# The need to move toward virtualized and more resilient disaster-recovery architectures

E. Bartholomy
G. Greenlee
M. Sylvia

*Due to growing concerns around natural disasters, information technology (IT) complexity, increasing cyber-attacks, and the sensitivity of financial systems such that corporations may lose millions of dollars per minute if key business processes are not available, corporations are finding the need to develop more resilient disaster-recovery (DR) architectures. For many years, corporations have critical business functions that have relied on tape methods for DR. However, due to pressures from regulatory groups such as the FFIEC (Federal Financial Insurance Examination Council), there is a growing requirement to recover business functions faster than offered by tape solutions. As a result, application owners are challenged with more aggressive recovery-time objectives that necessitate the development of recovery solutions that offer a faster, near-continuous recovery of critical business function. However, moving mission-critical workloads from tape to a near-continuous method can be very expensive. This is true when dealing with legacy, multisite, heterogeneous workloads that have a business process and governance model that prevents workloads from moving to a cloud-computing model. Nevertheless, to offset DR costs, emerging technologies such as cloud computing and virtualization can be used, along with existing underutilized server capacity, to form effective and affordable DR solutions that can accommodate heterogeneous and legacy workloads.*

## Introduction

IBM is constantly being challenged to improve its abilities to recover critical business functions in the event of a catastrophic disaster by utilizing higher levels of resiliency. In response to this challenge, IBM is exploring cost-effective ways to recover business function through innovative disaster-recovery (DR) solutions. Cloud computing is a relatively new model for delivering and consuming information technology (IT) that appears to promise affordable and flexible capacity. Currently, IBM is also assessing which of its DR workloads can more effectively utilize existing virtual capacity within the enterprise or evaluating whether existing DR solution patterns will work within new cloud constructs. However, the mission-critical nature of some production workloads raises the risk associated with cloud deployment, and in some cases, the risk may far outweigh the potential gains [1].

Within IBM, there is excitement around cloud-based solutions and the benefits the company will derive from it. The prospect of moving production and DR workloads to the cloud and reducing hosting costs by 25% is very motivating to management leaders seeking to reduce expenses [2]. However, the cloud model is still evolving, and hesitation exists about adopting cloud solutions because of concerns about service-level agreements (SLAs), scalability, and security. In addition, cloud solutions are often built on standardization and do not easily accommodate the large heterogeneous IT environments one typically finds within mature companies. Many large enterprises such as IBM can have hundreds or even thousands of applications that are

Digital Object Identifier: 10.1147/JRD.2013.2258759

deployed on open standard platforms such as IBM Power Systems* (UNIX**, Linux**, and Windows**/Intel systems) and a large cadre of appliances. Also, enterprises may have many legacy applications on back-level operating-system platforms and middleware that may not easily port to the cloud.

Given the attention and marketing around cloud technologies and other new replication techniques, there is a tendency to believe that DR can simply be addressed with technology alone. However, in larger enterprises with complex business solutions, in order to be effective and affordable, they must consider other aspects to solve their business-continuity and disaster-recovery (BCDR) needs. Companies should consider overseeing governance processes that classify application criticality, support consistent infrastructure building blocks, and promote design patterns that utilize virtualization and cloud-like capabilities.

## Current DR Landscape

Technology has become more integral to the operations of a company and the ability of a company to conduct business. For most companies today, if their enterprise systems and related business function become unavailable for extended periods, so does the ability to sell products, maintain customer relationships, ship products, and report earnings. Given the importance of critical business functions, companies need to plan for potential outages, including the following: increased end-user expectations of availability as well as business partners who demand high availability (24 hours a day, 7 days a week), pressure from financial and regulatory agencies for timely financial reporting and financial systems with operational integrity, and the need to address more frequent weather and natural disaster events such as the 2011 Fukushima earthquake. Depending on the size of the company and the length of the outage, a disruption could cost thousands or millions of dollars per minute, and this cost does not even include damages to reputation and brand image [3].

Most large companies understand that they cannot afford to ignore BCDR planning for an outage or disaster. Many companies use the terms *business continuity* (BC) and *disaster recovery* interchangeably. While BC and DR have overlapping elements, we should note that DR is a part of BC as it relates to recovery of IT systems, whereas BC deals with availability and recovery of the overall business functions in an effort to avoid an outage. Our primary objective within IBM is on maximum tolerable downtime (MTD), which is the business recovery time that relates to the question, "How long can the business withstand an outage of a given business function when a disastrous and disruptive event occurs?"

Business function and the financial impact to the company will dictate how a company plans for a particular disaster. Banks, credit card companies, and Wall Street brokerages employ BC planning to ensure their mission-critical business function is always available [4]. At the same time, there may be business functions a financial institution relies on that are not considered mission critical, such as an internally deployed social media tool that therefore may not require high availability. Like most large companies, IBM has a mixture of critical and non-critical business functions, along with mission-critical (hypercritical) business functions on which the financial and regulatory institutions rely. It would be unreasonable to treat all data and business functions the same and to insist that all functions (including social media functions) have a BC plan that requires it to be continuously available. No company would be expected to make a duplicate copy of the data of the entire corporation and have it continuously available—especially in the age of Big Data, in which companies generate terabytes of data at various rates.

In order to address the unrealistic prospect of backing up all pieces of data and applications, many companies classify their business function and data into categories based on enterprise governance and policies. As already suggested, data is generally categorized as follows: mission critical/hypercritical, business critical, and noncritical [5]. The "criticality" of the data or business function will dictate the availability requirement for the data during a disaster. Some data may need to be continuously available, whereas other data may need to be recovered over an extended period (hours or days). It might even be acceptable for some data to be permanently lost (e.g., certain data logs and archives). In other words, mission-critical business functions may require a "continuous" solution, whereas noncritical business functions could tolerate a slower "tape-based" solution.

There is a tendency for IT professionals to underestimate the challenge of BCDR. Many technologists assume they can make all business functions highly available by using extensive redundancy in the infrastructure. This thinking may be acceptable for a small business, but the solution is more challenging for enterprises that have rather large, complex business functions, which may span multiple data centers, with various costs of implementing BC via infrastructure, people, and processes. As a result, there seems to be a dichotomy when considering BCDR solutions in large enterprises. For example, data centers may consider either tape-based backup solutions that cannot deliver quick recovery times or highly resilient solutions that are very expensive. As mentioned, the need or requirement for quicker recovery times during a disaster event does not mean all critical business functions must migrate to a near-continuous or fault-tolerant solution. Instead, enterprises must have a method to determine the criticality of each business function and then move that workload to a DR solution with an appropriate recovery-time objective (RTO). Depending on the value that the enterprise places on its

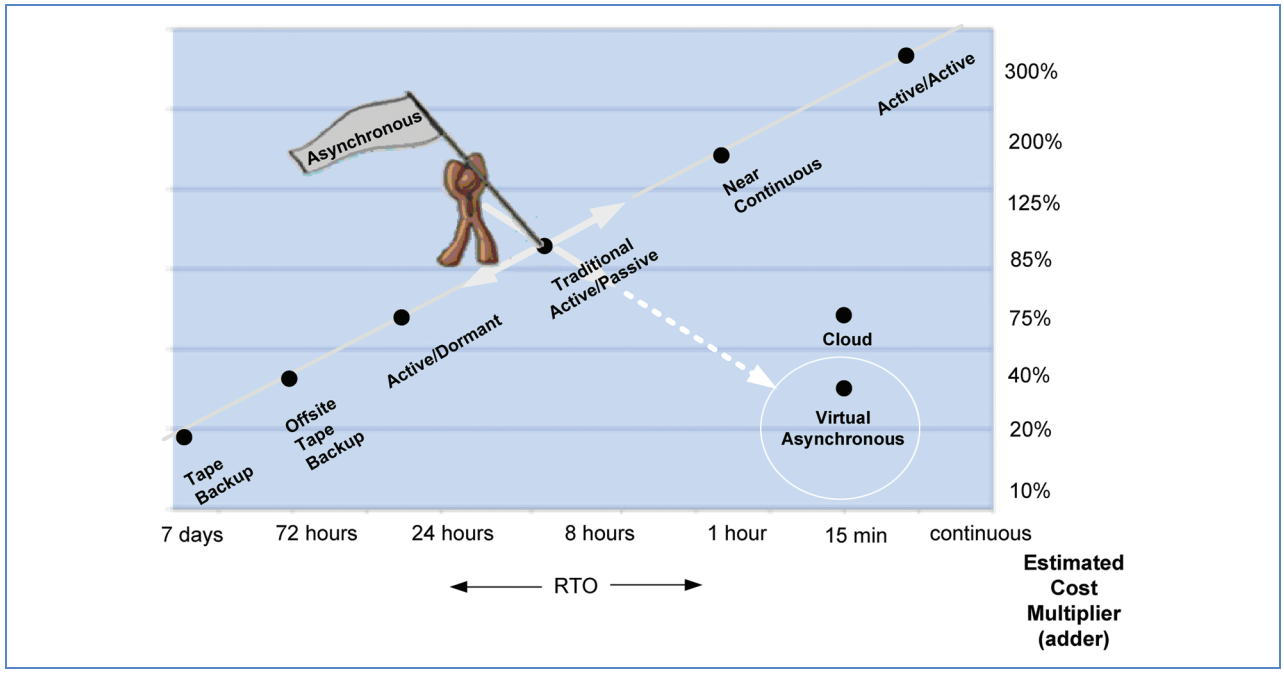


Figure 1

Disaster recovery cost/benefit model. (RTO: recovery-time objective.)

business functions, the company could implement a combination of solutions with various recovery and cost characteristics.

**Figure 1** shows the DR cost-benefit model. Traditionally, the BCDR cost models are linear, where the most rudimentary technique such as tape-based backup has one of the longest RTOs (see the bottom left of Figure 1).

At the opposite end of the scale, the most effective solutions, referred to as Active/Active (upper right of Figure 1), have one of the shortest RTOs. Notice that the better the solution and improved recovery times, the higher the cost. *Active/Active* generally refers to a two-site recovery solution, and more particularly, it is the term used for a highly resilient DR solution in which an application can exist in two geographically dispersed locations and in which user transactions are processed simultaneously. The approach can be extended to multiple sites.

In Figure 1, the person with a flag indicates where, within IBM, the Active/Passive solution (defined below) has been the traditional method for providing an effective RTO with a reasonable cost. However, there is now a new breed of solutions that no longer aligns with the traditional cost model because of dramatically different designs that use virtualization and cloud-like capabilities (dotted line). These technologies fall outside the linear relationship; these DR solutions have improved recovery times at a fraction of the cost of traditional DR.

The Active/Active solution can cost an estimated 8.6 times the standard IT service [6]. Within IBM we call that additional cost an "adder," which is actually a multiplier that we use to determine the cost of a particular solution. Active/Passive and Active/Dormant are examples of an asynchronous solution in which an active application resides in the primary site and can fail over to a standby/nonactive application that resides in a secondary site. This paper provides more detail in later sections regarding these solutions. Asynchronous solutions typically double the cost, because a copy needs to be made in a secondary site, and the cost would be a 100% adder (in addition to the current production costs). As shown in Figure 1, a number of DR solution approaches are available, depending on the criticality of the business function and how quickly the process must be recovered. Considering how high the costs are to ensure proper resiliency, it becomes obvious why it is so important to properly classify the criticality of the business functions of a company.

## Defining appropriate recovery objectives

As suggested, in order for an enterprise to properly determine the criticality of their business functions, it is beneficial to have a well-defined governance model. Corporate Governance includes corporate instructions and policies, along with business function resiliency classifications, an enterprise process framework, and ongoing business risk and

threat assessment. Under the enterprise process framework of IBM, the business operations are formally segmented into 12 different business processes that are then divided among three process categories: operating, enabling, and supporting. The entire internal application portfolio of IBM is aligned with the IBM enterprise process framework. The business owners of each process can use threat assessment and Business Impact Assessment (BIA) tools to determine the criticality of the business function. This evaluation is performed annually, because business function criticality may change from year to year. In order to determine how to mitigate the exposure of a disaster to a given business process, it is important to determine how long the business can tolerate it being unavailable. The enterprise must determine the business criticality of the function (mission critical, critical, noncritical, and other) by assigning a BCV (business criticality value). Using that criticality value, and through corporate policy, the business can dictate how quickly that business function must be recovered in the event of an outage to the applications supporting it. As suggested, the time it takes to recover a business process is commonly known as the RTO.

Within IBM, the actual DR solution is secondary to the classification of the business function. The solution mitigation strategy, or solution, is typically the last step of a carefully thought-out planning effort. As an example, if an enterprise has 2,000 applications, and 25% of them are deemed "business critical" and have a need for a fault-tolerant solution, it would be reasonable to develop a highly available solution only for those 500 applications—not the entire portfolio. Furthermore, the business may choose to build a three-site Active/Active/Active (continuous) solution for a subset of those 500 applications that they may term "hypercritical." Alternatively, they may choose to develop a less expensive solution, such as an Active/Active or near-continuous solution, for the remaining applications. Based on the criticality value of any given enterprise process (business function), along with any specific functional and performance requirements, IBM can guide the application to a well-suited architectural solution.

## Developing BCDR solutions based on standard patterns and technology

All of the IBM solution patterns for BC and DR are offerings that support the internal application portfolio. There is no one pattern that addresses the many different technology requirements, RTO requirements, and recovery-point objectives (RPO). An RPO is the maximum tolerable period in which data might be lost from an IT service due to a major incident [7]. When an application team or related business function team begins the planning process to address its BCDR requirements, the business validates the assigned classification as a part of its architecture. This validation, as part of the governance process of IBM, ensures the solution is neither underengineered nor overengineered. As an application enters the development process, a governance council or Architectural Review Council (ARC) recommends which solution pattern should be followed. For instance, for a mission-critical web application with a very high business criticality value, which requires fault-tolerant (continuous) execution, the council would recommend hosting the application within the Events Infrastructure environment of IBM—a premier three-site Active/Active/Active continuously available environment, which is a variation of Active/Active. Alternatively, if the application has a slightly lower criticality value, a near-continuous solution may be recommended within the strategic Application Hosting Environment (AHE) of IBM. Each hosting environment has its own unique architecture and requirements that the application must follow. The IBM standards and guidelines specify the hosting requirements that must be followed to ensure the application achieves its desired availability levels, so that it can be supported in the most cost-effective manner.

IBM offers a limited number of resilience solution patterns depending on required RTOs and RPOs. A pattern describes, in detail, a tested solution to a commonly recurring problem, addressing the objectives needed to be achieved. The specification of a pattern describes the problem that is being addressed, why the problem is important (the value statement), and any constraints for the solution. Patterns typically emerge from common usage and the application of a particular product or technology [8]. Furthermore, all of the IBM patterns consist of reusable technology called technology building blocks (TBBs), which are well-defined strategic enablers that make up the IBM global IT infrastructure. There are approximately 60 building blocks that make up a typical hosting center, and any one of the IBM BCDR solutions would use many of these preexisting components.

**Figure 2** shows some of the common IBM TBBs as they relate to the infrastructure of a hosting solution. If a web application deploys an asynchronous solution within two of five strategic IBM AHE data centers, for example, between Poughkeepsie (New York) and Boulder (Colorado), the application will be deployed and supported uniformly, because both data centers are built from the same uniform building blocks. Both web applications and mainframe applications are built upon these strategic and foundational capabilities so that when any one of the building blocks changes in terms of security patches or new functionality or technology refreshes, all of the BCDR solutions will be updated as a result. One of our strongest capabilities is the IBM Global Delivery Framework (GDF), which is an organizational construct that is used to support both the technical environments and the TBBs. The GDF team can work remotely at multiple delivery centers where they use

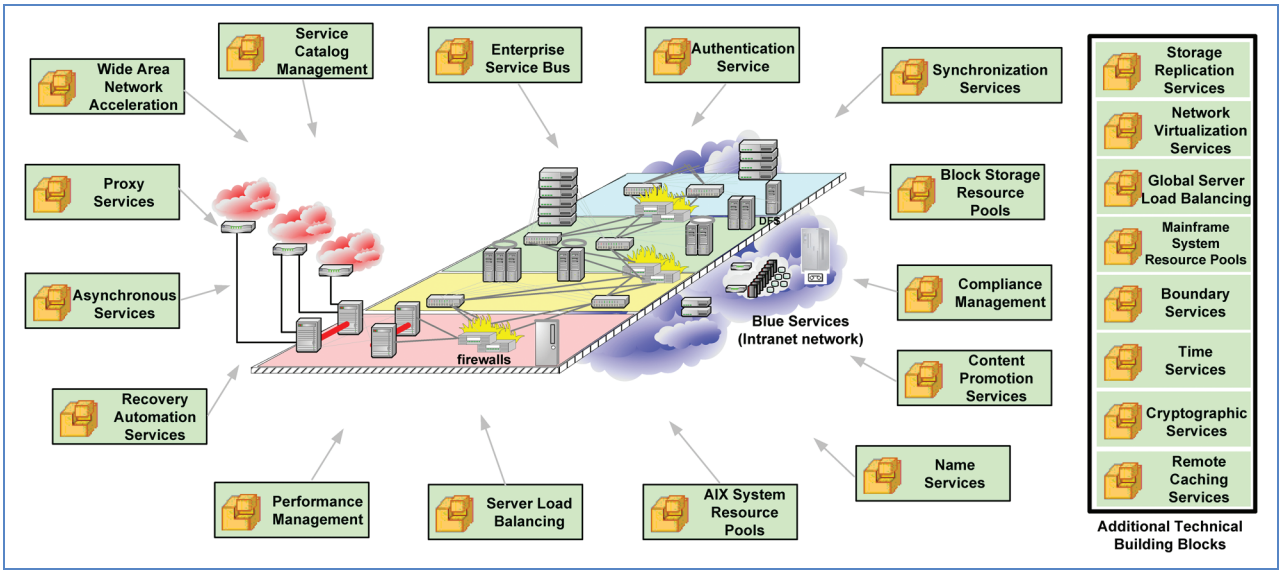


Figure 2

The IBM common technology building blocks (TBBs).

consistent tools, processes, and procedures. For example, a GDF in Dubuque, Iowa, can support data centers in Boulder and Poughkeepsie, using common tools and processes. Anytime the primary site is updated, the secondary site is managed consistently, which enforces the requirement for well-governed environments. This governance will help prevent or even altogether eliminate configuration "drift" between the environments. In addition, because the support teams are geographically dispersed, they perform work at any location; the arrangement helps mitigate pandemic concerns.

As previously mentioned, within IBM there is a tendency for both mainframe and web applications to migrate toward either end of the BCDR spectrum: tape based or continuous (e.g., see Figure 1). However, two new trends are developing within the BCDR space. First, the IT industry expects improved recovery times that tape can no longer deliver, because recovering from a tape solution can take between 24 and 72 hours on average [9]. Second, with the advent of cloud and cloud-like capabilities, IBM can now offer lower-cost solutions that deliver much quicker recovery times. IBM is one of the pioneers with respect to providing continuous availability for web applications within its continuously available Events Infrastructure hosting environment, but the cost of this can be more than double standard "internal hosting" costs. On the other end of the spectrum, the tape-based solutions are the most affordable option, with data being backed up to tape and sent offsite. When a disaster occurs, the data can be sent to an alternate recovery facility where the application is restored on reserved hardware. A typical offsite tape backup solution would have approximately a 25% to 40% adder depending on the number and capacity of servers and the amount of storage [10].

IBM has offered a standard asynchronous DR pattern for both web applications and mainframe applications for more than a decade. The demand for this particular pattern has gradually increased because of pressure for certain business functions to improve their recovery time. The biggest inhibitor to adoption is the cost of our asynchronous solutions that typically have a 100% to 125% adder. What we have learned is that most of our internal application owners would prefer to pay no more than a 25% adder—40% as the worst case. Any cost beyond 40% would be too high, and the application owner would likely adopt a less-expensive solution such as offsite tape backup. The attitude of our internal application community is no different from that of most enterprises where the planning for BCDR is executed to the extent that it makes financial sense. Ultimately, that analysis must consider the cost of planning versus the cost of failure [11].

## Lowering costs through virtualization

In anticipation of the demand for improved asynchronous DR solutions with lower RTOs, the IBM Advanced Technology Group within the office of the IBM CIO (Chief Information Officer) has developed a number of enhancements to its traditional asynchronous DR patterns—enhancements that bring the adder cost down below 40% (approximately 33% based on internal cost structure; see Figure 1). The DR capabilities have not

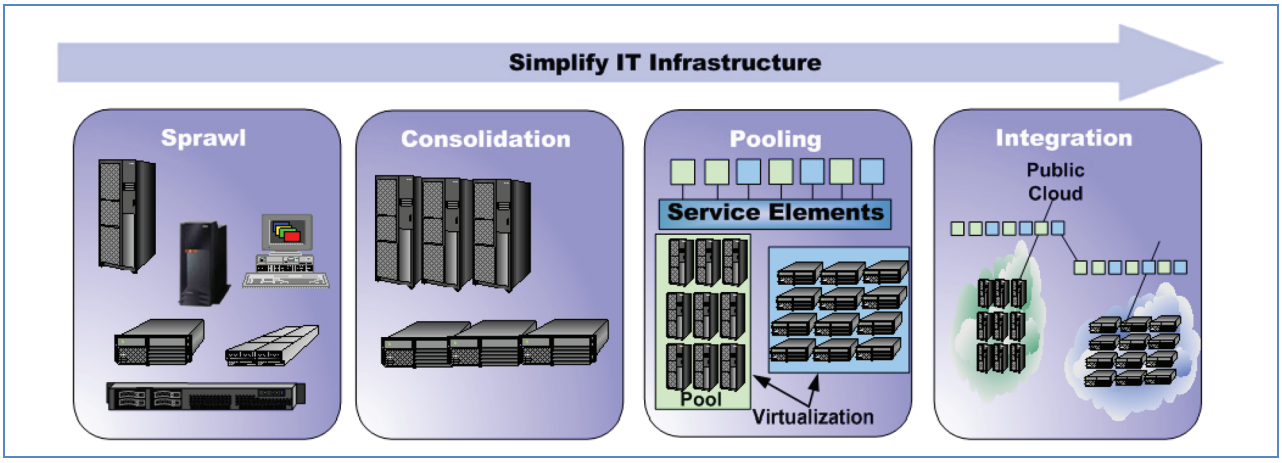


Figure 3

IT infrastructure transformation, from sprawl to integration. The Service Elements in the third block represent pooled services and/or IT components. Once resource elements, such as technology building blocks (TBBs), are pooled, they can be offered as services and used by multiple users, systems, or other services.

changed much throughout the years except for the use of virtualized server capacity. IBM began using virtualized capacity for DR long before the ubiquity of cloud system when, in 2003, IBM began to offer the On Demand Workplace.

**Figure 3** shows the IT infrastructure transformation, from sprawl to integration. Throughout the past decade, IBM has been steadily transforming its IT infrastructure to a point that it has almost all the properties of a cloud. Soon after the "sprawl" phase from the client-server period, when IBM struggled to manage all of the various web servers under the desks of its employees, IBM implemented the Global Web Architecture (GWA) organization to address architectural and support inconsistencies by consolidating servers and support. GWA eventually evolved into an AHE, where resources were pooled to support the "on-demand" hosting environment of IBM. It required more than 15 years to make this transition to consolidating and promoting a homogeneous platform with a uniform "look and feel" based on a standardized offering, which has significantly reduced the overall hosting costs of IBM. Over the past decade, IBM has also taken advantage of virtualization in order to offer resource pools for storage area networks (SANs), server pools of IBM System z* (MVS) resources, Power Systems (AIX*), and System x* (Windows/Intel). The pooled resources of SAN and virtual server capacity have become their own TBBs that applications can consume, and these resources are delivered consistently across the IBM Global Account worldwide. Any application requiring new server capacity must make use of these existing pools per our corporate CIO hosting standards. IBM has acquired and developed a number of provisioning technologies, and when integrated with virtualization and other TBBs, IBM is able to implement new technical solutions, including our new cloud offering called SmartCloud* Enterprise+ (SCE+). It is important to note that although IBM has cloud technology offerings, not all of its internal hosting environments are technically clouds. The National Institute of Standards and Technology (NIST) defines a cloud as having five essential characteristics: on-demand capabilities and self-service, broad network access, resource pooling, rapid elasticity or expansion, and measured service [12].

By design, our internal hosting model does not currently provide self-service provisioning for production servers. Although, we have highly automated processes to provision capacities, IBM still has a number of controls in place from a governance standpoint to ensure that resources are prioritized so that IBM can ensure a certain level of availability, security, and performance. The cost of implementing a provisioning system with all the required enterprise governance controls would be exorbitant while providing very little benefit to the enterprise. When the internal hosting approach of IBM is considered with respect to Figure 3, the hosting may be considered to be in the integration phase, and it is not fully integrated as is our SCE+ cloud environment. IBM internal operations use private cloud-like capabilities (at right of Figure 3) to keep the overall cost of deployment and hosting to a minimum, with the added benefit that our hosting environments are highly available, well governed, and secure.

IBM can now deliver an Active/Passive solution (and defined more fully in the following section), as shown in Figure 1, with a recovery time of less than 1 hour, with a lower overall adder cost for asynchronous DR patterns by

making use of virtualization capabilities that exist within our AHE internal hosting centers across the globe. Virtualization technology enables IBM to "crush" (or "deflate") DR workloads in the recovery site, which can then be inflated during a disaster event. The terms *crush* and *deflate* describe virtualization techniques used to temporarily de-allocate central processing unit (CPU) resources and memory resources. With integration tools and monitoring, IBM can reserve a large LPAR (logical partition/virtualized server) and then crush the resources to a bare minimum and charge the DR application owner for the use of a small or extra-small server (which is one fourth or one third of the cost of a typical large server). Instead of charging an application owner a 100% to 125% adder, the application owner would only be charged approximately a 33% adder. During a disaster event, the solution allows the DR workload to inflate, or expand, without affecting the availability of other mission-critical and critical workloads. Finally, if the application business-criticality value changes to a point where the application needs to have a near-continuous Active/Active DR solution, it can be transformed with a few simple configuration changes—assuming the application has the appropriate architecture.

## IBM improved Active/Dormant and Active/Passive patterns

IBM is in the process of transforming its internal DR approach to better balance today's business risk mitigation requirements with affordability and effectiveness. The asynchronous Active/Dormant and Active/Passive patterns address the need for a strategic multisite resilient solution for critical and vital business applications that are hosted in a standardized offering such as the AHE. These asynchronous solutions are DR techniques that use highly automated best-of-breed asynchronous replication to copy a mirror image of the application and its execution environment (e.g., operating system and middleware) into a secondary recovery site. Although this paper primarily focuses on the Active/Passive solution, it is important to understand that the primary difference between Active/Dormant and Active/Passive solutions is that with Active/Passive, the secondary site is "alive" (i.e., operational) but consuming a minimum amount of capacity. The Active/Passive solution employs replication technologies such as high-availability DR (HADR) message queue replication (MQ rep). The Active/Dormant solution relies on storage replication. The Active/Dormant solution does have images allocated in terms of LPAR profiles (i.e., IBM Power Systems); however, the LPARs are not actually alive or online because they exist as LPAR profiles (predefined execution footprints). As a result, Active/Dormant LPARs do not have a requirement to be patched or monitored, and any change that takes place in the primary site automatically gets replicated to the secondary site, which therefore reduces the cost. However, the Active/Dormant solution requires that a skilled team be available to break storage replication mirrors and instantiate (bring online) LPAR profiles that can require additional precious time and require expensive highly trained technicians when there is a DR event.

Both the Active/Dormant and Active/Passive patterns use existing underutilized virtual capacity that typically can be found on many of the servers. On average, server utilization within the IT industry ranges from 10% to 15% [13]. The workload in these server environments fluctuates and naturally leaves underutilized server capacity, which in turn can be capitalized for DR activity. Server resources and capacity needed to host the recovered application in the secondary site are identified and reserved through the use of a proprietary reservation tool based on Tivoli* products. The deflated DR LPARs at the secondary site, where an image is running with reduced resources, are allowed to run with reduced capacity until a disaster is called at the primary site, thereby keeping utilization and costs low. The dormant secondary site recovery solution is sized according to the full production capacity and criticality value of an application. However, for the deflated image where the recovery LPARs are running in a non-DR mode, the server will be consuming only 10% to 30% of its full capacity in order to run replication in the background. Tivoli brand tooling and virtualization techniques will ensure that the LPARs in the recovery site never exceed a *small allocation* during non-DR operations. When a DR event is declared, the LPAR is inflated to 100% of its entitled capacity, and in turn the application owner's bill is adjusted accordingly.

**Figure 4** is a graphical depiction of the DR workload of an application as applied to a primary and secondary site, before and after a recovery event. The active production servers are depicted by the orange servers on the left in Figure 4 and are to be replicated to a secondary site and stored in a dormant (or passive) state until there is a DR event. The dormant deflated servers are represented by the smaller orange server on the right. The deflated LPARs will run at approximately 10% to 30%. This is accomplished through virtualization configuration parameters relating to crushing (deflating), fencing memory, I/O (input/output), and CPU capacity to allow sufficient capacity to replicate data, code, and patches.

In the event of a DR for a failing application, workload at the secondary site is transformed to accommodate the failing application with "five methods of accommodation." When the dormant LPARs in the secondary site inflate, the necessary capacity is attempted to be attained and afforded by each of the following five levels:

1. Available unused capacity is used. This unused capacity is referred to as white space or headroom capacity and is reserved for growth or rare usage spikes. In reality, most LPARs are rarely if ever used above 50% capacity.

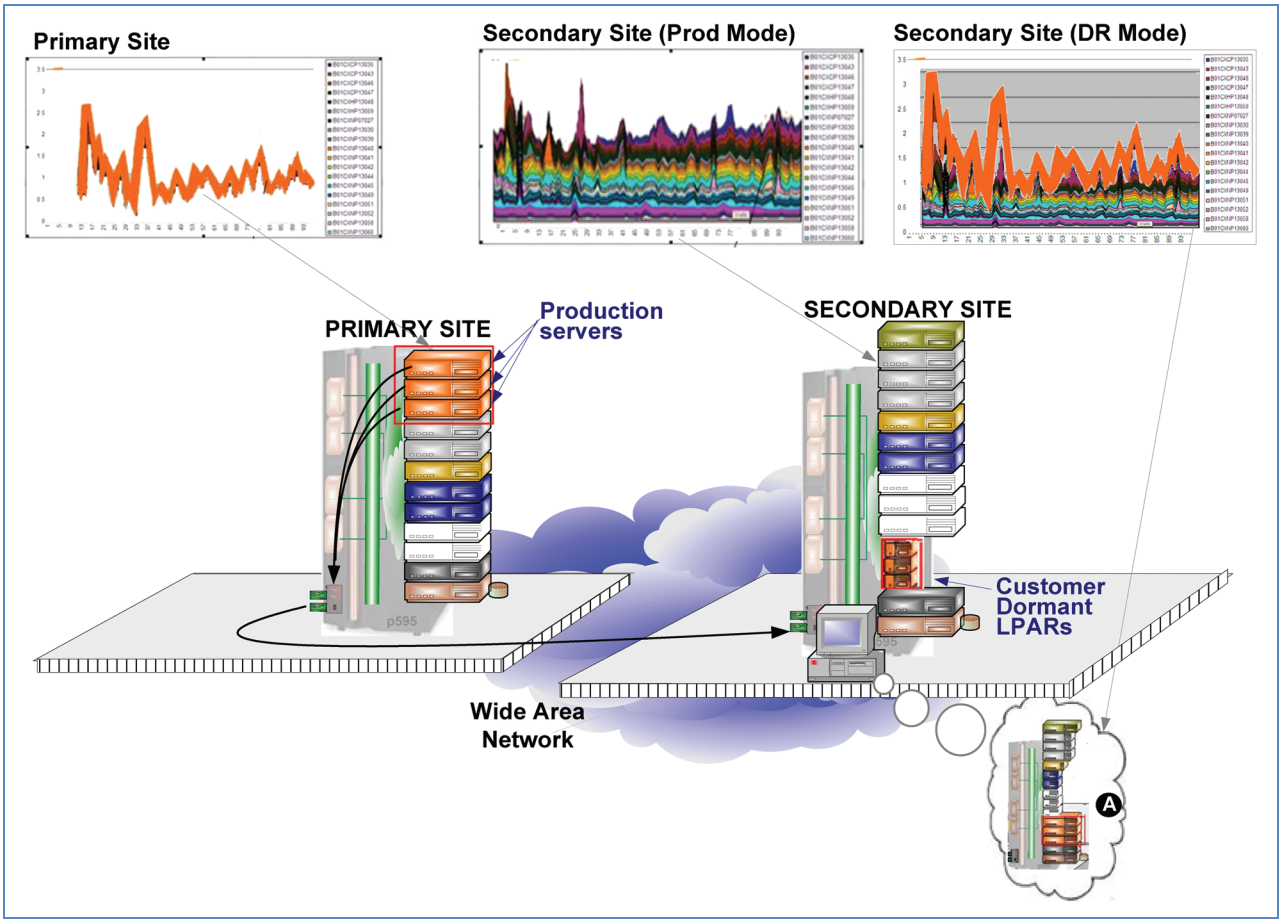


Figure 4

Application workload, as it applies to a primary and secondary site, before and after a disaster-recovery (DR) event. A wide area network connects the primary site and secondary site. (Prod: production.)

The dormant DR LPARs at the secondary site will inflate to the entitled capacity of the application.

2. In the event there is insufficient unused capacity, the system assesses the opportunity to crush or deflate (i.e., reduce the server resource allocations) the lower-priority and/or less-critical application workloads on the basis of their business critical value (BCV) classifications. The recovered capacity is reallocated to critical workloads as well as the DR LPARs.
3. If the frame containing DR LPARs is still found to be constrained, the system assesses the opportunity for shutting down lower-priority and/or less-critical application workloads on the basis of their BCV classifications. Again, the recovered capacity is reallocated to critical workloads as well as the DR LPARs.
4. If the frame is still found to be constrained, the system assesses the opportunity to reduce the recovered server resource allocations of the application. The recovered capacity is reallocated to other critical workloads, permitting the critical and recovered applications to run in an acceptable manner. However, the noncritical applications will have reduced capacity and may result in degraded performance.
5. If, after the previous measures have been executed, the recovered application is still unable to run in a reasonable manner at the secondary site, noncritical applications will have to be shut down, suffer an outage, and possibly be recovered on different server(s) and/or at different locations.

Note that as the system and/or computers are more constrained, it is beneficial to have a mixture of both critical and noncritical workloads. This works to the overall favor of a methodology because there is an adequate reserve of noncritical workloads that can be crushed.

To further understand the asynchronous Active/Passive solution, we consider an example and refer to Figure 4.

In our example, we assume that we have an IBM application hosted in a standardized environment in AHE. This particular topology has three LPARs (or servers): two clustered WebSphere* servers and one DB2* server.

As a part of the boarding process, our Tivoli tooling captures key utilization metrics of each of the servers in the primary site and develops a workload baseline through time in order to establish a trend. At the same time, Tivoli tooling will continuously trend the utilization of virtualized capacity at the secondary (DR) site. The trending can be seen by the graphs at the top left of the Figure 4. The orange line in the graph on the left shows the application utilization during a period of time. Also in this example, Tivoli tooling is trending all of the LPARs on a given frame within the secondary site and can provide information about the amount of available capacity for new applications that include DR LPARs, which is shown by the graph at the top center of Figure 4.

The system can forecast available unused capacity on the basis of how various LPARs are deflated given their criticality (see label A in Figure 4). This tiering mechanism will provide additional available capacity by taking the capacity from the deflated noncritical LPARs and reallocating the capacity to the more vital LPARs, including the DR servers. Once the trending analysis has been completed on the production workload, the tool can project the workload on top of the adjusted (tiered) DR capacity, if needed, in the secondary site. The graph at the top right projects the utilization load of a server in the secondary site with the application DR LPARs running, represented in orange, with the deflation of the non-vital workload.

The tooling functions help the hosting centers plan and adjust DR allocations in any potential secondary site. In this example, there is sufficient unused capacity in the secondary site to introduce production workloads from the destroyed primary site in the event of a disaster. This tooling can also be used to help manage workloads and allow the delivery teams to move workloads to less busy servers.

## Disaster recovery steady state and invocation

When the production application is copied over to the secondary site, both sites would be capable of running the application simultaneously. However, because the host name in the DNS (domain name system) points to the primary site, users would not be able to log into the secondary site [via our single sign-on (SSO) technology], and partner applications would not be able to transmit updates and feeds. Because both sites would have their own unique cluster and/or IP address and host name, with a simple change to Global Server Load Balancing and SSO, it allows the Active/Passive solution to transform to an Active/Active solution where traffic can be distributed between sites, assuming the application is architected, or designed, for multiple sites. The solution pattern does address partner connections and communications to other upstream and downstream applications. The corporate governance of IBM suggests that applications communicate to other applications via web services over a service-oriented architecture (SOA) enterprise service bus. Assuming the partner feeds flow over the enterprise service bus in the event of a disaster, upstream and downstream feeds will be redirected as well.

After the deployment of the DR application has been tested and accepted by the application owner, all access to the application would be turned off and the LPARs are then deflated to minimum entitlement. However, ongoing replication updates occur in the background to ensure changes and updates are made to operating system, middleware, and application code and include security fixes, patches, and customer data. These ongoing updates will run in the background, ensuring that the applications in the primary site and secondary sites are synchronized. This is shown in Figure 4 by the orange LPARs in the secondary site, on the right, which are deflated at this point.

In compliance with IBM corporate instructions, the creation and updating of a DR test plan and the periodical DR testing must occur. In conjunction with the test plan, if there is a catastrophic event and DR failover is required, the hosting offering owner, CIO security, and application team would jointly declare a disaster. The hosting center engages the GDF DR SWAT team to failover the application. At this point, the DR LPARs in the secondary site are still deflated. The frame/servers in the secondary site and all of its LPARs will adjust, if need be, to the tiering policy established by the CIO. As suggested, the noncritical workload will be deflated and the critical workload of the DR LPARs will be inflated (label A at the bottom right in Figure 4). The DR LPARs represented in orange will be inflated to their full entitled capacity. The application continues its application forward recovery, and the DR SWAT team makes necessary failover changes including changes to SSO that allow users access to the secondary site. Likewise, the partner feeds are redirected to the secondary site over the enterprise service bus. A key part of any DR plan is to determine when and where to start building out a new secondary site and how it needs to use the same TBBs and DR solution pattern.

## Conclusion

Through the use of reusable TBBs, standardization, and virtualization, companies can quickly recover heterogeneous workloads after a disaster by using an asynchronous solution as a more cost-effective way than traditional Active/Active or tape-based approaches. In addition, using a governance process to properly classify workloads, the company can determine the most cost-effective way to recover workloads on the basis of the criticality of the business function. The asynchronous solutions outlined in this paper provide the best cost-performance ratio for critical workloads that do

not require a continuous (Active/Active) solution. The governance process described in this paper has an added value of allowing the enterprise to evaluate and assess the criticality and the ability to de-provision noncritical workloads for DR. Not only does this governance model provide a wide variety of DR deployment patterns based on the criticality and RTO values of a given business function, but it provides an intuitive and graduated path to improved recovery times. As described, the architectural patterns are highly configurable, such that TBBs can be assembled in a fashion that best accommodates the unique DR requirements of an application. In addition, if a current business function or application has an elevated business criticality value and is currently deployed as Active/Passive, the solution can become Active/Active with little more than some configuration changes (to global load balancing and the inflation of secondary-site LPARs).

As enterprise governance models change and applications are built to execute in the cloud (or cloud-like) environments (e.g., such as SCE+ or the new IBM PureSystems*), migrating mission-critical functionality will become easier over time. The architectural transformation from legacy applications to cloud-based applications may require a number of years. However, the transformation will provide large enterprises the time for the cloud to mature in terms of availability and security. The approach outlined in this paper will enable companies to quickly recover from disaster events by making use of underutilized server capacity and using cloud-like capabilities across their enterprises, which will significantly reduce their DR costs. Readers who want to become acquainted with further background information in the technical literature on multi-site DR solutions may consult references such as [14] and [15].

*Trademark, service mark, or registered trademark of International Business Machines Corporation in the United States, other countries, or both.

**Trademark, service mark, or registered trademark of The Open Group, Linus Torvalds, or Microsoft Corporation in the United States, other countries, or both.

## References


1. M. Sylvia and B. Peterson, "Success in the cloud: Why workload matters," IBM Global Services, Somers, NY, USA, Mar. 2012, IBM Thought Leadership White Paper. [Online]. Available: http://www-935.ibm.com/services/in/igs/pdf/CIW03082USEN.PDF
2. F. Etro, "The economics of cloud computing," *IUP J. Manage. Econom.*, vol. 9, no. 2, pp. 7–12, 2011.
3. S. Snedaker, *Business Continuity & Disaster Recovery*. Burlington, MA, USA: Syngress, 2007.
4. F. Arduini and V. Morabito, "Business continuity and the banking industry," *Commun. ACM*, vol. 53, no. 3, pp. 121–125, Mar. 2010.
5. M. Mithani, M. Salsburg, and S. Rao, "A decision support system for moving workloads to public clouds," in *Proc. Global Sci. Technol. Forum*, 2010, pp. 7–8.
6. B. Malik and D. Scott, "How to calculate the cost of continuously available IT services," Gartner Inc., Stamford, CA, USA, Gartner Res. Paper–G00163539, Apr. 22, 2009. [Online]. Available: http://www.gartner.com/id=944921
7. T. Wood, E. Cecchet, K. K. Ramakrishnan, P. Shenoy, J. van der Merwey, and A. Venkataramani, "Disaster recovery as a cloud service: economic benefits & deployment challenges," in *Proc. 2nd USENIX Conf. HotCloud Comput.*, 2010, p. 8.
8. IBM Publications, Armonk, NY, USA, ver. v. 7.0.05, Jul. 17, 2007. [Online]. Available: http://publib.boulder.ibm.com/infocenter/wmbhelp/v7r0m0/index.jsp?topic=%2Fcom.ibm.etools.mft.doc%2Fac67850_.htm
9. M. Klein, "How the Cloud Changes Disaster Recovery," *Industry Perspectives*, Data Center Knowledge, Lawrenceville, NJ, USA, Jul. 26, 2011. [Online]. Available: http://www.datacenterknowledge.com/archives/2011/07/26/how-the-cloud-changes-disaster-recovery/
10. H. Burton, "How Much Does Tape Backup Cost," Cloud Direct, Bath, U.K., 2010. [Online]. Available: http://www.backupdirect.net/how-much-does-tape-backup-cost
11. S. Snedarker, *Business Continuity and Disaster Recovery Planning for IT Professionals*. Burlington, MA, USA: Syngress, 2007.
12. P. Mell and T. Grance, "The NIST Definition of Cloud Computing," Nat. Inst. Of Standards and Technol., Gaithersburg, MD, USA, Special Publication 800-145, 2011. [Online]. Available: http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf
13. S. Wexler. (2009, Nov.). Gartner Says Server Replacement Not the Answer. Channel Insider. [Online]. Available: http://www.channelinsider.com/c/a/Spotlight/Gartner-Says-Server-Replacement-Not-the-Answer-248207
14. S. Sengupta and K. M. Annervaz, "Planning for optimal multi-site data distribution for disaster recovery," in *Proc. Econom. Grids, Clouds, Syst., Services, Lect. Notes Comput. Sci.*, 2012, vol. 7150, pp. 161–172.
15. E. Bauer, R. Adams, and D. Eustace, *Beyond Redundancy: How Geographic Redundancy Can Improve Service Availability and Reliability of Computer-Based Systems*. Hoboken, NJ, USA: Wiley, 2012.


*Received November 19, 2012; accepted for publication December 14, 2012*

**Erik Bartholomy** *IBM Office of the CIO, Advanced Technologies, Boulder, CO 80301 USA (erikb@us.ibm.com).* Mr. Bartholomy joined IBM in 1998 and led the National WebSphere Practice until 2001 when he joined the IBM Global Web Architecture team, where he designed disaster-recovery solutions for some of the largest internal applications of IBM. Mr. Bartholomy is a Senior Certified IT Architect and author or coauthor of more than 10 patents.

**Gordan Greenlee** *IBM Global Business Services, Endicott, NY 13760 USA (greenlee@us.ibm.com).* Mr. Greenlee joined IBM in 1984 and has worked in various areas of the business and with a wide variety of technologies. Since the late 1990s, his focus has been on security and identity management services working with the IBM CIO security team on the IBM internal directory (BluePages) and other related infrastructure strategies and solutions. Mr. Greenlee is currently an IBM Master Inventor and a Certified IT Architect.

**Michael Sylvia** *IBM Office of the CIO, Advanced Technologies, Sacramento, CA 95833 USA (msylvia@us.ibm.com).* Mr. Sylvia is a Distinguished Engineer and Director of the Infrastructure Management and Optimization Architecture team in the Office of the IBM CIO. He joined IBM in 1983 and has since performed in various consulting and services, leadership, IT delivery, and technical sales support roles. He has a B.S. degree in computer science from California State University and is an Open Group Certified Architect. In 2000, Mr. Sylvia was named an IBM Distinguished Engineer and was elected to the IBM Academy of Technology.