

White Paper: Using Reserve Systems for Business Continuation

By Michael Miora, CISSP

President & CEO

ContingenZ Corporation

mmiora@contingenz.com

Table of Contents

Establishing a Strategy: Continuation or Recovery	2
A Survey of Strategies	3
Cold Sites	3
Hot Sites.....	3
Reciprocal Agreements	4
Internal Redundancy	4
Mobile Data Centers	5
Priority Replacement Agreements	5
Commercial Recovery Services.....	5
Reserve System Strategy.....	7
Summary of Strategies	8

ContingenZ
Corporation™

Training ▪ Education ▪ Consulting

ContingenZ Corporation

www.contingenz.com ▪ info@contingenz.com

Incident Management: Recognize, React, Respond

ESTABLISHING A STRATEGY: CONTINUATION OR RECOVERY

In this white paper we will discuss the specific strategies to continue the most time-critical function, called Category I functions, identified during the Business Impact Analysis. We will briefly review some traditional recovery strategies, such as hot sites, warm sites, and cold sites, and compare them to the more modern technique we call Reserve Systems. We will also describe how to make good use of Internet and client/server technologies, and of high-speed connections for data backup, for making electronic journals, and for data vaulting.

Focusing in on continuation rather than recovery.

Recovery strategies are defined differently from continuation strategies. The recovery strategy is the overall plan for resuming a function, or set of functions, in a near-normal work mode. The continuation strategy is the plan for immediate or nearly immediate operations of a key function, even if it results in significantly degraded performance, or limited capability. The continuation strategy is most often applied to a small set of functions rather than to a more complete set, and it is often only temporary. Although the recovery strategy may be used for several weeks or longer, the continuation strategy typically survives for hours or days.

The BIA defines the timeline for recovery for each category of functions. The recovery strategies depend on factors such as the complexity of the functions, the amount of information required to carry them out, the number of people involved in performing the function, and the amount of interaction between this function and other functions. The strategy of choice – continuation or recovery – for each function is based on the timeline and the specific factors relating to it.

Traditional recovery strategies include hot sites, warm sites, and cold sites. The most common recovery strategy uses a commercial service provider to maintain an on-demand operating capability for well-defined systems with internal redundancy and backup capability for a variety of other systems.

A continuation strategy usually depends on internal resources rather than on commercially available providers. Continuation strategy requires that equipment be available instantly, with current data. Instant availability means that the systems must be local, or at least connected, yet survivable. The continuation strategy is becoming more common each passing crisis.

The continuation strategy often relies on reserve systems, a concept pioneered over the past five years. This strategy, which is the focus of this paper, uses high-performance microcomputers, PDAs, and other small computers, residing off site at homes, branch offices, and other locations easily accessible to employees. The reserve system provides immediate, but limited, functionality so that operations can be continued after a disaster, while awaiting full restoration or recovery.

The reserve system provides immediate, though limited, functionality.

The continuation phase begins at the time of the disaster with the goal of supporting the Category I functions as best as possible, and within the time frames defined by the BIA.

A SURVEY OF STRATEGIES

There are numerous strategies of dealing with the restoration of functions through a recovery process, but only one for continuing operations with minimal interruption. This section surveys some of the strategies for achieving the former; the next section describes the strategy for achieving the latter.

Cold Sites

A cold site is a room or set of rooms in a “ready-conditioned” environment. It is a fully functional computer room with all of the required facilities, including electrical power, air conditioning, raised floors, and telecommunications, all in operating condition and ready for occupancy. Missing from this fully functional room, because of high costs, is the computer equipment, including the processors and all required peripherals. Cold sites can be owned by one company or shared by several companies.

The major advantage of a dedicated cold site, somewhat diminished in the shared case, is simply that for a relatively low acquisition or leasing cost, the site is guaranteed to be available over the long term to the owner or lessee of the site. To be effective, a cold site must be distant enough from the main facility so that a disaster that makes the primary facility unusable will likely not affect the cold site.

Low cost but with limited effectiveness.

There are several disadvantages to cold sites. First and foremost is that ordering, receiving, installing, and powering up the computer system can take many days. Once the system is functional, the cold site becomes a dedicated facility ready to perform all functions in a controlled environment, but the time required to achieve this state can stretch into a week or more. Few organizations can rely on a cold site as their primary vehicle for disaster recovery. A secondary disadvantage of cold sites is the inherent inability to test the recovery plan. An untested plan is an unreliable plan—testing a cold site would require obtaining all of the requisite equipment, installing the equipment, and performing operations on the new equipment. Few organizations can afford the costs associated with such testing. There are also hidden pitfalls with this strategy. Key equipment may be unavailable for immediate delivery. Communication lines may be untested and unreliable. This strategy is generally not a desirable first line of defense. It can, however, be a part of a larger overall strategy that includes other types of sites.

Hot Sites

A hot site is a facility ready to assume processing responsibility immediately. The term usually refers to a site that contains equipment ready to assume all hardware functions, but requiring massive restoration of data, and an influx of personnel to operate the equipment. The hot site cannot, therefore, begin total processing instantaneously. In actuality, the site is only warm, with an ability to get hot fast. A hot site can be dedicated to one organization, shared by several organizations, or leased from a specialty company, called a

Fast for most functions but with high cost.

commercial recovery service provider. The provider option is described in a later section.

The primary advantage of a hot site is the speed of recovery. The time to resume processing is defined as the time to reach the facility with people and backup media, plus the time to restore data and programs from the backup media, to test the operations, and to go live with a fully functional system. For larger systems, this period can range from less than two days to almost a week.

The primary disadvantage of a hot site is the cost associated with acquiring and maintaining the fully equipped site. For this reason, most organizations choose to share the cost of a hot site with other organizations, who may be sister companies or only neighbors. In any event, there is a large cost associated with maintaining the site, and ensuring that it is updated with every change to each of the participant's requirements, while still maintaining compatibility with all of them. One of the most common solutions to this problem is to use a service provider.

Reciprocal Agreements

Reciprocal agreements were often used in the earlier decades of disaster recovery planning, but are uncommon today. A reciprocal agreement is an arrangement between two or more companies in which each agrees to make excess capacity available to the others in case of a disaster. The major advantage is the apparent low cost of this solution. The major disadvantage is that these arrangements rarely provide the needed computing power. A major issue with these arrangements is maintaining compatible systems. If one company changes a processor, it may find that its [partners] systems cannot perform adequately, even in degraded mode. If the host company faces a crisis or deadline of its own, a reciprocal company may find itself with no computer power at all. These arrangements are seldom testable, because it is the rare company willing to shut down operations to help a partner perform a disaster recovery test.

**Not in widespread use
any more.**

Internal Redundancy

A strategy of internal redundancy requires that a business have multiple facilities, geographically dispersed, with similar equipment in each site. If there are data centers at several sites, then the alternate data centers may be designed with excess capacity to support a potential failure at another site.

The major advantage of internal redundancy is that the organization maintains complete control of all equipment and data, without relying on any outside company to come to its aid. The excess capacity at the various alternate sites must be carefully protected, and management must exercise diligence in budgeting and operations. Careful intra-company agreements must be crafted to ensure that all parties are aware of, and agree to, the backup arrangements. Internal redundancy can be an effective solution in cases where temporarily degraded processing can still provide sufficient support to meet time-

**A management challenge
and difficult to test.**

line requirements. If degraded performance is not an acceptable option, then the cost of the excess capacity will probably be too high. If reasonable degradation is acceptable, then those costs can be manageable.

Internal redundancy can also be difficult to test. Testing requires that processing be shifted to a recovery mode. Unlike external, separate computers, all of these redundant systems would be operational. Testing one disaster recovery plan requires affecting a minimum of two corporate locations. A failed test that causes a system crash or other problem can have damaging consequences.

Mobile Data Centers

Mobile data centers are transportable units such as trailers outfitted with replacement equipment, air conditioning, electrical connections, and all other computer requirements. Mobile data centers are most often used for recovery of midrange and PC LAN systems. The primary advantage of the mobile data center is that it can be activated quickly at reasonably low cost. The primary disadvantages are the expense of testing such a facility, and the possibility that a local or regional disaster will prevent successful activation. Deploying a mobile data center requires careful planning. Land must be available to accommodate the transportable units, with outside parking lots as the most common resource. Local government and municipal regulations must be researched in advance to ensure that such units do not violate ordinances and can arrive as certified for immediate occupancy. External power and communications hookups also must be available.

Difficult to test and most often used for a building-specific incident.

Priority Replacement Agreements

Some computer vendors support priority equipment replacement agreements. These are arrangements in which the vendor promises to ship replacement equipment on a priority basis. For midrange systems, this is often an agreement to send the "next off the line" system to the priority customer. The major advantage of this strategy is its low cost. However, if the vendor is not currently manufacturing the required system, or if the assembly line is down for any reason, and if equipment stocks are depleted, there may still be a significant delay in receiving the equipment. This is the major disadvantage. This strategy also assumes the disaster recovery plan makes an alternate facility available in case the primary facility is damaged along with the equipment being replaced.

No reliable or fast enough for key functions.

Commercial Recovery Services

Commercial recovery service providers can support a combination of the strategies discussed above. These companies generally provide three major benefits: cost sharing, reduced management needs, and diminished risk of obsolescence. First, they spread facility costs across multiple subscribers so that each subscriber saves as compared with building and maintaining a comparable, privately held capability. Because all subscribers share the same physical

Common and effective but expensive.

space, each pays less than would be needed to maintain such a site independently.

The second major benefit of using a commercial provider is that these companies eliminate the need for the subscriber to manage backup resources. Management and maintenance of such a site by an individual business could be a heavy burden, but the provider's primary focus is on managing, maintaining, and upgrading the equipment and sites. The subscriber company can be assured that the equipment is exercised and serviced regularly, and that peripheral equipment, power systems, and facility support structures are properly maintained. A properly run site will also provide security, safety, and compliance with evolving rules and regulations, and competent staffing. The provider assumes full responsibility for these functions during normal operations, and continues support during times of crisis. The subscriber brings its technical personnel, while the provider leaves its facilities staff in place.

The third major benefit centers around today's fast pace of hardware evolution. A subscriber company will typically lease a hot site or other service for a five-year period. During that time, hardware platforms will evolve. The subscribing company can protect its lease investment by ensuring that system upgrades are reflected in the leased equipment configuration for reasonable extra charges. A business that provides its own hot site must upgrade the hot site whenever hardware, and sometimes software, changes are made to the operational systems.

The disadvantage of commercial recovery services is in the obvious risk that the hot site may not be available in an emergency. Indeed, if there is a local or regional disaster that affects numerous subscribers, there could be significant contention for the provider's resources. To address this issue, providers typically maintain hot sites in geographically dispersed areas. Although it is likely that in the case of a local or regional disaster a subscriber would need to use a hot site further away than planned, it is unlikely that the subscriber would be left completely without the prearranged resources.

RESERVE SYSTEM STRATEGY

The newest of strategies is the reserve system, which is a small replica of a portion of an operational system meant for use during the first few days following a disaster. The reserve system provides continuation of key functions for short durations, although in degraded mode. The reserve system usually resides off site at an employee's home or at another corporate office. Another version of the reserve system is also kept on site. A reserve site may be equipped with a microcomputer or a minicomputer ready to assume functioning in case the primary system becomes unavailable. This meets the important criteria of a reserve system, which must be fast and easy to activate, simple to move, low in cost, testable, available, and highly reliable.

The reserve system concept was not feasible until client/server technology emerged, and Internet telecommuting with high-speed communications became accepted and readily available. Web-centric processing and remote application server technologies can provide powerful reserve systems for disaster recovery.

Not really feasible until the last few years.

Proper security precautions must be taken to protect proprietary, confidential, and critical information stored on these reserve systems. Strong encryption safely protects against theft while redundant systems protect against other losses. The reserve system is a quick-response, short-term solution intended to solve the problem of immediate continuation even in the case where employees may be unable to travel outside their immediate residence areas.

Over the past three decades, the very nature of the threats we considered and the preparations we made have changed dramatically. In the early days, major fears included equipment failures and minor geographic events such as storms and small floods. In the even of a major catastrophe, we felt confident that our plans would unfold as well as our neighboring companies plans would unfold – perhaps there would even be some cooperation between companies.

In the decade of the 1990s and in the beginning of our new millennium, however, the criticality of systems combined with an upsurge in man-made disasters forced us to rethink our practices, scenarios and procedures. The first bombing of the New York World Trade Center in 1993 to the bombing of the Murrah Federal Building in Oklahoma City in 1995 to the heinous destruction of the World Trade Center twin towers in 2001 have finally led us to the point where we must consider "artificial" disasters in our planning as much as natural ones. The Disaster Recovery planner is now in a position of planning for disasters and building recovery scenarios that include terrorism as well as tornados. The job is difficult to perform, but the techniques and technologies are the same as before.

Highly effective, low cost, and available fast for key continuation functions.

SUMMARY OF STRATEGIES

The table below summarizes the advantages and disadvantages of each of the strategies described the paper.

Strategy	Activation	Cost	Testability	Availability	Reliability
Reserve Systems	Very Fast	Low	Excellent	High	Good
Internal Redundancy	Fast	Medium	Poor	Medium	Good
Commercial Providers	Fast	Varies	Excellent	High	Excellent
Hot Site	Fast	High	Excellent	High	Excellent
Mobile Data Centers	Medium	Low	Medium	High	Good
Reciprocal Agreements	Slow	Low	Poor	Low	Poor
Priority Replacement	Slow	Low	Poor	Low	Poor
Cold Site	Slow	Low	Poor	High	Poor