Title:	**Business Continuity Planning in the Aftermath of 9/11/2001**
Author:	Joseph T. Balsamello, M Moser Associates
Subject:	Economics/Financial
Keywords:	Economics Life Safety
Publication Date:	2003
Original Publication:	CIB / CTBUH Kuala Lumpur 2003 Conference
Paper Type:	1. Book chapter/Part chapter 2. Journal paper 3. **Conference proceeding** 4. Unpublished conference paper 5. Magazine article 6. Unpublished

BUSINESS CONTINUITY PLANNING IN THE AFTERMATH OF 9/11/2001

JOSEPH T. BALSAMELLO

M Moser Technology International – New York, NY USA

Abstract

The purpose of this white paper is to address the various ways to develop and implement improved business continuity planning in the financial services industry and general corporate sector. This paper will identify several vulnerabilities and identify a set of "sound practices" to address these concerns.

Keywords: business continuity, disaster recovery, disaster planning, business continuation

1. Introduction

The physical destruction, loss of life and the widespread dislocation of physical operations and personnel resulting from the terrorist attacks on the World Trade Center in New York City on September 11, 2001 were unprecedented. Despite this disaster, the U.S. financial system continued to perform its vital economic functions. As a result, there has been a significant increase in public interest ensuring that a large-scale disaster (whether natural or man-made) would cause no systemic disruptions to the financial markets and that companies/consumers maintain confidence in their ability to complete transactions and access to their funds.

The events of September 11 highlighted the vulnerabilities of the Financial Services Industry, especially in the clearing and settlement infrastructure. It was clear that decisions made by companies regarding their individual level of preparedness for disasters affected many others, both directly and indirectly. We now know that a coordinated approach to enhanced business continuity planning is vital to the overall health of the economy and the entire country.

The ultimate financial cost of 9/11 is still being tabulated and may never be completely known. The New York City Controller released a report on September 4th, 2002 that estimates the economic cost of the attack to New York City would be upwards of $95 billion. The final cost will ultimately depend on how many jobs are lost based on companies relocating out of New York City.

2. The Impact of the Events of September 11

During the week of September 11, the widespread destruction of physical infrastructure in and around the World Trade Center and extensive telecommunication breakdowns throughout the region caused mass confusion in the financial and corporate communities. U.S. equity markets were closed for four days and most bond trading, including government securities trading, halted for two days. There were significant disruptions in the clearing and settlement mechanisms for government securities, repurchase agreements and commercial paper. The core financial clearing systems continued to work well until operational failures and telecommunications breakdowns among major financial institutions led to significant liquidity bottlenecks for several days. During this period, the Federal Reserve and other major payment systems remained open well past their normal closing times to accommodate companies that were attempting to send funds or were waiting to receive funds.

Some institutions could not ascertain the financial positions of their company or of their customers for several days. Various regulatory requirements were relaxed in the face of credit and liquidity disruptions. Some transactions were lost and had to be reconstructed, a laborious and inexact process.

Vaults containing physical certificates were destroyed and records identifying these certificates were not always readily available. Reconciliation of disrupted transactions still continues now, over a year later. At the same time, extraordinary levels of cooperation by market participants in the aftermath of the September 11 events helped overcome limitations in the scope of firms' business continuity planning. Some firms were physically accommodated by other organizations (including competitors) in the New York City while obtaining office space for staff. A good example of this was Merrill Lynch which loaned office and trading space to JP Morgan immediately after the attacks. Companies extended credit to customers hampered by liquidity shortfalls despite the inherent uncertainty and lack of reliable information in the marketplace about their current financial condition. A large number of people inside and outside the financial industry worked long hours to restore communications links that had failed. In fact, the most oft-cited lesson learned from the tragedy is the importance of people, including considerations for their personal safety as well as their dedication and the critical role they play in keeping institutions functioning in times of crisis.

3. Major Vulnerabilities

The systemic effects of 9/11 highlighted several important vulnerabilities that may not have been widely appreciated previously. First, it was clear that business continuity planning had not taken into account the full potential for wide-area disasters and for the major loss of or inability to access critical staff. Contingency planning at many institutions generally focused on problems with a single building or system. Some firms arranged for their backup facilities to be in nearby buildings based on this assumption, for example, a fire that might incapacitate or destroy a single facility. Very few planned for an emergency disrupting an entire business district or city. As a result, some firms lost access to both their primary and backup facilities in the aftermath of the September 11 events, severely disrupting their operations. Institutions also generally had not considered the possibility that transportation of personnel could be significantly disrupted to preclude the relocation of staff to alternate sites.

Second, market-based and geographic concentrations intensified the impact of operational disruptions. Financial institutions are significantly concentrated within the geographic area of New York City most severely affected by the devastation at the World Trade Center. Over recent years some institutions have further consolidated their staff in one or two locations for efficiency and financial purposes. In addition, some critical market functions, particularly in the clearing/settlement of funds and securities rely on a small number of companies with operations in a concentrated geographic area. These market based and geographic concentrations further intensified the impact of operational disruptions.

Significant vulnerability in telecommunications capabilities resulting from similar concentrations became evident when telecommunications failures affected numerous institutions including primary and secondary backup sites in the same region. Many firms believed they had achieved redundancy in their communications systems by making arrangements with multiple telecommunications providers or by contracting for diverse routing. They quickly discovered during the disaster that this was not an effective redundancy plan because all of the lines traveled through any of several now well-known single points of failure.

Cantor Fitzgerald (CF) lost over 700 employees and its largest data center located at the North Tower. Despite this loss, they were up and running in less than 47 hours after the disaster. They were able to do this because CF had a disaster recovery plan to follow, which, ironically was put into effect only a few months prior to the attacks. ESpeed, a freestanding business unit and wholly owned subsidiary of CF, built the trading/clearing systems on a architecture that replicated all machines, connections and functionality at the World Trade Center, Rochelle Park and with a third facility in London. Gardner Group, a research think tank, estimates that two out of five companies that experience a disaster go out of business within five years. Without their disaster recovery planning it is safe to say that CF would not be in business today.

Third, the events of September 11 graphically demonstrated the interdependence of financial system participants, regardless of location. Although organizations outside New York City were less directly affected, the impact of the disaster was widespread and experienced on a global scale. Most companies outside New York City lost connectivity to banks, broker-dealers or other organizations in lower Manhattan. This breakdown impeded the ability of companies not physically impacted by the disaster to

conduct business and determine whether transactions had been completed as expected. Some customers were affected by actions of institutions with which they did not even do business when funds or securities could not be delivered due to operational problems at third party institutions. American Express, for instance, had to shift all credit card clearing out of New York City to their Phoenix location causing sporadic outages on a global scale.

4. Business Continuity Models

The events of September 11 has lead to changes in the way that companies plan for emergencies as well as changes in their ongoing operations. It is helpful to evaluate the basic models for business continuity planning and how well they performed during the recent crisis. This will allow us to plan accordingly for the future by evaluating what worked, what needs improvement and what completely failed.

4.1. Traditional Active/Backup Model (Active/Backup)

The most common or traditional model of business continuity is based on an "active" operating site with a corresponding backup site for critical data processing and operations. This strategy relies on relocating staff from the active site to the backup site and maintaining backup copies of technology and data at the alternate site. This model has an inherent dependency on the staff at the active site and their ability to get to the backup site. An adequate "desktop" recovery strategy - one that contemplates the movement of, at minimum, core employees to fully functional backup office space - is also a critical element of this model. This approach tends to restrict geographic separation of primary and back-up facilities to limit relocation time. Common approaches for the backup of technology infrastructure and data processing also rely on keeping data, hardware and software current at the backup site while having resilient and diverse services (including telecommunications and electric power) at each site. While the traditional active/backup model has been considered cost-effective and practical for many purposes, this model presents an internal and industry-wide challenge to ensure that the combination of primary and backup sites across diverse counterparties are compatible, well understood by all relevant parties and have up-to-date technology and procedures.

In the traditional model, backup capabilities are generally assured through planning and testing. Even with regular testing, it is often difficult to determine the effectiveness of backup sites, staff and systems that are not routinely used for production purposes. During the week of September 11, many institutions found that disaster recovery plans of particular business lines were not always accessible or up-to-date which, caused delays in restoring operations. [1] Merely placing outdated equipment at the recovery location to "meet" minimum requirements is no longer acceptable. Many companies now realize the importance of proper planning and investment in Disaster Contingency planning.

Other vulnerabilities of the traditional model were evident during the week of September 11. Some firms send back up files to offsite storage facilities on a weekly or bi-weekly basis. After losing their primary offices they were forced to devote substantial resources to reconstruct records that had not yet been transferred to their backup facilities. This suggests that recovering critical real-time processing operations from backup tapes is generally not realistic for large institutions' or for critical high-volume processing activities. Most larger institutions now employ data "mirroring" or remote real-time transaction logging technologies through which transactions are transmitted immediately to a second (and in some cases third) site. However, even in some of these cases, problems such as out-of-date software, reduced systems capacity and inadequate telecommunications at the backup site often were not discovered until operations were in the process of being recovered causing a whole new set of problems.

Institutions using the active/backup model may also rely on the services of a third-party to provide the backup facilities. Many companies, particularly small to mid-size ones, have contracted with third-party disaster recovery vendors, such as Sungard, for backup space for staff and computers. During the days following September 11 some disaster recovery vendors found they were unable to accommodate all of their effected clients. As a result many institutions found themselves without the anticipated backup facilities that they had contracted and paid for. Furthermore, the small number of disaster recovery

vendor sites supporting a large number of major financial institutions across the country may represent yet another more widespread vulnerability.

4.2. Split Operations Model (Active/Active)

An emerging business model, which is beginning to be used by some firms with national or global operations, is to operate with two or more widely separated active sites for critical operations. This "active/active" model creates inherent redundancy for each site that are often located hundreds of miles apart. For international firms, routine workloads can be shared among sites in different countries. Each site has the capacity to absorb some or all of the work of another for an extended period of time. This strategy can provide nearly immediate resumption capacity depending on the systems used to support the operations and the operating capacity at each site. The "active/active" method addresses many of the key vulnerabilities in the active/backup model. It eliminates dependency on availability and relocation of staff at any single location, reduces likelihood of telecommunications single points of failure, supports maximum geographic separation and assures business continuity through actual use, rather than infrequent and less than complete testing.

A facility does not have to be struck by a catastrophic event, like 9/11 to experience a total failure. A newly commissioned 100,000 square foot Mission Critical Internet Data Center in Sao Paulo, Brazil lost complete network connectivity due to a simple misconfiguration of a core router that created cascading problems for the local network. This data center was the only one of its kind in Brazil and incorporated the latest in security technology including primary and secondary bomb blast areas, retina scanners with PIN ID's, trip wires and motion activated cameras. The site was designed with the criteria of six nines of reliability (less than 32 seconds of downtime per year) and no single points-of-failure. Additionally, equipment included a 3000,000 gallon external condenser tank, a 1.2 million gallon thermal storage tank, four 675 ton chillers, six 2.5 megawatt generators, eight hundred thousand gallons of diesel fuel storage, multiple 800 kilovolt amps UPS's, duel electrical feeds and multiple fiber entrances.

Luckily, this company incorporated an "active-active" back-up strategy with a "self healing" ring topology. The closest failover facility was Buenos Aires, Argentina, approximately 1,000 miles away. The Buenos Aires site immediately became the primary active processing point for Brazil since they had a full duplication of staff, bandwidth capacity and servers. Buenos Aires saw a traffic increase of approximately 40% and personal onsite were unaware they became the primary distribution point for most of Brazil. It is safe to say that the "active-active" model of disaster recovery, in this particular situation, was cost effective, prevented negative publicity, functioned as designed and was completely transparent to clients as well as employees.

The "active-active" approach is not, however, without limitations. There are significant costs associated with maintaining excess capacity at each site and increased operational complexity. Depending on the sophistication of the function involved, it also may not be practical to maintain appropriately trained staff at multiple remote sites. After the World Trade Center attacks some firms with offices in other U.S. cities redirected workloads there as a short-term solution. In some cases this arrangement worked well while in others the remote offices lacked sufficient personnel trained to perform these functions.

Even with the "active/active" model, current technological limitations also preclude wide separation of data centers using full real-time, synchronous data mirroring backup technologies. However, emerging technologies are in use by some institutions that permit much more distant replication of data at multiple sites, if some time lag between sites can be tolerated. As technology advances and other techniques become more robust, greater geographic diversification of technology operations may very well become practical and financially feasible for many firms.

4.3. Other Models

There are other business continuity models that can provide a high degree of resiliency by utilizing various aspects of the "active/active" and "active/backup" models. For example, some institutions employ a variation on the above models in which a backup site periodically functions as the primary site for some pre-defined period of time.

The September 11th events demonstrated that response to a disaster is enhanced when records are kept electronically, allowing even the most current records to be replicated and recovered at backup sites. In addition, various institutions have noted that by increasing automated processing, backup arrangements are more straightforward because they do not depend as much on large-scale staff relocations. Recent events included the grounding of all air transportation, bridge/tunnel closing and disruptions to mail delivery further prompting movement away from physical, paper-based transactions and recordkeeping systems in favor of electronic methods. Although electronic records help protect against loss of a physical site, dependence on electronic records increases vulnerability to cyber attacks and to defects in hardware and software.

5. Developing Sound Practices for Business Continuity

In the face of more realistic assessments of the types, severity and probability of potential threats to US businesses, the cost-benefit balance of enhancing Business Continuity in response to these potential threats has clearly shifted post-September 11. There are a number of steps, described below, that may help achieve a common general view of sound practices for business continuity in mission critical environments.

5.1. Define the Scope

A core question when conducting disaster recovery contingency planning is the range and scope of scenarios that financial institutions realistically need to consider when conducting business continuity planning. There are a number of scenarios that would affect entire geographic areas such as explosive devices, biohazards and natural disasters. Such scenarios could render a large area inaccessible and could harm or disperse an organization's critical employees. Other scenarios might deal with "cyber terrorism" which is aimed at computer networks and systems rather than a particular physical location. This threat is the focus of other efforts within the public and private sectors and is not covered here. It is painfully clear that preparing for impacts to a single facility are no longer sufficient.

Scenarios must also encompass targeted attacks on entire elements of the financial system. The financial services industry must now consider how to achieve greater geographic diversity among major financial institutions and clearing and settlement providers in order to withstand events of greater geographic scope than previously considered. Many now see the need to plan for extended periods of inaccessibility of more than one operating site within the same area. Citywide disruptions may be the minimum benchmark for planning purposes going forward and the ability to withstand disruption of an entire metropolitan area or region must also be considered.

Expectations have also changed regarding the length of time an event may incapacitate an area, increasing the depth of required business continuity. For example, institutions whose operations or data centers were destroyed in the World Trade Center attack were often left operating at a backup site indefinitely, often without adequate redundancy for the backup site.

5.2. Establish Business Continuity Objectives

Business continuity objectives must be developed to be consistent with cost-effective, sound business operations and that take into account the impact that one critical institution's operations can have on another. These objectives should cover issues such as:

- Recovery time expectations for critical operations
- Recovery capacity or volume expectations
- Sound business continuity practices to support these objectives

Recovery time expectations will differ depending on the severity of the scenario. The expected times for institutions to recover from a localized power outage will be much less than that of a regional disaster with loss of life. The near-immediate "fail-over" capabilities provided by current technologies can support this objective: The events of September 11 demonstrated that institutions that had planned for and tested their ability to recover critical processing operations prior to the attacks fared significantly better in resuming normal operations than those who had not prepared as thoroughly. [2]

Operating capacity recovery objectives should be reassessed in the light of data collected post 9/11. Many institutions' backup arrangements were based on plans to handle a reduced volume of activity during a disaster scenario. We now know this may not be a valid assumption as Monday, September 17, was an exceptionally high-volume-trading day and the shifting of settlement timeframes for securities led to wide fluctuations in day-to-day settlement volumes[3] Some institutions found that systems and telecommunications backup lines were designed to operate at significantly lower volume, severely hampering their ability to complete all processing during the disaster.

5.3. Identify Key Elements

A coordinated industry-wide approach to business continuity planning requires identification of the critical operational components of the system that must achieve a high level of business continuity preparedness. It is also useful to identify the types of operations that may require the highest level of operational resilience, such as credit card processing and fund transfers, which could have worldwide implications. This should involve identification of the core markets and essential functions supporting these markets (e.g., trading, brokering, transaction execution etc.)

Sound practices for the financial sector necessarily include planning with non-financial institutions such as telecommunications, utilities and other vendors of infrastructure services. Third party providers of business continuity must also have a comprehensively tested business recovery plan. Institutions are already exploring methods of ensuring that diversity of telecommunications lines is achieved and single points of failure are eliminated. Establishing diverse telecommunications methods and moving toward wider geographic diversification of operations in the longer term may be more effective in addressing these vulnerabilities.

5.4. Testing and Crisis Management

Finally, the effectiveness of business continuity strategies needs to be assured through planning, testing and regular use. Testing and planning absorbs resources; many companies have found ways to integrate business continuity tests into their routine operations by actively switching live operations to alternate sites periodically. Some institutions found that routine testing of their business continuity plans as frequently as monthly or quarterly helped considerably in dealing with the crisis as compared to annual or less frequent testing.[4]

Coordinated testing of business continuity plans between institutions and their customers or suppliers has not been a common industry practice. Joint testing exercises are needed going forward to truly ensure preparedness. While some companies had conducted frequent back-up site testing prior to 9/11 the tests generally were conducted with the primary sites in operation handling the bulk of the coordination. It would be safe to say that these tests were not reliable and would not demonstrate preparedness for a real disaster. In the wake of September 11 many participants who found themselves operating from their backup sites discovered connectivity and communications problems in addition to over-crowding and a lack of resources and leadership personal. Simple things like taking notes became increasingly complicated without required resources including office supplies.

Institutions found that their sound business continuity planning and testing prior to September 11 helped in locating and communicating with staff during the initial hours of the crisis, making operational decisions and quickly restoring relatively normal operations. The importance of accurate and clear information flows, both internally and externally, was particularly evident during the week of September 11. During a crisis proactive, ongoing and honest communication regarding operational status to customers and counterparties can help others to make informed decisions.

The industry must determine whether a more coordinated approach to crisis management and communication needs to be developed. Since September 11, several public and private sector initiatives have begun to address the issue of coordinated crisis management communication within the industry and with regulators.

6. Discussion Questions

Initial queries suggest that companies are aware of the most important vulnerabilities to the common methods of DR planning. Most firms appear ready to commit the necessary resources to strengthen the industry as a whole, particularly if their industry peers adhere to similar standards.

1. What range of scenarios should companies be expected to consider when planning?

As noted above, there is a growing consensus that the industry must plan for events of wider geographic scope and greater physical disruption than in the past, including the loss or inaccessibility of critical staff or of widespread disruption to telecommunications or other services. Citywide disruptions may be the benchmark for planning purposes going forward and regional disruptions also need to be considered.

2. What are appropriate sound practices for business continuity planning?

Sound practices should provide for very rapid resumption of critical operations at increased volumes of activity following a wide-scale, regional disruption that could result in loss of life or the inability of staff to access at least one major operating area.

3. How broadly should the sound practices be applied?

Clearly, companies collectively need to plan for very robust business continuity objectives for the essential functions they provide to support core markets. In addition, the level of business continuity planning and investments made by one firm has consequences for other institutions, as a company is only as well prepared as their least prepared critical supplier.

Conclusions

The events of September 11, 2001 tested the disaster recovery plans of many companies both located in New York and outside the directly affected area. We now know that disaster recovery plans need to be revised to be more effective.

Input from the various corporate sectors and a well-established firms specializing in Business Continuity is essential to determine the most effective way to ensure that common sound practices are implemented as widely and as quickly as possible. The need for supervisory and/or regulatory standards, including government subsidizing to support the application of sound practices across various industries, is now evident. The industry may also be able to provide guidance on how to provide a high level of confidence through ongoing use or robust testing, including coordinated testing, that the plans implemented by individual institutions are effective and compatible across the industry.

References

[1] Data Center Institute Survey of 422 Senior Data Center Managers dated January 1,2002

[2] The Disaster Recovery Journal in volume 15, 2002

[3] The Wall Street Journal September 18, 2001.

[4] The Business Continuity Association, Fall 2001