Comments on Draft Interagency White Paper on Sound Practices to Strengthen the Resilience of the U.S. Financial System

December, 2002

SunGard Availability Services
http://www.sungard.com/availability.

Introduction

SunGard Data Systems Inc. ("SunGard") appreciates the opportunity to provide comments to the Board of Governors of the Federal Reserve System, Office of the Comptroller of the Currency, and the Securities and Exchange Commission on the *Draft Interagency White Paper on Sound Practices to Strengthen the Resilience of the U.S. Financial System* ("the White Paper"). This response reflects SunGard's perspective as both a supplier of business continuity/information availability services through SunGard Availability Services and as a processor of systems supporting "critical activities" for "core clearing and settlement organizations" through SunGard Investment Support Systems. SunGard shares some of the concerns expressed by a number of financial services firms commenting on the White Paper. These points include the time span to plan, implement and test revised continuity programs, the time objective allowed for "recovery," the degree of geographic separation between primary and secondary sites, and the requirement for splitting labor pools of subject matter expertise.

Based on our discussions with business leaders and government agencies and our more than 25 years of experience in business continuity, we offer technology-neutral commentary on specific areas of the White Paper. We seek to facilitate the ongoing industry discussions through this response.

Summary of Key Points

1. Notwithstanding the criticality of certain large institutions to the stability of global financial markets, the two conditions proposed in the White Paper, accelerated intra-day recovery/resumption with zero data loss, and a separation of 200 miles between primary and secondary sites, are technologically incompatible at this time. Additionally, more than 90% of financial institutions would have to replace their IT infrastructures to meet such distance parameters. Such massive reengineering represents decades of technology spending. SunGard recommends that the joint agencies clarify these requirements in order of importance: first, intra-day recovery and zero data loss; second, a prudent and technologically feasible separation between the primary and secondary sites.
2. SunGard recommends data mirroring in synchronous mode as the most cost effective continuity strategy for most market leaders. Separation of the mirrored sites by 10 to 20 miles will fit the risk profile of most companies except when they represent a significant portion of a national financial utility. The mirrored sites must provide segregation of key infrastructure components such as power, telecommunications, transportation and water supply. Tertiary backup solutions such as tape backup must also be required to protect against data corruption and cyber attacks.
3. SunGard recommends that financial firms strengthen their telecommunications infrastructure, to include diverse physical routing, multiple local central office access, and redundant network services beyond the central office level.
4. SunGard believes that cyber-attacks, which represent clear and present danger to all financial institutions, are not sufficiently addressed by the White Paper. SunGard recommends that contingency plans address potential cyber attacks by including an assessment of security threats and a strategy for enterprise-wide coordination of risk mitigation and response.
5. SunGard recommends that, in assessing staffing issues, companies should differentiate IT requirements from those of business operations departments. For example, outsourcing IT professional services for data center testing and recovery is very feasible, while relocating financial subject matter experts is problematic. If the institution's business model and risk profile does not support splitting subject matter experts, then it should deploy business recovery positions at a prudent and commutable distance of 10 to 20 miles.

6. SunGard strongly supports the joint agencies' agenda for more rigorous testing. SunGard recommends the implementation of common metrics for industry participants so improvements can be made to ensure collective survival in a regional outage. SunGard welcomes the opportunity to work with industry groups to create a program tailored to support these common metrics and testing initiatives for our customers.

Background

SunGard was deeply involved in the response to the September 11th attacks. Within minutes of the first plane hitting the World Trade Center (WTC), our Crisis Management Team was in action. SunGard supported 77 separate customer disaster declarations and was put on alert status by more than 125 additional companies. As of this writing, there are still companies in recovery mode at our facilities as a result of the attacks.

We also share the joint agencies' concern about widespread disruptions affecting the financial markets. However, all risk management involves the examination of probability or true likelihood of an occurrence. SunGard has performed an internal study of all of its disaster declarations for the past five years, including those from terrorist attacks. The causes are illuminating and as follows:

- 56% hardware, software, power, telecommunications problems
- 24% natural disasters such as fire, flood, earthquake
- 20% malicious intent including 9/11

If those attributable to 9/11 are isolated, then an even larger percentage of disaster declarations affect single institutions and are caused by simple-to-recover equipment failures. Analysis of recent disasters including the WTC collapse show that as one moves away from the epicenter of the disaster and the immediate collateral damage, the degree and type of damage diminishes in direct proportion to the distance. While a worst-case scenario must be taken into account, contingency plans should be weighted toward the *most likely* scenarios.

Definition of Relevant Terms

The White Paper states that the resilience of the U.S. financial system rests on the rapid recovery and resumption of critical financial markets and the activities that support them.

If an institution suffers a disruption, whether a hardware failure or natural disaster, before it can resume normal business operations, it must first recover to the point of failure. The following explanations will serve to frame the ensuing discussion:

The recovery time objective (RTO) is defined as the amount of time that has elapsed from the point at which a disruption has occurred until the specified business operation is restored and current business transactions can be applied.

The recovery point objective (RPO) measures the amount of potential data loss in number of hours from the time of interruption.

An additional consideration for strategic planning is prioritizing applications that can commonly be divided into different tiers of mission-criticality and recovery staging order. Most financial institutions will perform a business impact analysis (BIA) to establish these tiers for business applications; what is less obvious is that those RTO's and RPO's also must drive the choice of continuity technology strategies.

Tier 1 applications are the most mission-critical and require an RTO of 0 –2 hours and an RPO of zero data loss. They typically include funds transfer, online banking, trading and ATM applications.

Tier 2 applications require an RTO of less than 24 hours and have an acceptable RPO ranging from zero data loss to less than 2 hours. They typically include advisory services, reporting, email functions and call center.

Tier 3 applications have an RTO of greater than 24 hours and a typical RPO between 24 and 48 hours. They typically include administrative applications such as payroll, general ledger, and development or R & D efforts.

Because the operational efficiency of the financial services market is inextricably tied to the availability of information systems, any discussion of recovery and resumption strategies must begin with technology.

Driven by the need for quick response times, high volume transaction processing and zero data loss for Tier 1 applications, most large institutions have moved away from slower batch processing. Software and hardware products can be combined to reliably mirror data over a given distance, virtually eliminating data loss and data currency issues. These techniques are known as *High-Availability* (HA) data mirroring and/or replication. For market leaders with high throughput, HA eliminates recovery reliance solely on tape batch processing and the associated data loss risk.

Synchronous versus Asynchronous Replication

Whether using mirroring or replication, synchronously or asynchronously, organizations must strike a balance between data integrity, bandwidth requirements and cost as well as application performance.

Synchronous mirroring provides high levels of data integrity by confirming that a transaction has been completed at both source and target machines before an additional transaction can take place. The cost of obtaining exact copies is driven by the resultant bandwidth requirements to complete timely confirmation of successful write operations. Because the process is gated by the speed at which data travels back and forth over the network, application performance erodes dramatically as the distance between the primary site and the secondary site increases. Most experts agree that synchronous data mirroring works most effectively within 60 miles (or 100 kilometers.)

Asynchronous replication alleviates the performance issues with synchronous mirroring over distance by allowing transactions to build up in a cache when load levels are too high that are later sent as permitted by the network infrastructure. The tradeoff, however, is the risk of reduced data integrity and the possibility of lost transactions. If data has built up in a buffer under heavy loads, all transactions that have not been transmitted upon failure of the source system are lost. In addition, because individual transactions do not take place after confirmation of those immediately before it, there is concern about the consistency of data and write order.

Recommendation #1 – Clarify the Priorities

When choosing the appropriate recovery strategy, the joint agency white paper poses a conundrum: the agencies and, indeed, market participants themselves, seek a standard of intra-day recovery and resumption (within as little as two hours for some institutions.) To achieve this time objective, companies must employ technology solutions that are presently incompatible with the suggested 200-mile distance recommendation. SunGard recommends that the joint agencies clarify their requirements in order of importance to the industry: intra-day recovery and zero data loss, followed by prudent and technologically feasible distance requirements.

To understand the incompatibility of these two directives, examine the key measurements by which any recovery strategy in this arena must be judged: recovery point and recovery time for the protected data.

An institution synchronously mirrors its data to a near-region secondary site (within 60 miles.) Data will be written to the target before the write operation completes on the host system. This provides the highest possible level of data currency—at any point in time, the secondary site has the same data as the source. The goals of an RPO of zero transactions lost and intra-day recovery can be met. However, this configuration does not pass the proposed minimum separation test for distance.

If the institution splits the primary and backup sites by 200 miles, and compensates for performance degradation with an intermediate hop, it then loses the ability to achieve zero data loss. It cannot resume business until it has recovered to the point of failure and completed pending transactions. The RTO is likely to stretch beyond the window of intra-day recovery and that delay may prevent other market partners from meeting their own settlement obligations. Because data loss is an unacceptable risk to the financial sector, many large institutions have chosen near-region synchronous technology.

Additionally, if the 200 mile limit is upheld, more than 90% of financial institutions would have to replace their IT infrastructures. Such massive reengineering represents decades of technology spending and would significantly constrain capital availability for other production and growth requirements.

Recommendation #2 – Set Prudent Distance Requirements and Require Segregation of Key Components

Assuming that the intra-day recovery of clearing/settling operation is first priority of importance, coupled with zero data loss, then present technology constraints dictate compromise on prudent separation of primary and secondary sites. SunGard respectfully submits that regulatory guidance cannot be predicated upon technology promises of the future but should reflect current reality. We believe that a 10 to 20 mile minimum distance will meet most institutions' risk profile, except where such companies represent a significant portion of national financial utilities. The chosen strategy must provide segregation of key infrastructure components.

To best examine the implications of the technology and continuity strategies under discussion, we offer three simplified examples:

Case #1 – Active/Synchronous Geography with Dedicated Disk Configuration

A configuration that features active/synchronous replication and dedicated storage is ideal for institutions looking to manage operational performance, cost efficiency and risk reduction in their contingency programs. This is the most common strategy chosen for providing zero RPO in the recovery environment. The duration of RTO is largely tied to whether dedicated or shared processors are employed.

The primary processing site runs in full active mode, replicating data to a secondary site within 60 miles that utilizes dedicated mirrored disk. This site may be an alternate location of the financial institution or provided by a third-party vendor. To be effective, the secondary site must be located on a different power grid and use dissimilar water supply and telecommunications provisioning. SunGard believes that this is a minimum requirement that the agencies should address for "near-region" scenarios.

At time of disaster, connectivity is established with contingency processors and peripherals that can be either dedicated or shared with other potential users. One of the major benefits of this solution type is that it is possible to customize it to an institution's degree of acceptable cost/risk ratio. As dedicated pieces such as hot stand-by are added, the recovery time and risk goes down, however the overall cost goes up.

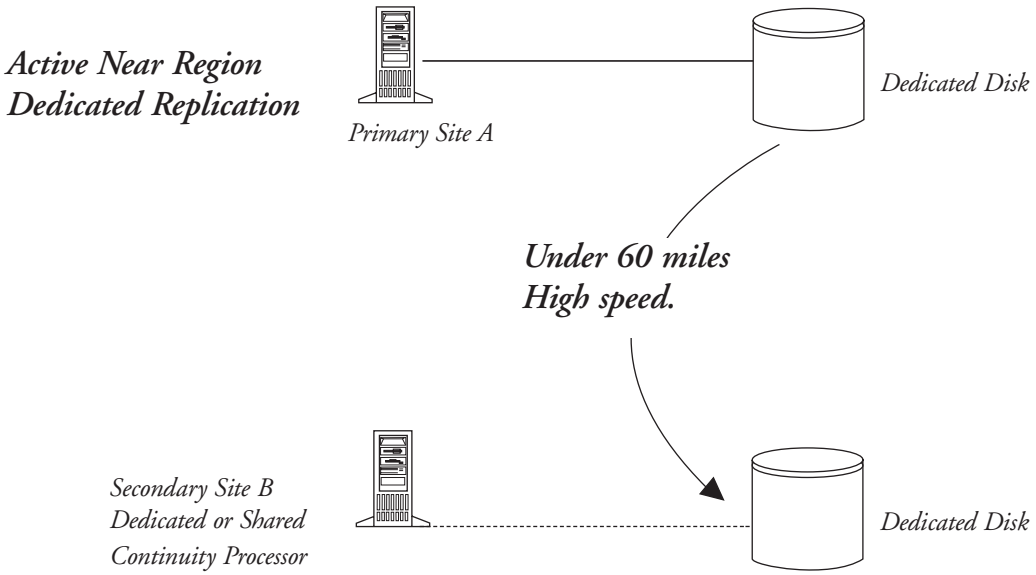

The 9/11 attacks on the WTC were the first wide-scale and real-life test of dedicated near-region synchronous data mirroring. Institutions employing secondary synchronous sites outside of Manhattan, across the Hudson River or farther, were successful at data recovery and prevented data loss where they employed this technique.

The joint agency white paper also requires that whatever technology solution is chosen, it must be tested to meet RTO and RPO objectives. Using this configuration, multiple copies of the same data must reside on the dedicated disk so that production data is preserved separately from test data. It is not a testable solution if the configuration is not linked to hotsite or alternate processing location capabilities.

This configuration meets the joint agencies' requirement for intra-day recovery and can be implemented at a prudent distance of less than 60 miles. It often strikes the best balance for business drivers, operational cost and regulatory requirements. It also enables the combination of dedicated disk and processor solutions covering data loss and data integrity for Tier 1 applications, as well as shared solutions for tier 2 and tier 3 applications. Many lesser applications can tolerate 24 hour recovery, in which case companies can employ conventional (and cheaper) recovery methods. Co-locating in the same facility assists data interaction across applications, thereby improving the timing for business resumption. finally, if institutions use third party vendor hotsite solutions, they can easily upgrade subscriptions when server technology changes. However, they must also have a tertiary backup to guard against malicious intent cyber attacks and the potential of simultaneous database corruption.

Pros	Cons
• **In most cases, will support intra-day recovery and resumption** • **No data loss** • **Recovery timeframe is very short, depending upon processor boot duration, (mainframe = less than 2 hours, open systems between 2 and 12 hours)** • **Less costly than alternatives that require duplicate servers** • **Accommodates a hybrid solution with synchronization of both dedicated and shared recovery solutions (tier 1, tier 2, and tier 3)** • **Proven solution, well accepted by the industry** • **Most scalable for growth and business line changes**	• **By itself, not a regional disaster solution since the distance between production and recovery centers is limited due to performance implications** • **Remains serious information security concern since this strategy does not protect against data corruption or cyber attacks** • **Requires tertiary data protection such as traditional tape backup and storage.**

Case #2 – Active/Active Synchronous Geography with Dedicated Servers/Disk Configuration

The dual-active configuration requires duplicate processors and additional network, making it the most expensive strategy to implement. As companies utilize high volume 24 x 7 processing, this strategy is commonly used to improve throughput. The primary processing site is tied through a synchronous network to a duplicate secondary production site. The closer they are, the greater the performance benefit. Thus, the two sites are usually very close to one another. This is not an effective recovery strategy without tertiary backup.

The Active/Active Synchronous configuration can represent one of the greatest contradictions in planning for adequate recovery. Before September 11th, many companies in the financial sector believed that smoke-and-rubble disasters were extremely unlikely. Consequently, they used a clustering configuration—multiple servers at one site writing to dedicated disk, a solution designed to protect against server outage. While yielding extremely dynamic performance, this strategy is unacceptable because it does not address redundancy of either data or location.

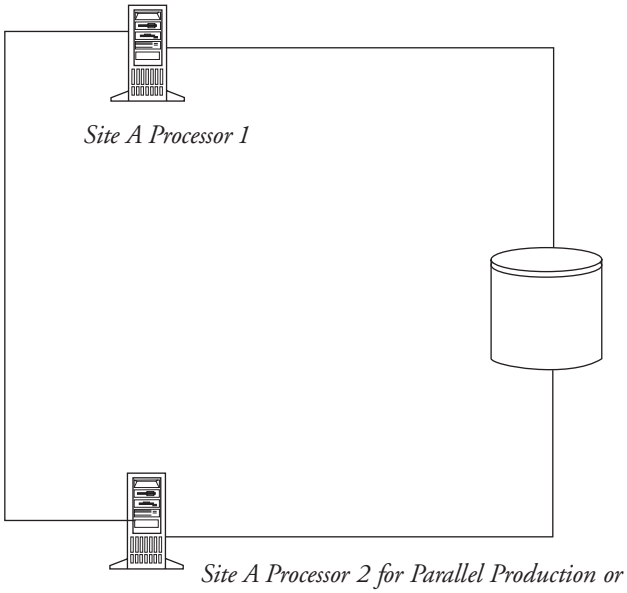

Clustered Site with Server Failover

Clustering is a common high availability application or performance strategy that frequently gets confused with having a recovery strategy.

Other permutations have the second server remaining in very close proximity, housed on a different floor within the same building, in an adjacent building or within the same campus location. But, as the 9/11 outages in lower Manhattan proved, these scenarios proved undesirable. Both locations were subject to power and communication outages even if they were not directly affected by the collateral damage of the towers' collapse. As a result, clustering in proximity will prove unacceptable no matter what distance the agencies determine in their final regulatory guidance.

Active/Active Synchronous Replication

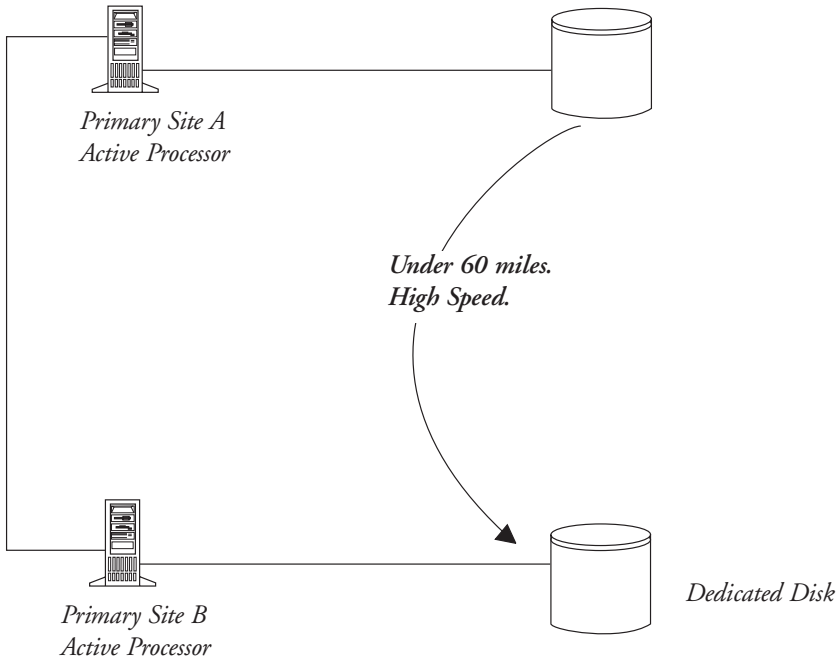

Because there are additional stresses put on active/active synchronous technology and the performance degrades significantly, it is usually a near-region solution. In most cases, this geo-plex solution operates well within a 60-mile radius telecommunications link. The secondary location may belong to the financial institution or a third-party outsourcer.

Active/Active can be used as a load balancing strategy for production servers. But in most instances, a master/slave relationship exists due to the difficulties of database sharing. There is also concern over testing an active/active strategy. While it is certainly possible to perform failover testing, there is inherent risk because the test could precipitate an outage.

This strategy also meets the intra-day recovery objective when implemented at a prudent distance. Due to cost considerations, it is an option only for those institutions with very large budgets in addition to exceptional availability and throughput requirements. Our experience has shown that companies who undertake the considerable expense of the active/active cost structure do often not provide for the tertiary backup requirements that would protect them in the event of a regional outage or malicious-intent cyber attack.

Pros	Cons
• **Supports intra-day recovery and resumption** • **No data loss** • **No recovery timeframe—immediate failover** • **Load balancing capability** • **Flexibility for scheduled equipment maintenance, etc.**	• **Highest cost solution requiring:** – **Exact matches for equipment and upgrades** – **Large network requirements** – **Additional software licensing fees** – **Capacity to support all applications as tier 1 priority (there are no distinctions between dedicated and shared solutions)** • **No protection in the event of a regional disaster** • **Hard to protect against a rolling disaster by splitting off a copy of the data (point in time)** • **Active/Active is the most vulnerable to cyber attack and insider malicious intent since both copies of the data are exposed simultaneously** • **Requires tertiary data protection such as traditional tape backup and storage at additional cost** • **Distance affects production response and throughput** • **Cost is prohibitive for many companies** • **High degree of complexity for day to day management** • **No scalability (everything is 1 to 1)** • **Difficult to preserve the capacity required for recovery while in production mode.**

Case #3 – Active/Multi-Hop Intermediate/Alternate Secondary Site

In the Active/Multi-Hop configuration, the primary processing site runs in full active mode and mirrors data to dedicated disk at an intermediary location (B) well within 60 miles of the primary. Multiple copies of the data (at different points in time) are stored at the intermediary which forwards net changes to dedicated disk at the secondary site that can be located anywhere. The intermediary location helps to maintain processing efficiencies for the primary site while data is being moved to its ultimate destination. If the joint agencies mandate a 200+ mile distance between the primary and secondary sites that proscribes synchronous solutions, this configuration answers the need to go beyond a 60-mile distance and yields the least data loss. However, timing requirements are not consistent with the desired goal of intra-day recovery.

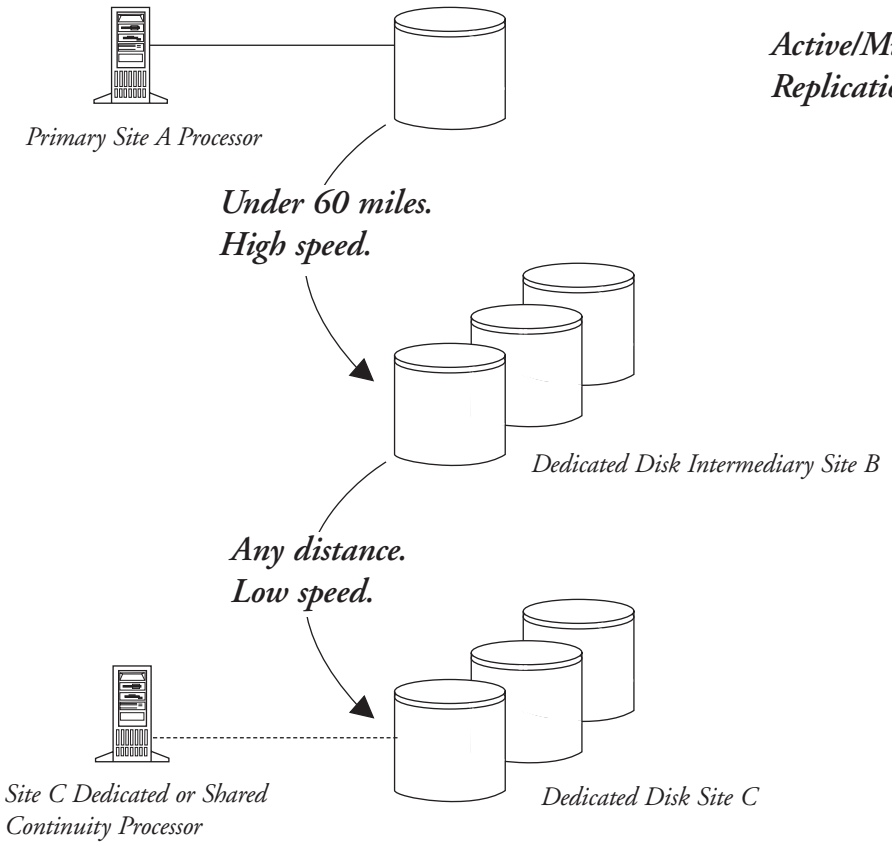

To function as the main resumption site, the destination site will have to be provisioned with processors and higher bandwidth network. For the most common failure scenario (e.g., hardware, software, power) where the primary fails, but the intermediary site remains intact, data would not be lost while bringing the recovery processor online at the secondary site and re-pointing production to the intermediary location. For a regional outage, however, data loss becomes more probable and the recovery timeframe extends out.

At time of disaster, connectivity must be established with contingency processors and peripherals that can be dedicated or commercial shared subscription. This strategy is most useful when the organization has a specific need for geographic distance away from a damaged area. For instance, the second hop moves the copy of the data closer to a large commercial hotsite with expanded capabilities. As with Case #1, organizations must determine the acceptable degree of cost/risk. When adding dedicated processors and network, the recovery time and the risk goes down but the overall costs go up.

Active/Multi-Hop is an acceptable solution for institutions that have distributed operations. For institutions concentrated in New York, however, it becomes problematic and expensive as an internal solution. In addition, this strategy requires complex synchronizing of the applications which is usually not included in the design at the front end and which usually requires an intimate knowledge of them most often found in the primary production facility. It is easy to set up for mainframe, Unix and NT platforms; other platforms are more difficult. Finally, the solution is testable only if multiple copies of the data are maintained at both the intermediary and secondary sites and there is access to processors and sufficient network connectivity.

Pros	Cons
• **Zero data loss for non-regional disaster if the intermediary site is not affected (and configured correctly with at least three copies of data)** • **True regional disaster solution if both primary and intermediary are lost** • **Recovery timeframe can be relatively short depending on how much is dedicated and the platform allocation between mainframe and open systems (a minimum 2 hours, up to 24+ hours)** • **Good solution for redundancy** • **Good mix of synchronous benefits with regional protection** • **Offers protection against self corruption of the database and cyber attacks** • **Enables hybrid solutions of dedicated recovery and traditional shared recovery solutions for less critical applications.**	• **Higher cost (compared to Case #1) requiring:** – **Multiple (3 or more) copies of data in both intermediary and secondary remote locations)** – **Additional long distance network** – **Multiple processors, multiple software licenses (if self provisioned)** – **Additional staff for process management** • **Data loss for regional disaster is minimum of 2 to 4 hours if intermediary is lost** • **Will probably not meet intra-day recovery and resumption requirement** • **Much more complex to synchronize applications** • **May be more difficult to connect with value chain partners**

There is an extremely high concentration of financial companies in the greater New York City area that has historically enjoyed a relative immunity from long-term regional outages. However, provisioning for regional disaster is not foreign to other areas of the country where vulnerability to extensive natural disasters is common. As such, SunGard has supported regional recovery efforts from major hurricanes, floods and earthquakes.

Our own infrastructure has been constructed to disperse risk with power independence, telecommunications diversity, and facilities distribution. In the event of another NYC outage, we can utilize a roll-back strategy to access a broad recovery inventory of facilities, equipment and skilled personnel knowledgeable in system restoration and operation.

If the joint agencies support prudently separated near-region solutions, we expect to see greater inspection of comprehensive planning for all value chain participants, metropolitan incident management programs and emphasis on eliminating single points of failure for power, water, transportation and telecommunications.

Recommendation #3 – Require Telecommunications Redundancy

SunGard believes that the financial and telecommunications regulatory bodies should work closely together to formulate joint recommendations of how financial firms should provision for telecommunications redundancy. Examples include critical local telecommunications links from the primary processing facility with diverse routing, more than one central office for redundancy and examination of circuit designs to avoid single points of network failure beyond the central office.

In order to sustain function in a regional disaster, financial firms must take measures to strengthen their telecommunications resiliency in addition to IT and labor resources.

Disruption at a telecommunications central office or transmission circuit can interrupt functions at the firm's primary processing location just as much as a failure within the processing location itself. Such network components as local central offices and "last-mile" local loop circuits are, by design, in close proximity to the primary processing facility and exposed to the same risks as the processing location. Therefore, telecommunications redundancy at the primary processing facility must be a mission critical element of overall infrastructure.

Similarly, telecommunications redundancy for secondary/tertiary sites must be incorporated into a contingency plan and must provide the means to establish connectivity with other market participants involved in the clearing and settlement process as well as second points of aggregation for access to buy/sell participants. Network links associated with this contingency processing must be completely independent of services supporting primary site processing. To this aim, SunGard recommends that organizations explore such redundant solutions as connection to multiple "network pooling agents."

Our many years of recovering clients from every kind of disaster have taught us to design our own telecommunications infrastructure to meet these same precautions. Our architecture standards include diverse physical routing, multiple local central office access, and redundant network services beyond the central office level. This resilient telecommunications infrastructure allows the redirection of large, complex networks to alternate facilities should one SunGard facility be compromised by the same disruption affecting our customers.

Recommendation #4 – Protect Against Cyber-Attacks

SunGard believes that cyber-attacks represent clear and present danger to all institutions within the economic infrastructure and are under-represented in the agencies' paper. Contingency plans should include a heightened understanding of security threats and enterprise-wide coordination of risk mitigation and response.

Attacks against information targets happen everyday: denial of service attacks to disrupt the flow of information, attacks against confidentiality by compromising sensitive data, or attacks against integrity involving deliberate falsification or corruption of vital records. In addition to the startling increase in security incidents recorded by the CERT® Coordination Center, both FBI and CIA security officials warn that attacks against western economic targets are likely. In their view, various terrorist groups are developing the necessary technological expertise to consummate an effective strike that could have global repercussions.

The first requirement of defending against a "logical disaster" is the ability to detect and identify information security attacks. Proactive security management is essential to identify and stop ongoing attacks and establish new barriers to prevent future attacks. Procedures should be developed so that all employees know what to do should an information security attack occur.

Institutions must also recover any lost data and have disaster recovery strategies in place so that any damage to networks and systems as a result of the attack does not disrupt normal business operations. Information security professionals must integrate their efforts with the institution's continuity program.

Recommendation #5 – Differentiate IT Staffing from Subject Matter Expertise for Recovery Strategies.

For staffing issues, SunGard recommends differentiating IT requirements from those of business units. For example, IT professional services for data center testing and recovery can be outsourced, while relocating financial subject matter experts remains problematic. If the institution's business model and risk profile does not support splitting subject matter experts, then it should deploy business recovery positions at a prudent, commutable distance of 10 to 20 miles.

Business continuity plans must include people as well as technology. Prior to 9/11, companies planned for their data centers but did not explore the necessity of keeping people and information connected.

After the 1993 WTC bombing, many financial companies implemented distributed system architecture and moved data centers out to less expensive, less threatened real estate. And many already have efforts underway to create even greater geographic dispersal by separating data center staff from business subject matter experts.

SunGard encourages the agencies to take a more flexible approach for staffing issues. A one-size-fits-all mandate will cause significant disruptions to an industry that must be responsive to pressures from changing economic cycles. Instead, regulatory guidance should focus on examining whether an organization's business recovery plan addresses its specific risk profile and whether its recovery strategy stages both equipment and personnel in priority matching the appropriate degree of mission-criticality.

Outsourced IT professional services for data center testing and recovery can supplement an institution's strategy. For example, SunGard has professionals all over the world with extensive continuity expertise who spend every day configuring systems and connectivity for clients. To facilitate intra-day recovery of mirrored solutions, SunGard also offers managed services to stage the process switchover. At time of disaster, our personnel can pull results of the last test of a configuration and begin restoration of systems, reducing the time until applications are available to users.

Finding an out-of-region recovery solution for business operations departments remains problematic. Complete real estate redundancy is prohibitively expensive. We also believe that business processes within the financial services industry benefit from the "synergy of proximity" in metropolitan areas.

There are many arguments against the split labor model for financial subject matter experts. It is not only an expensive proposition to change a company's business model; it is also more expensive to run split operations. For operational efficiency, most institutions drive to combine "like" processes in one location. Split operations will require new layers of control and will invariably introduce new risk. Additionally, such change cannot happen as quickly as proposed in the agencies' white paper.

Institutions who have chosen to concentrate their operations in metropolitan areas like New York report that it is easier for them to attract talent. Staffing difficulties will undoubtedly escalate under split-labor pool requirements.

The post-9/11 continuity market will drive the evolution of new business unit or work group solutions. Consideration of operational resilience has expanded from fiscal capital and interest risks to encompass the ability of a company's infrastructure and business processes to withstand unexpected interruptions. SunGard sees a trend toward blending both dedicated and shared work group strategies, with perhaps 25% of the recovery seats dedicated to such Tier 1 applications as funds transfer, and the remainder in shared solutions connected by a redundant network. We also foresee the implementation of more flexible technology for porting market data feeds to virtually any seat.

Recommendation #6 – Require Testing with Common Metrics

SunGard recommends the development of common testing metrics to improve the ability to survive an outage since it is important that continuity arrangements of critical market players be both effective and compatible with others in their sector. Testing parameters must include the ability to recover the back office business processes as well. Alternate telecommunications connectivity should be tested regularly in conjunction with any IT or labor pool contingency plan testing.

To ensure a high degree of confidence in recovery plans and the resumption process, SunGard strongly supports the joint agencies' agenda for more rigorous testing. Enhancing the testing process will accomplish the following:

- Confirm the ability of chosen recovery strategies to match business objectives
- Identify vulnerabilities and exposures so that they can be remedied
- Address ongoing change in the organization's infrastructure.

SunGard welcomes the opportunity to work with industry groups to create an ad hoc program and metrics tailored to support such testing initiatives for our customers. Scheduling a cross-industry test for specific calendar dates when the market is closed will eliminate the need to dilute individual customer test time for what may be only one portion of their continuity program.

With more than 14,000 tests in 2002 alone, SunGard is highly experienced at identifying ways to enhance a testing program such as:

- Properly defined scope and objectives
- Development of realistic timeframes
- Identification of interdependencies
- Documented and detailed recovery procedures
- Improved methodologies for data backup
- Verification of hardware/software configurations
- Assignment of adequate test resources
- Implementation of end user connectivity

Testing is the ultimate assessment of operational resilience, the ability of a financial institution to continue to function through elimination of risks associated with technology, processes and people.

Conclusion

While we acknowledge that the recovery paradigm has changed since September 11th 2001, we believe one size does not fit all when it comes to choosing the correct strategy for business. financial institutions must practice principles of good risk management: identify exposures, mitigate where possible and then choose an appropriate solution to address the unexpected.

SunGard has built a global infrastructure that can be leveraged to assist our clients no matter what the final resolution. However, we believe that a reasonable near-region solution can be an acceptable best practice if it includes continued vigilance and thoughtful planning to include "what if" scenarios for regional disaster.

Although the White Paper is aimed specifically at the very largest institutions in the commercial and clearing sectors, SunGard believes that smaller financial institutions will eventually adhere to the same standards in order to maintain their partner relationships.

SunGard supports the agencies' efforts to safeguard the financial sector by improving expectations for business continuity. The challenge within our industry will be to identify the necessary requirements and develop cost-effective strategies to assist financial agencies in reducing systemic risk and promoting rapid recovery.

About SunGard Availability Services

SunGard Availability Services is the pioneer and leading provider of information availability services, helping to ensure that more than 10,000 clients in North America and Europe have uninterrupted access to their business-critical information systems. With over 3 million sq. ft. of hardened facilities, it offers a complete range of information availability services for more than 30 technology platforms, from 48 hour disaster recovery hotsites to always-on, high-availability infrastructure, co-location and electronic vaulting services. SunGard also provides technology and systems management services for application and data center outsourcing, as well as business continuity consulting services and planning software.

SUNGARD®

AVAILABILITY SERVICES