



## **Disaster recovery strategies for internal recovery**

---

<b>Contents</b>	
<b>2</b>	<b><i>Introduction</i></b>
<b>3</b>	<b><i>The Y2K effect</i></b>
<b>4</b>	<b><i>Fall of distributed systems</i></b>
<b>4</b>	<b><i>New challenges</i></b>
<b>5</b>	<b><i>New suite of enablers</i></b>
<b>6</b>	<b><i>Internal recovery considerations</i></b>
<b>7</b>	<b><i>Critical strategy considerations</i></b>
<b>8</b>	<b><i>Direct expenses</i></b>
<b>10</b>	<b><i>Indirect expenses</i></b>
<b>11</b>	<b><i>Summary</i></b>

***Introduction***

In the past three years, dynamics for disaster recovery has changed more dramatically than in the previous 40 years combined. Until recently, internal recovery focused on mainframe recovery or failover and included a limited number of servers used to access the mainframe. Data was replicated between two local sites and in the event of an operational outage or facility failure, processors picked up the additional work load of failed processors, or alternate processors were brought up to begin processing in as little as two hours. Distributed systems used as open systems and Wintel servers implemented a few years ago were not deemed critical, and not included in the internal recovery process.

Many of the original considerations for an internal recovery site are no longer relevant. Among other technology improvements, asynchronous data replication and server virtualization has dramatically changed the fundamental underpinnings for local or failover sites. The continued evolution of servers, data replication, data volume and the direction, in which they continue to move, require companies to reexamine strategies that were developed and implemented as recently as two years ago.

This paper will discuss the elements that influenced the decisions of the internal recovery strategy used in the past. It will look at the evolving direction for strategies going forward, and focus on key trends such as: data center capacity, emerging data replication technologies, regulatory direction, and how recovery and continuous or high availability fit into strategic discussions.

---

Highlights

---

***In many cases, disaster recovery used the same hardware, software and technical skills that were left over from Y2K investments.***

***The Y2K effect***

Internal recovery traces its origins back to the period immediately following the end of Y2K. One of the many technology benefits of Y2K was the maturation of data replication solutions. Synchronous replication tools such as IBM's Peer to Peer Remote Copy (PPRC) and EMC's Symmetrix Remote Data Facility (SRDF) allowed users to create duplicate Direct Access Storage Device (DASD) environments quickly and efficiently in their own data centers. With the close of the Y2K event, customers were left with significant investments in redundant hardware, software, and technical skills with a timely need to utilize these resources for disaster recovery.

The initial application of this technology was relatively straight forward and mainframe centric. A second data center located within approximately 25 miles, the effective distance for synchronous replication, was created or acquired. The second mainframe processor was acquired, and data was replicated synchronously between the two sites. Since these solutions tended to be mainframe only, a second lights-out data center provided a cost effective means of backing up a mainframe environment. At that time, growth of disk and mainframe processing was relatively stable and well suited for this kind of solution.

Considering synchronous data replication was the only technology available during this time that ensured minimum data loss, recovery sites were required to be within synchronous distance, usually less than 25 miles. Internal mainframe recovery was the only means to meet this objective. This strategy offered a good solution for mainframe recovery and data protection, and met the regulatory requirements of that time frame.

It also provided a solid foundation for the two-site strategy combining both high availability and disaster recovery. These two sites allowed for split operations or failover capability, as well as protection against a facility or component level failure.

---

---

Highlights

---

---

***Explosive growth in using open systems and Wintel servers has created new challenges for IT environments.***

***Fall of distributed systems***

As companies were in the process of implementing internal recovery strategies for their mainframe environment, several significant changes were occurring in the world of technology. As open systems and Wintel server's experienced tremendous growth, companies began to recognize this growth, and then the migration of these platforms into data centers began. However, unlike the mainframe that had a long history of strong controls with processes in place, open systems and Wintel environments lacked established operating disciplines and processes. Process and procedures were developed at the platform level and lacked the formal, tested structure of the mainframe. Back-up and recovery were disciplines that did not exist or were developed by individual users. Backup was often designed at a platform level, requiring many different tools and procedures. This lack of operational discipline, explosive growth, and integration into the critical production environment created a rapidly moving challenge from a recovery perspective.

The conditions that enabled companies to easily implement a second-site recovery strategy for their mainframe environment were completely absent for the open systems arena. Even as companies were successfully completing the implementation of second-site mainframe-only strategies, the world was rapidly moving to a different level.

***New challenges***

In the years between 2005 and today, the IT industry witnessed more systemic changes to the IT environment than in the previous 40 years combined. Open systems and Wintel servers became critical to the business process and completed their journey into the data center. The explosive growth of these systems and storage continued. The integration of these systems with the mainframe has become so pervasive that uncoupling applications is increasingly more difficult and in many cases almost impossible to accomplish.

---

---

**Highlights**

---

---

Business recovery replaced mainframe recovery. This meant not only recovery of data to the point of failure for all systems, but in more and more instances, recovery of a large number of platforms in the data center.

Secondary data centers that were designed to support mainframe-centric recovery are stressed to meet the infrastructure demands of today's dense-packed server farm. Regulators have moved aggressively to require out-of-region recovery capabilities and increased the focus on data protection. They are also introducing the concept of alternative staffing into the recovery equations.

Due to the expense of comprehensive recovery programs, the exposure of moving data off-site and business requirements integrating high availability and disaster recovery, executive accountability is now more clearly defined than ever before. Disaster recovery which previously accounted for one to two percent of the overall data center budget began to grow into the range of eight to ten percent. No longer can a company simply report that they have conducted an annual hot site exercise. Disaster recovery and data protection has become a strategic component of all data center operations.

***New suite of enablers***

***Advances in technology provide tools that increase efficiency and decrease the cost for recovery.***

Fortunately, technology has also evolved, providing new and exciting options. Enterprise-class storage devices emerged providing greater management and ease for data replication. Asynchronous replication has matured, enabling us to move data out-of-region routinely with only seconds of data loss. Bandwidth capacity has increased, while prices have fallen. Tools for system configuration have improved, eliminating many of the time consuming challenges and complexities of recovering open systems and servers.

---

---

**Highlights**

---

---

Hot site services that previously provided customer's access to data and system recovery equipment for their employees have evolved as well. Today, vendors provide comprehensive coverage where the vendor's staff assumes the responsibility for loading data and operating systems. Operational staff is also provided, eliminating the need for many customers to even send staff to the recovery site.

As data protection has become the driving factor, the cost of recovery has increased dramatically. Data replication has replaced tape-based recovery for companies looking to reduce their recovery time to within less than 24 hours. The cost of the hardware, software and network to support replication are now a significant part of any data center's production budget. The growth of data farms and the need to protect data have moved data management to a strategic imperative as opposed to a regulatory nuisance.

***Current estimates project that data centers will increase power utilization by more than 20 percent a year over the next five years.***

***Internal recovery considerations***

Evaluating the real cost of internal recovery represents a critical step in the decision making process. As the very nature of the primary data center evolves and changes, it will impact the secondary site equally. The continued growth and critical importance of open systems must be factored. The enormous power and cooling requirements experienced with the primary site will impact the secondary site as well. In addition, the uncertainty of the cost of power will represent one of the biggest concerns in the coming years. Current estimates project that data centers will increase power utilization by more than 20 percent a year over the next five years, a trend often overlooked in developing cost cases.

In the past, inaccurate assumptions and missed requirements have lead companies to drastically underestimate the cost of building and maintaining a second data center for recovery.

---

Highlights

---

***If you are considering internal recovery solutions for your organization, carefully weigh the benefits against the commitments required.***

***Critical strategy considerations***

Consider the following points when deciding if internal recovery is right for your organization:

- Am I able to achieve the appropriate level of high availability in a single data center? The improvements in raid protection, system virtualization and facility hardening have dramatically improved overall availability. Does a dual site strategy still make sense both technically and financially?
- How can I accomplish an out-of-region recovery strategy? Generally, out-of-region is construed as being a minimum of 100 miles. Current trends and regulatory direction are requiring out-of-region recovery capability.
- Can my management team commit to the expense of maintaining parallel centers, even during times of economic downturn?
- Does my company's workforce possess the commitment and energy to support multiple data centers?
- Are our business and operating processes mature enough to support this strategy?
- Do our data centers have the capacity and expandability to support this strategy?
- Do I have sufficient access to necessary information to make a 10 year strategic commitment?

***Direct expenses***

Direct expenses are those related exclusively to building and maintaining a second out-of-region data center. The major items are:

- Acquisition facility cost – The initial cost of acquiring or fit out of a data center. In addition to the actual cost of the data center, there is a growing expense associated with the fit out of the site to bring it up to the power and cooling standards now required. Data centers built as little as two years ago were provisioned to support a power load of 40 to 50 kilowatts per square foot. With the density of open systems previously not found in the secondary data center, the power load now required is in the area of 100 to 150 kilowatts per square foot. Recent industry studies are predicting that power draw will increase to over 200 kilowatts per square foot by 2010. While most analysis capture the initial cost and fit out expense, very few are including the incremental cost associated with increasing the power, distribution, and back-up systems, as well as the heating ventilating and air conditioning (HVAC) systems required to support these densely packed facilities. Out-year expenses should be carefully analyzed and anticipated. If a facility is being depreciated, depreciation expenses should be included in the analysis.
- Ongoing operating expense – As a result of the increased load referenced above, current predictions are that power will become the most expensive component of running a data center in the near future. Expenses in a normal forecast at current expense levels do not accurately reflected projected costs. Care should be taken in projecting not only the increase of actual kilowatts used, but also include adjustments for both annual increases and periodic price jumps of fuel and power, similar to those present during the first half of 2008.

- Hardware expense – All critical hardware, including open systems and Wintel servers should maintain parallel growth in capacity between the servers on the production floor and the servers at the recovery site. Experience has shown that during the activation and early stages of a recovery, system capacity will be taxed until recovery is achieved. Attempting to recover on lesser capacity machines can cause and extend disruptions. Analysis must include the original cost of the hardware, as well as periodic refreshes of both production and recovery of hardware. To effectively maintain an adequate recovery, upgrades or hardware replacement should occur simultaneously at both the production site and recovery site.
- Labor – Based upon mainframe only recovery, past assumptions often provided no increase in labor expense to support the solution. Unlike mainframes with a single footprint, open systems and Wintel servers require almost continuous care and feeding. If the second site is internally managed, it requires additional dedicated staff to plan, implement, maintain and test. A full time staff may be required to support the second site, considering it will most likely be located several hundred miles from the primary site. Projections should also include routine rate increases in personnel expense.
- Software tools – Managing two sites will require additional tools for inventory control, layout and design and data migration between sites. Additional tools may be required for system configurations and verifications.

- Management personnel – Planning, building, executing and maintaining a second site will require additional management and management support staff. Any time allocated to supporting the second site, including administrative time, should be captured for analysis. These include management personnel, financial support staff, change management, audit and compliance, test and procurement.
- Miscellaneous Expense – The second site may require backup media, documentation, education and training.

***Indirect expenses***

Considerations should be taken to capture indirect expenses as well. These hidden expenses are often not included on the spreadsheet analysis but represent a significant cost to the company. The following provides examples of indirect expenses:

- Lost opportunity – Capital allocated to the acquisition of a new facility and the upgrade of that facility is capital not available for other projects. A comparison of how that money could have been allocated is important to measure the overall impact of that decision on the business.
- Repurposing of staff – During the initial phase of project development and implementation, as well as on an ongoing basis, IT management and staff will devote a significant amount of time and attention to the project. This is time that could have been spent on other business related projects. Similar to the impact of Y2K or the migration to a new data center, implementing an internal recovery center can funnel IT attention away from core business support and into project support for an extended period of time. Considerations must be given to business initiatives that can be delayed or not completed as a result of the strategic decision to create a second site for internal recovery.

## **Summary**

Internal recovery remains a viable, but increasingly more challenging option. The next several years may be particularly volatile as issues such as virtualization, green data centers, and power consumption change the fundamental underpinnings of IT. Companies concerned with making commitments to internal recovery should carefully examine options such as colocation and hybrid solutions offered by traditional recovery vendors. These options allow companies to reduce capital expenditures for both facilities and hardware while achieving similar recovery points and recovery times of six hours or less. With system configuration tools providing even quicker turnaround times for shared platforms, this approach could provide acceptable levels for recovery when combined with a data replication scheme identical to that used with an internal solution.

When paired with a strong strategy for continuous or high availability at the primary data center, the cost effectiveness and flexibility provided by a hot site vendor or colocation provider can provide a more than acceptable recovery alternative.

### **For more information**

To learn more about disaster recovery strategies for internal recovery, visit: **[ibm.com/services/continuity](http://ibm.com/services/continuity)**



© Copyright IBM Corporation 2008

IBM Global Services  
Route 100  
Somers, NY 10589  
U.S.A.

Produced in the United States of America  
09-08  
All Rights Reserved

IBM, the IBM logo, and [ibm.com](http://ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

Other company, product and service names may be trademarks or service marks of others.

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates.