same. This was only detected when the backhoe highlighted this single point of failure.”

Minassian agrees with Corkery: “Avaya has had experience with dual circuits from different carriers traversing the same fiber backbone. This was especially apparent in transatlantic and transpacific links where many providers sublease capacity.” Minassian recommends that network professionals “demand



detailed route paths from your providers and know what they carry themselves and what they sublease.”

7. Technologies that Enable Contingency Solutions

This section of the document will briefly describe a number of technologies and explain how they can be a part of the solution to a company's contingency requirements. No attempt will be made in either this or the next two sections to create anything close to an exhaustive list of relevant technologies, techniques and processes. The goal of these sections is merely to exemplify the vast array of technologies, techniques and processes that can enable contingency planning.

7.1 Load Balancing for Servers and Contact Centers

Many companies have implemented load balancing across servers and/or contact centers for the obvious operational efficiencies. However, this technique also facilitates contingency planning. This follows because the use of load balancing means that a company is less vulnerable to a single outage, be that the outage of a server or of an entire contact center.

However, Corkery has some concerns about this approach. “While load balancing is a desirable function to offset excessive hardware and operational costs, our experience with this technique is mixed. In a simple transaction format it works well and achieves its business objectives. But when the transactions are complex — specifically secure multithreaded transactions — we have had difficulties in implementing this capability. The difficulty was largely because of the security issues surrounding the use of PKI. Though it turned out not to be a network issue, for several months the network folk bore the brunt of the blame.”

7.2 Wireless Solutions

Wireless solutions such as wireless LANs, cellular phones, and 3G mobile offer the potential to provide backup for those situations in which a wireline solution has failed due to a catastrophe. However, the experience of one company, in the agriculture industry, points to some of the potential limitations of such an approach. The firm had, as part of its contingency plan, the backup of wireline voice services with cellular voice services. What the company found out during a flood, however, was that for the first few days after the flooding started their local cellular service provider could not come close to supporting the extra volume of calling. As such, their contingency plans for providing voice services was largely ineffective.

7.3 Expert Systems

The focus of expert systems is to minimize or eliminate Type 1 incidents; i.e., normal outages of a system or a component of a system, due to factors such as a hardware or software failure, or

bad change management. The idea is that the expert system would be able to identify faults before they impact service. While it is difficult to quantify the value of these types of systems, some service providers are incorporating this type of approach into their service description. In particular, by use of expert systems, some service providers are writing Service Level Agreements that promise to identify and eliminate 90% of all network faults before they impact service.

7.4 MPLS Enabled Frame Relay WAN Services

A number of service providers are enhancing their existing Frame Relay services by incorporating Multi-Protocol Label Switching into their core network. Since IP routing is performed in the MPLS core, each site on the Virtual Private Network (VPN) can now communicate “directly” to each other site on the VPN. This means that through one enhanced Permanent Virtual Circuit (PVC) the customer gets any-to-any connectivity. This any-to-any connectivity can be used as part of a contingency plan because it enables a site to easily reach other company sites, as well as to reach stand-by sites.

It could be argued that there is nothing that a company can do with an MPLS Enabled Frame Relay service that it could not do with a more traditional Frame Relay service. For example, in a traditional frame relay network it is somewhat common to deploy multiple PVCs originating at each site. With this type of a design, a company can provide for connectivity from that site to other production sites, as well as to hot and/or cold stand-by sites.

However, an MPLS Enabled Frame Relay service has a distinct advantage over a Frame Relay service relative to the cost of providing a high degree of connectivity among a company's sites. And, as discussed in Section 6, cost is a key factor in the deployment of any contingency plan.

The cost advantage of an MPLS Enabled Frame Relay service in a highly meshed network comes from how MPLS Enabled Frame Relay is priced. An MPLS Enabled Frame Relay service typically has the same port costs and the same access link costs as a traditional frame relay service. The only difference in the pricing is the cost of the MPLS Enabled Frame Relay PVCs. While pricing information is somewhat sketchy, for planning purposes it is reasonable to assume that a PVC in an MPLS Enabled Frame Relay service costs between one and two times the cost of a PVC in a traditional Frame Relay service.

7.5 Content Distribution Networks

In the traditional computing model, all of the data for a given application typically resides in a single data center. This is clearly a single point of vulnerability, and is the motivation for why at least some companies implement the techniques outlined in Section 4 of this document.

The primary motivation for deploying a content distribution network has little to do with contingency planning. The motivation is to store content closer to the users and by doing this, to both reduce the cost of the WAN and to improve the user experience.

However, once there are multiple copies of content distributed around a company's WAN, that company has eliminated its single point of vulnerability for this content. Looked at this way, the deployment of a content distribution network directly enables contingency planning if there are tight linkages between the ongoing management of the content distribution network and the contingency plan. In particular, in order for a content distribution network to play a significant role in a company's contingency plan, somebody on an ongoing basis needs to identify exactly what content is distributed throughout the company's network. Then, plans need to be developed for backing up other critical content, as well as ensuring that users have access to all of this content should a catastrophe occur.

7.6 Back-Up Power

When most network professionals think about installing back-up power, they typically are thinking about installing it inside of a data center. However, Minassian believes that that thinking is somewhat flawed. "Most data centers have back-up power, but unless each piece of equipment that resides between the data center and the critical end users are also on back-up power, you do not really have a contingency plan." Similar to the advice provided by Endry and Corkery, Minassian believes that "having a data center running, but not having users being able to access the data center, is of no value." Network professionals, he adds, need to identify all their company's critical users, or user locations, and make sure they also have the necessary power."

7.7 Re-direction of Phone Calls

An earlier section of this document commented on the tactic of deploying hot stand-by office facilities. For this to be a viable component of a contingency plan, it is necessary to be able to easily take phone calls that were destined for the primary office site and re-direct them to the stand-by office site.

This ability to re-direct phone calls can be done either as part of a company's private network, or by utilizing a service provider. In particular, some service providers are offering call-forwarding services that are intended to be part of a contingency plan. These services promise to provide to an IT organization the ability to quickly reroute incoming voice calls to an alternative site, or to individual alternate telephone numbers.

Kevin Lopez, Director of Telecommunications at professional services firm Grant Thornton, says that "Re-direction of phone calls offers flexibility in call routing in cases of disaster recovery or telecommuting associates. Initiating this feature can be done

in a variety of ways, among which are user programmable and administrator programmable. Using transfer supervision, calls are re-directed off the network but still follow their coverage path back into the user's company-provided voicemail box. This allows companies to maintain the integrity of client communications on a company's network and also monitor and decrease instances of toll fraud."

7.8 VoIP

A number of companies are deploying Voice over IP. The typical reason that a company deploys VoIP is some combination of cost savings and the ability to more easily deploy new integrated applications or offer sophisticated voice functionality in offices where that was not previously economically feasible.

However, some companies are deploying VoIP in part because of how it enables contingency planning. For example, Corkery says that "Voice over IP, or IP Telephony, is becoming increasingly common. We are looking to replace our Centrex services because there is a business case to do so — not only from a financial perspective — but also it will result in improved contingency plans and increased security. Instead of having to deal with multiple networks, a robust, single IP network can simplify the planning and operational issues."

One component of many VoIP implementations is the use of a Softphone. A Softphone is a software package that emulates the look, feel, and function of a phone on a PC. Lopez notes that "IP Softphones are excellent tools for the telecommuter, personnel working from home, or in the case of disaster recovery. The IP Softphone allows the utilization of one phone line for voice and data connectivity, while also allowing the user to control incoming and outgoing calls. Users can utilize all PBX features (i.e., transfer, hold, conference, and auto-dial) and have access to inter-company dialing while being offsite."

7.9 SRDF

SRDF (Symmetrix Remote Data Facility) provides for the synchronization between symmetrix devices. EMC's Dymek uses this functionality to implement the type of hot stand-by data center that was described in Section 4. Dymek says he has two data centers that are connected over a SONET ring that runs at 600 Mbps. Any data that is written to one storage device is automatically written to the other. "If a disaster were to strike one data center, the other will take over operations, and we will never miss a beat," he says.

8. Other Techniques that Enable Contingency Solutions

This section of the document will briefly describe a number of general techniques and explain how they can be a part of the solution to a company's contingency requirements.

8.1 Network Simplification

Many IT infrastructures have grown out of control in terms of the number of different objects in the network. In this context, the term “object” refers to things such as:

- A type of hardware, such as a router or a PC;
- A version of software, such as a given application, or the operating system running a mainframe, a server, or a PC;
- A type of LAN;
- A particular WAN transmission technology.

A number of IT organizations have begun efforts to reduce the number of objects in their infrastructure. For the most part, these efforts are motivated by the desire to reduce cost. For example, one CTO at a large financial services firm explains that when he benchmarked the cost of his company’s IT infrastructure vs. similar companies, his costs were running roughly 30% below similar companies. He attributed this difference to the fact that they have standardized most of the IT infrastructure.

However, simplification of the IT infrastructure also facilitates implementing a continuity solution. This follows because it is notably easier to get an IT infrastructure back up and running after a catastrophe if there are relatively few options that have to be accommodated.

8.2 Geographical Diversity

If all of a company’s resources are at a single site, then that company is very vulnerable to a catastrophe that affects that site. Because of this, a number of organizations are reassessing their real estate policies. In particular, a number of companies are now implementing policies that limit how many employees can work in the same building.

8.3 Hosted Solutions

A number of businesses offer services in which they host applications for companies. One of the advantages of using a hosted service is reliability. In particular, these services are typically offered from a data center that is designed to survive a number of disasters. In addition, the data center typically has access to multiple, diverse fiber routes.

Note that there is nothing that these hosting vendors provide for customers that these customers could not provide for themselves. However, since these vendors are providing these services for a vast array of customers, they tend to have an economy of scale. Based on this economy of scale, these hosted services tend to be more cost effective than similar services that the customer could build for themselves, and hence are more likely to get management approval.

9. Processes that Enable Contingency Solutions

The previous sections of this document addressed some technologies and some general techniques that can be part of an effective solution to a company’s contingency requirements. This section will detail some of the processes that need to be part of that solution. The cost of these processes needs to be included in the business case that is developed to evaluate the financial viability of deploying some form of contingency plan.

9.1 Link to Other Key Processes

In order for an IT contingency plan to be impactful, it must be linked to other key processes. For example, as was emphasized by both Endry and Corkery, if IT is working on a contingency plan that does not have tight linkages to corresponding businesses processes to recover from a catastrophe, then IT is wasting both time and money.

In addition, an IT organization’s contingency plan needs to be linked to its security plan. This follows for two reasons; First, a security incident such as a Denial of Service attack has the effect of keeping a company from fully functioning; Second, there is considerable overlap between what has to be done to create an effective security plan and what has to be done to create an effective contingency plan. Linking these two plans will tend to both minimize cost and maximize the effectiveness of these plans.

9.2 Perform a Regular Inventory Analysis

Minassian says that in order to deploy a robust contingency plan, network professionals need to “know what you have installed — everything from the overall design to the latest revision of firmware on each device. Too many network teams have the design on a Visio diagram sitting on an engineer’s PC.”

9.3 Have Arrangements for Back-Up Equipment

The CTO at one New York City financial firm says that they were able to have 3,000 affected employees fully functioning when the stock market reopened after September 11. He attributed their ability to do this in part to the fact that they have a standard desktop environment. Given this standard environment, it was relatively easy for them to load the appropriate software onto the backup desktops. The most difficult part of this operation was getting 3,000 new PCs from their suppliers.



9.4 Create an IT Crisis Management Team

In order to quickly respond to a catastrophe, an IT crisis management team, and a set of alternates, must be in place prior to the catastrophe occurring. The team needs to have a well-defined set of duties. A key part of ensuring that this team is effective in a short period of time is establishing and maintaining up-to-date contact information for everybody on the crisis management team, as well as for their alternates.

9.5 Identify and Document Methods of Efficient Communications

In order for the crisis management team to be effective, they need efficient communications. Corkery notes that “Critical to managing in a crisis is excellent communications. The use of new technologies such as Research-in-Motion’s Blackberry family of devices has proven to be highly effective. Last spring, a fire in a street power vault caused our building to close. We were advised of the problem as most of the senior management team were in the middle of their morning commute. By the time we got there, arrangements had been made for activating alternate office space and rerouting the help-desk calls. To the business user it was transparent but it could never have been done without the Blackberry text messaging.”

9.6 Establish a succession plan

In order for an IT organization to flourish under normal conditions, they need to do succession planning for their human resources. The need for this kind of planning is magnified when doing contingency planning. That follows because a company can well have a somewhat hidden single point of vulnerability if all of the knowledge of some key aspect of the network, such as the IP addressing scheme, resides largely inside of the minds of one or two individuals.

Minassian recommends that network professionals “make sure that you have more than one person who knows your network. Too many times people rely on one employee, or even worse rely on a single contractor, only to have the person leave with all the knowledge. Get a buddy system or formal succession plans in place immediately.”

9.7 Educate the Employees

In order for any contingency plan to be effective, there must be a high level of awareness of the plan on the part of the company’s employees. Formal training, both at new employee orientation and on an ongoing basis, can create this awareness. There are a variety of other techniques that can also be used. For example, a number of firms in New York City have created emergency information cards. These are typically laminated cards that fit inside of a shirt pocket that detail what actions should be taken when and if a catastrophe occurs.

9.8 Testing the Plan

Corkery says that “The development of a disaster recovery or continuity plan is immaterial unless it is tested and updated. We test ours four times a year — three scheduled and one ‘surprise.’”

Corkery notes that in a prior role, during the surprise test, they used to leave envelopes on people’s desks telling them they or a co-worker had been disabled in the fire/bomb blast/whatever. This technique caused the test to be a truer simulation of what would really happen in an emergency situation. Kirk also emphasizes the need to keep the call out numbers up to date — in his experience he has “found that by normal movement 20% are likely to be wrong after 90 days.”

PricewaterhouseCoopers’ Brown says that “Your test plan is one of the most critical elements of any disaster recovery plan. This plan needs to cover everything from standby hardware availability to does the diesel generator have fuel and function correctly.” Brown notes that when he was in a previous position, “We found that someone had rewired a piece of the building’s physical plant and as a result, the diesel would not start. This was discovered in a routine quarterly test, which clearly demonstrates the value of period testing.”

Minassian recommends that network professionals also regularly check their out of band management access. “In many instances you have technicians disconnect links for the use of a modem and then never reconnect them. You only find this out when it’s too late and you need access to respond to an emergency.”

9.9 Recovery

As was mentioned previously, in 1984 iSERV’s Corkery was involved in the recovery of a major data center for a large telco following a fire. As a result of that experience, Corkery learned that a contingency plan needs to include what happens once the disaster is over. “The second major lesson that I learned as a result of that fire was that although we had planned on how to get to the alternate data center very well, and our plans had been executed almost flawlessly, we did not have a plan on how to migrate back. It would literally have been easier for us to declare a second disaster to move everything back but instead took an additional month to migrate it all slowly and carefully back to whence it had come.”

10. The Funding of Contingency Planning

Given the plethora of technologies that can be used to implement a contingency plan, it is reasonable to assert that the key issue in designing a contingency plan is not the lack of technology, but the lack of funding. This section of the document will dis-

Discuss some of the issues relative to getting funding to design and implement a contingency plan.

10.1 IT Governance

As was mentioned in Section 5, one of the key sets of questions that needs to be answered relative to the development of a contingency plan is:

- Who is the decision maker for the contingency plan?
- What is their willingness to pay for the identified levels of protection?

There are many models used by companies to fund Information Technology. However, it is common across most IT organizations that funding typically occurs only once a year. If something develops during the year, such as the need to deploy a contingency plan, there are a couple of options that can be used to pay for the new development, neither one of which is very desirous. One option is to cut activities that had been in the budget in order to fund the new activity. The second option is to go back and seek additional funding. While this second option may work in some companies, it will not work in all companies. One CIO described the viability of the second option in his organization as “You never go back to the well – never.” Given the issues surrounding both of these funding options, network executives need to build contingency planning into their yearly budget cycle.

How IT is funded has a significant impact on what steps network professionals need to take in order to get funding for either developing or implementing a contingency plan. This section will describe one common funding model and discuss the impact of this model on the implementation of a contingency plan. For the purposes of this document, this funding model will be referred to as The IT Funding Model.

The IT Funding Model calls for the network organization to provide a core set of functionality to most, if not all, employees of the company. This core set of functionality could include services such as dial tone, 4 digit dialing, voice mail, internet access, e-mail, and remote access. The funding for this core set of services comes from the Sales, General, and Administrative (SG&A) budget line. The appropriate level of funding for these core services is determined in negotiations between the CIO and the IT steering committee. The IT organization recoups these costs through a combination of usage-sensitive chargeback mechanisms, for features such as dial tone, and some form of allocations for services such as Internet access.

The IT Funding Model also allows for the network organization to provide services that are only used by a small subset of the company's employees. In this case, the negotiation for the funding to support these services takes place between executives of

the network organization and the appropriate Business Unit Managers.

Given The IT Funding Model, getting the funding to develop and implement a contingency plan would have the following primary steps:

- If the proposed contingency plan were focused on just one or two business units, then those business units would be funding the development and implementation of the plan. As such, the network organization would need to negotiate directly with these business units.
- If the proposed contingency plan supported most, if not all, of the company's organizational units, then funding the development and implementation of the plan would have to be approved by the IT steering committee. As such, the network organization would need to negotiate with all of the organizations that make up the IT steering committee.

J.D. Edwards' Endry says that when managing the budgeting process, “Network managers need to make sure that business cases proposing the implementation of new applications or services include the incremental expenses associated with changes that might be required in the contingency plan. The best way to do this is to include the costs up front in the business case. If they are shown as a separate line item, people involved in approving the business case will be tempted to reduce the cost of the business case via line item veto. If the new application or service truly impacts the contingency plan the expenses must be approved or the entire business case needs to be revised.”

10.2 Business Case Analysis

The conventional wisdom in the IT industry is that 40% of companies that experience a disaster will go out of business within a few years of the disaster. While this statistic may be compelling enough to induce an organization to evaluate deploying a contingency plan, it typically will not be compelling enough to convince management to spend a significant amount of money to implement a contingency plan. To accomplish that goal, network executives need to work with the decision makers within their company to do a business case analysis that is creditable in their environment. That business case analysis begins with an identification of the cost of down time.

It is worth mentioning that for the majority of business functions there is not a general recognition that if the function was down for a few hours, or perhaps even a day, that there would be a clear cost to the business. This does not say that people would not be extremely upset if the function was inoperative. What it does say is that there would not be a willingness to pay any additional amount to reduce the chance of an outage.

However, in many, if not most, companies, there are at least some business functions for which it is recognized that there is a cost of downtime. Examples of such functions include:

- Brokerage trading;
- Online sales;
- Just-in-Time Manufacturing;
- Online reservations.

For functions such as these, there are a number of factors that go into determining the cost of downtime. These factors include a loss of:

- Revenues;
- Customers;
- Creditability.

In order to do a business case for deploying a contingency plan, the IT organization should do an analysis to create a table similar to Table 2. The table is created for either a few key sites, a few key functions, or possibly for the entire operation. One viable approach is to start with the business situations that are the most likely to have a recognizable cost of downtime within a company. If business management is not willing to fund a contingency plan for these situations, it is likely to be a waste of time to expand the analysis to other business situations.

For the sake of example, assume that the focus of Table 2 is the trading floor of a major brokerage firm. To complete the example, assume that the specific issue being analyzed is the economic viability of deploying a contingency plan that would enable the firm's traders to continue to have access to market data in spite of a catastrophe.

The purpose of creating Table 2 is to position the IT organization to be able to present alternative solutions to the person or people who will ultimately decide if it is worthwhile to fund a contingency plan. In this context, a solution is a combination of the types of technologies, techniques, and processes outlined in the preceding sections of this document.

In order to construct Table 2, an IT organization has to evaluate a number of different potential solutions based on two key factors. Note that as part of evaluating the solutions, the IT organization also develops an estimate of the cost of implementing the solution. The cost of these potential solutions is a critical component of Table 2.

One of the factors that goes into the creation of Table 2 is the amount of protection that will be provided as part of the contingency plan. In terms of this particular example, two options were explored. In one option, the traders never lose access to the

trading data. In the second option, the traders could lose access to trading data for up to four hours.

The second factor that goes into the creation of Table 2 is the severity of the catastrophe for which the plan is intended to provide contingency against. Note that a methodology for categorizing types of catastrophes was detailed in Table 1. In this particular example, all three types of catastrophes are considered.

Table 2: Comparison of Solutions

Solutions	Amount of Protection	Type of Catastrophe	Cost (in millions)
1	Back up 4 hours	Type #1	\$1M
2	Back up 4 hours	Type #2	\$5M
3	Back up 4 hours	Type #3	\$25M
4	Hot stand-by	Type #1	\$5M
5	Hot stand-by	Type #2	\$25M
6	Hot stand-by	Type #3	\$125M

The use of a tool such as Table 2 accomplishes the following:

- It better positions the IT organization as a partner to the company's business and functional managers. This follows because the IT organization is presenting alternatives for the business decision makers to consider.
- It introduces the element of cost into the decision-making. To paraphrase Endry, without introducing the element of cost, it may be assumed at first that a requirement of a business continuity plan is to duplicate 100% of the company's e-mail. After quantifying the cost of duplicating 100% of a company's e-mail, it may well be determined that it is acceptable, from a business perspective, to only duplicate 50% of the company's e-mail.



11. Summary

In part due to the events of September 11, Disaster Recovery and Business Continuity have become hot topics. While the effects of the attack on the U.S. might fade in the minds of some people, Disaster Recovery and Business Continuity will remain important to a wide range of companies. This follows for two reasons. One, for many companies, the effects of the attacks of September 11 are not going to fade away. Second, Disaster Recovery and Business Continuity are concerned with more than just protecting against events of the magnitude of September 11.

Table 3: Classes of Catastrophes

Type of Catastrophe	Characterization of Catastrophe	Expected Length of Outage
Type 1	Normal outages of a system or a component of a system, due to factors such as a hardware or software failure, or bad change management	24 Hours or Less
Type 2	Force Majeure that impacts a building or a campus; i.e., fires, loss of electricity, sabotage	24 hours to 7 Days
Type 3	Force Majeure that impacts a region; i.e., wide spread power outage, floods, hurricane, terrorism	8 Days or Greater

Given the lasting importance of this topic, this document developed a framework that can be used by network executives to create an IT contingency plan that reflects their company's particular situation and needs. For example, some companies will want to take some small steps in order to get started with contingency planning, while others will want to role out a more comprehensive contingency plan. In either case, in order to be effective, there must be tight linkages between an IT contingency plan and two other activities: the company's security planning, and the business processes that are being implemented that correspond to the IT contingency plan.

The framework that was developed in this document suggests that when developing a contingency plan, network professionals look at the three classes of catastrophes that are shown in Table 3.

The document presents some guidelines for how to design highly available LAN and WAN infrastructures. However, as with virtually any aspect of contingency planning, the implementation of these guidelines will most likely increase the cost of networking. The implementation of these guidelines can also result in a complex network infrastructure.

The framework presents the solution to a company's contingency requirements as a combination of technologies, general techniques, and processes. As the document points out, there is a wide array of technologies that can be used to create these solutions. Some of these technologies, such as back-up power and the re-direction of phone calls, are deployed primarily to enable contingency planning. However, a number of other technologies, such as VoIP, and Content Distribution Networks, are typically deployed for reasons that have nothing to do with contingency planning, even though they can be used for that purpose.

Some techniques that many companies are implementing in order to reduce their vulnerability to a catastrophe are:

- Simplification of the IT infrastructure;
- Diversification a company's real estate holdings;
- Utilization of hosted solutions.

There are also a wide array of processes that can be used in part to create a contingency solution. Some of these processes, such as performing a regular inventory analysis, can be classified as just good IT management. However, the focus of the majority of the processes is just to support a contingency plan. As such, implementing these processes represents incremental work and cost.

A number of senior IT executives who provided input for this document were in agreement that it is not possible to have a meaningful discussion about contingency planning without also discussing the relevant funding issues. This document suggests that the IT organization that is just getting started with contingency planning should begin with business situations that are the most likely to have a recognizable cost of downtime within their company. If business management within their company is not willing to fund a contingency plan for these situations, it is likely to be a waste of time to expand the analysis to other business situations.

Business Continuity at Coca-Cola Enterprises

According to John Ridley, Senior Enterprise Network Architect at Coca-Cola Enterprises (CCE), CCE is currently developing a comprehensive Disaster Recovery/Business Continuity (DR/BC) plan. The goals of the plan are to provide:

- 7/24/365 availability for their core ERP applications;
- Seven to fourteen day recovery in the event of a catastrophic event;

- Technology, processes and procedures which produce the longest return on capital investments and the lowest practical operating costs to the business.

The CCE DR/BC plan is intended to protect against all levels of application availability outage. The emphasis of the plan is weighted in proportion to CCE's estimation of the probability of occurrence of an event which would impact availability of their core applications. As a result, much more emphasis has been placed on the RAS (resiliency, availability, scalability) principles of design than with the consequences of acts of war or other cataclysmic events. However the goal is that if a meteor strikes, Coca-Cola will be available to be served to the masses who are assisting with the recovery effort.

CCE's Corporate IT Department has worked with industry experts and many vendors over the last year to formulate various options, complete with cost estimates and risk assessments. The most important guiding principle throughout the process has been to not allow fear, uncertainty, and doubt (the FUD factor), or vendor or technology religions, to outweigh logical and factual reasoning.

The principle technology building blocks have included the use of:

- MPLS as a WAN technology with geographically diverse POP and Switch entrance points to the cloud;
- A single state-of-the-art level nine central data center;
- A diverse routed SONET underground ring;
- Dual carrier class WAN routers supporting multiple physical and logical connections to the WAN with BGP-enabled load balancing;
- Dual non-oversubscribed, high-speed ASIC-based aggregating switches with multiple OSPF enabled load-balanced connections to the WAN routers and the down stream redundant switch clusters;
- Dual and redundant switch clusters supporting non-revertive fail-over of multiple host resident NICs;
- Strategically placed optical taps between the three layers


for non-intrusive monitoring, intrusion detection, and diagnostic purposes;

- Aggregating switch-layer, ASIC-based, load-balancing appliances;
- Redundant network appliances such as Stratum 1 GPS based time servers and DNS/DHCP/WINS servers;
- State of the art SAN technology;
- Creation of an integrated systems and network monitoring and troubleshooting group.

Business Continuity at EMC

The most important consideration in business continuity planning is identifying how current the data used by a particular application is after an outage. EMC Corporation has adopted a methodology whereby any component of EMC's business that ultimately impacts its customer base must have immediate restore capabilities in the event of an outage — ensuring virtually zero data loss. Such a methodology ensures the highest levels of information protection for most of EMC's business, including: product development, sales operations, manufacturing, and customer support.

At its headquarters in Hopkinton, Massachusetts, EMC utilizes two data centers and installs a mirror image of the networks, servers, and storage required for all associated applications. Production operations are run from the environment in one data



center with real-time updates sent to the remote data center. This remote data center is also populated with IT personnel and is used to perform backups, test new applications, and perform data warehouse updates. By using this type of data replication, the two environments are always in sync. When an outage occurs, production can be “flipped” to the other data center with virtually zero data loss. Both centers in Massachusetts are also connected to EMC data centers in Europe and North Carolina, providing an extra layer of protection.

Another important aspect of business continuity planning is people. One has to consider where the human assets are located and, if forced out of their usual placement, where they can go to continue conducting their business operations. Having multiple facilities is good for the general population, but there are specific groups of people that require specialized services or equipment in order to conduct business. The needs of these identified groups must be considered when establishing a plan. For example, EMC's Customer Service centers require some unique voice and data services that the rest of the company does not. They cannot pick up and relocate to just any facility in the event that a disaster occurs. As a result, multiple call centers have been established around the world, allowing work to be shifted if one center is no longer able to function. This is accomplished through the rerouting of calls and customer cases to the next available center.

It is through a careful assessment of your company's operation that a viable business continuity plan is developed. Each business process should be examined in detail to determine its sensitivity to the loss of current data and the ability of people to operate from various locations. Once that analysis has been performed, the appropriate business continuity solution can then be implemented.

