CLABBY ANALYTICS

Advisory

Business Resiliency in the New Enterprise Data Center

Introduction

In November, 2002, IBM announced "on demand computing" — an initiative that provided information technology (IT) executives with insights and guidance related to the building of information systems infrastructures that could transparently and dynamically support efficient business process flow. On demand computing called for IT strategic planners and designers to:

1. Consolidate and virtualize their information systems;
2. Implement a service-oriented architecture (SOA) to execute service requests;
3. Automate the management of systems, storage, networks, databases, and applications; and, ultimately to
4. Focus on the orchestration of business process flow over virtualized, SOA-enabled information systems.

Now, six years later, IBM has issued further guidance. Based upon changing market conditions (such as greatly increasing energy costs, global expansion, and rising IT management costs), as well as based upon its experience in helping thousands of its customers deploy SOA-enabled infrastructure and consolidated, virtualized systems, IBM is calling for the transformation of the existing distributed computing data center model into what it calls the *"new enterprise data center"* (NEDC) model.

> ***Like on demand computing, this new enterprise data center model focuses strongly on enterprise information systems architecture and resource virtualization. But the new enterprise data center model also focuses on energy efficiency; infrastructure optimization; facilities management; business driven service management; security; and business resiliency.***

This *Advisory* takes a closer look at business resiliency in the new enterprise data center. PART 1 describes what a new enterprise data center is, the problems that it is meant to solve, and three phases of NEDC implementation. PART 2 defines business resiliency and describes four critical needs that business resiliency needs to address. This section also describes how business resiliency maps into the three phases of NEDC. And PART 3 concludes by providing step-by-step advice for IT designers and planners who are looking to implement resilient business services within a new enterprise data center environment.

PART 1: What Is a "New Enterprise Data Center"?

Due to poor design practices, data centers all around the world have become highly inefficient and fragmented. Lax distributed system design practices have made data centers extremely complex and costly to manage, as well as difficult to secure.

To address these issues, IBM has introduced a new model for data center design (NEDC) — a model that focuses on controlling data center costs through more efficient information

systems and data center designs, as well as through better resource utilization, improved program-to-program communications, and more advanced service management practices.

The Problems with Older Data Center Designs
The major problems with older, distributed data center designs include:

- *Access point proliferation* — distributed designs proliferate network access points, creating tens, hundred, or thousands of ports that need to be protected from security intrusions. Scale-up, centralized designs (as recommended when designing new enterprise data centers) can greatly reduce the number of access points — reducing management complexity while mitigating associated risks;
- *Inefficient resource utilization* — Systems and storage in distributed data centers are notoriously under-utilized, resulting in wasted capacity (in many cases, distributed application servers run at 10-20% of total capacity in order to leave headroom for computing peak periods). Improving utilization can lower acquisition costs (IT buyers can get more computing out of existing systems), simplify management, improve availability, and so on;
- *Inefficient power usage* — distributed system constantly burn power. Power supplies waste energy as they convert from DC to AC current. Fans need to continually blow air to cool CPU, memory, and other components. And network interface cards (NICs) and associated network switches, hubs, and routers also burn power — even when not in use. Scale-up designs increase system utilization (resulting in less power waste by having fewer machines consuming electricity while delivering an equivalent amount of computing). Further, scale-up designs significantly reduce the number of NICs and associated network equipment that needs to be powered (because a lot of "distributed computing" takes place over internal backplanes or busses inside a given systems enclosure);
- *Program model complexity* — older data center designs rely heavily on tightly-coupled programming models that are highly-inflexible and difficult to maintain, and that prevent enterprises from capitalizing on efficiencies that can be derived by adopting newer loosely-coupled "service-oriented" programming models; and,
- *Labor-intensive process flow management* — managing the flow of business processes in old-style data centers requires a lot of human intervention, and is therefore labor intensive. Better, automated *service management* — as called for in NEDC design — can greatly reduce human-intensive labor costs.

Addressing These Problems: The NEDC Model
To help enterprises overcome these distributed computing design complexities — as well as to help enterprises dramatically reduce IT operations costs — IBM NEDC model focuses on:

1. Improved information systems resource utilization through resource consolidation and virtualization;
2. Energy efficient, green and optimized infrastructure and facilities;
3. Business driven service management;
4. Improved security and business resilience; and,
5. Adoption of an information infrastructure strategy addressing information availability, retention, security, and compliance.

Three Phases of NEDC Design
In general, there are three phases (stages of adoption) involved in moving toward the new enterprise data center model of computing (see Figure 1):

1. Simplified;
2. Shared; and,
3. Dynamic.

Figure 1: The Three Phases of NEDC Design

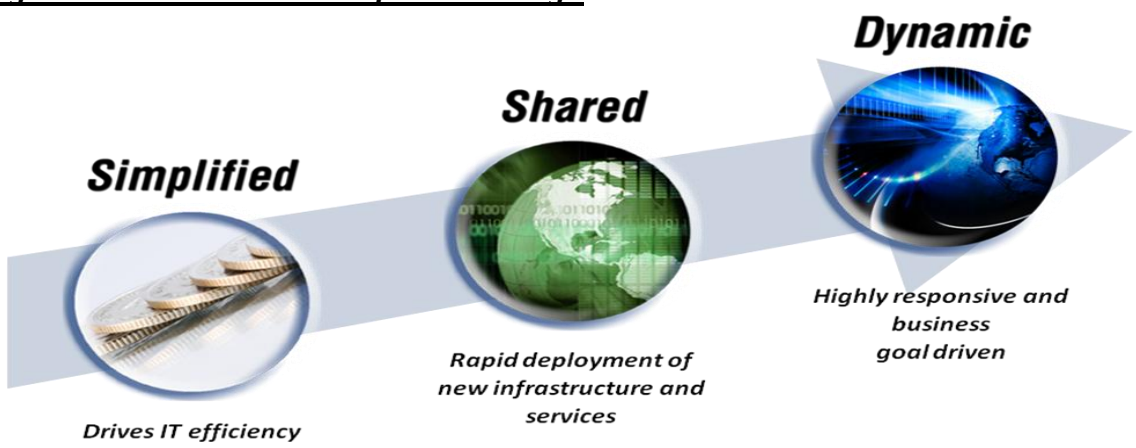

Source: IBM Corporation — July, 2008

The *simplified phase* of adoption concentrates on driving IT acquisition and operational costs down through consolidation and virtualization. In this phase of adoption, enterprises start to concentrate distributed resources into fewer data centers, servers, storage, and network devices (to simplify management). And enterprises also virtualize resources (create logical groupings or pools or physical servers, storage, and network devices) to increase utilization rates.

The *shared phase* changes the focus of IT administrators and managers from the management of individual systems to the aggregation and sharing of "like" resources. These like resources are clusters or "ensembles" that are virtualized and aggregated resources of the same type with management aggregated and elevated to managing the shared pool of resources as if it were a single system. This phase also introduces the concept of managing service flows. By breaking down the siloed walls between *like or similar* resources (such as x86 servers, or siloed storage devices), regrouping these devices into logical resource pools, and then deploying end-to-end systems, storage, and network management — enterprises can radically reduce IT operational costs. And by moving to a service-oriented infrastructure, enterprises can now focus on orchestrating process flows and managing services rather than managing disparate system/storage/network devices.

A shared infrastructure permits movement, control and balance of workload and data across "like" environments. In this phase of adoption IT managers and administrators start to concentrate more on process flow and service management (because the management and provisioning of underlying information systems becomes — to a large extent — automated).

The *dynamic phase* of new enterprise data center adoption completely obscures the physical tie between service delivery and the characteristics of the underlying IT infrastructure. Information systems become an IT "cloud" (a collection of ensembles) where users request services/service levels without having to worry about the complexities of the underlying information infrastructure.

In this phase, because services are no longer tied to specific, siloed technologies, IT managers can modify resource assignments to ensure that IT systems are running as efficiently as possible without disrupting the general flow of business. Standardization and automation of policies, processes, and procedures serve to further streamline operations in this phase of the new enterprise data center.

From a business resiliency perspective, highly-virtualized systems, storage, network resources contribute to improved resiliency (because back-up, redundant, available resources can be found easily in the cloud). And all of the efforts that have been made in the previous phases to improve availability, to establish fault tolerance, and to enable heterogeneous and homogeneous business continuity all meld to set the stage for cost effective, fully autonomic business resiliency.

> ***In the dynamic phase information systems become highly responsive the needs of a given business. Information systems no longer dictate business flow — instead they automatically serve business flow according to predetermined business rules. And these rules can be modified on-the-fly to adapt to changing market conditions or competitive pressures.***

PART 2: Building Business Resiliency Into the New Enterprise Data Center Model

The term "business resiliency" is used to describe a *comprehensive strategy that allows enterprises to rapidly adapt and respond to risks, as well as opportunities, in order to maintain continuous business operations, be a more trusted partner, and to enable growth.* A more in depth description of business resiliency can be found in "Beyond Disaster Recovery: Becoming a Resilient Business" by Richard Cocchiara, IBM's chief technology officer for business continuity and resilience services at:

http://www-935.ibm.com/services/us/bcrs/pdf/wp_becoming-a-resilient-business.pdf

When designing a resilient environment, *the ultimate goal is to deliver a fully recoverable, highly available, fault tolerant architecture that provides continuous computing services despite any kind of business disruption.* (Outages can be as simple as an application response time issue, an internal or external security breach...). In its final phase of implementation, this fully resilient environment should be automated to ensure rapid response to failures, as well as to significantly reduce costs related to IT management labor.

> ***The gating factor when it comes to business resiliency is investment capital cost (the best practices know-how, products, and related services for implementing a fully autonomic, resilient business environment exist today). Enterprises need to weigh their business needs versus the costs for business resiliency features and functions.***

The implementation of a resilient business environment can be accomplished in steps (see Figure 2).

Figure 2 — Incremental Investments Leading to Fully Automated Business Resiliency

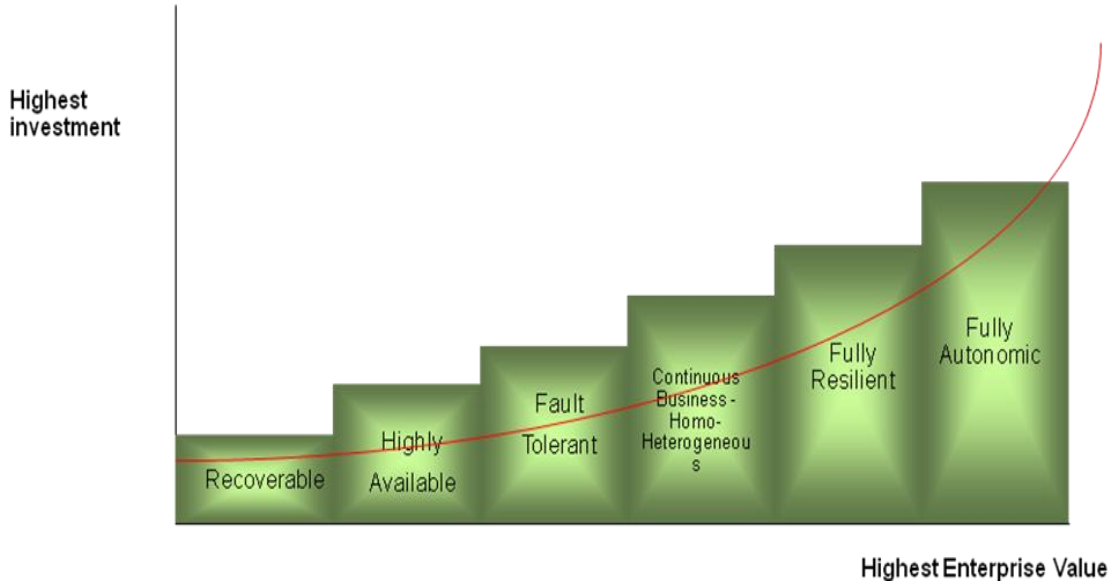

Source: IBM Corporation — July, 2008

Addressing Four Critical Client Needs When Building Resilient Systems

Each phase of implementing a NEDC requires different levels of business resiliency. The simplified phase, for instance, requires the ability to recover from failures (disk failures, power outages, etc.) at a bare minimum. The far more advanced dynamic phase requires fully resilient information systems that are highly available and/or fault tolerant to meet business service demands. The dynamic phase may require automated management services (policies and procedures) to simplify the management of dynamic resources.

In general, building a resilient new enterprise data center entails putting in place:

1. *Disaster recovery plans, products, policies, and procedures*. Disaster recovery plans, and associated policies and procedures, enable businesses to recover quickly from failures. These plans minimize the impact of unforeseen events and enable businesses to respond quickly to failures while minimizing the costs, time, and risks associated with such failures.
2. *High availability and/or fault tolerant systems/storage/network environments*. Availability requirements differ by industry — but plans need to be put in place to ensure that the right level of information systems reliability is put in place to meet 24x7 internal or regulated recovery and uptime demands.
3. *High availability/fault tolerance monitoring/management systems*. Enterprise availability needs to be monitored and managed. Increased resilience is provided through availability management at the component business system level. Monitoring resources helps reduce outages; while simplifying management reduces management labor overhead. And, ultimately,
4. *An environment that minimizes disruption by ensuring business continuity*. All of the preceding activities help enterprises maintain business operations in the face of a disruptive event — as well as meet compliance requirements.

Figure 3 illustrates these four "critical client needs".

Figure 3: Building Business Resiliency: 4 Critical Client Needs

Source: IBM Corporation — July, 2008

Mapping Business Resiliency into the Three Phases of New Enterprise Data Center Design
As described earlier, there are three phases of implementation for building a new enterprise data center. Each phase carries with it different requirements for business resiliency.

The *simplified phase* focuses on consolidation and virtualization; energy assessment, the deployment of energy efficient systems, and power monitoring; and basic service management (monitoring of systems/storage/networks/applications/databases; auto-discovery of resources; and workload automation).

> ***In this phase, the primary foci from business resiliency/security perspective are to implement back-up and recovery systems, server high-availability, and security encryption services.***

The *shared phase* stresses the organization of systems and storage into more easily managed ensembles (groupings of like resources) that can be further virtualized and better managed. In this phase enterprises move to a service-oriented infrastructure — and optimize that infrastructure for service delivery. Green activities move beyond the assessment and monitoring of systems environments to encompass the entire data center (chillers, power management systems, air conditioners, etc.). And, because a service orientation delivery approach is leveraged in this phase, a service management strategy needs to evolve and rudimentary business service management functions need to be introduced. It also makes sense to introduce a change and configuration management database (that serves a critical role in simplifying the monitoring, control, and management of resources as an information systems environment evolves into the dynamic phase).

> ***In this phase, business resiliency and security activities need to focus on ensuring that geographically dispersed clusters can be kept up-and-running; isolation/integrity/identity security systems need to be in place; and business continuity and resiliency services also need to be put in place.***

In the *dynamic phase*, ensembles are grouped into cooperative clouds where resources can be made available dynamically (dynamically provisioned) to address business needs. The whole IT environment essentially becomes a service point for process flows. Advanced energy efficient practices can be implemented in this phase including power optimization driven by types of workloads, just in time capacity delivery (unused resources can be turned off when not in use), and the Internet can be used to deliver additional scalability. Service management in this phase is heavily automated: data center operations operate in a lights out mode, automated scheduling of resources is orchestrated, and composite applications (applications that are "mashed" together to create a business result) can be managed transparently.

> ***Business resiliency in this phase features continuous availability, continuous data protection, and automated archiving.***

The entire phased approach to building business resilient NEDCs is illustrated in Figure 4.

Figure 4 — The Phased Approach for Introducing Business Resiliency into the NEDC

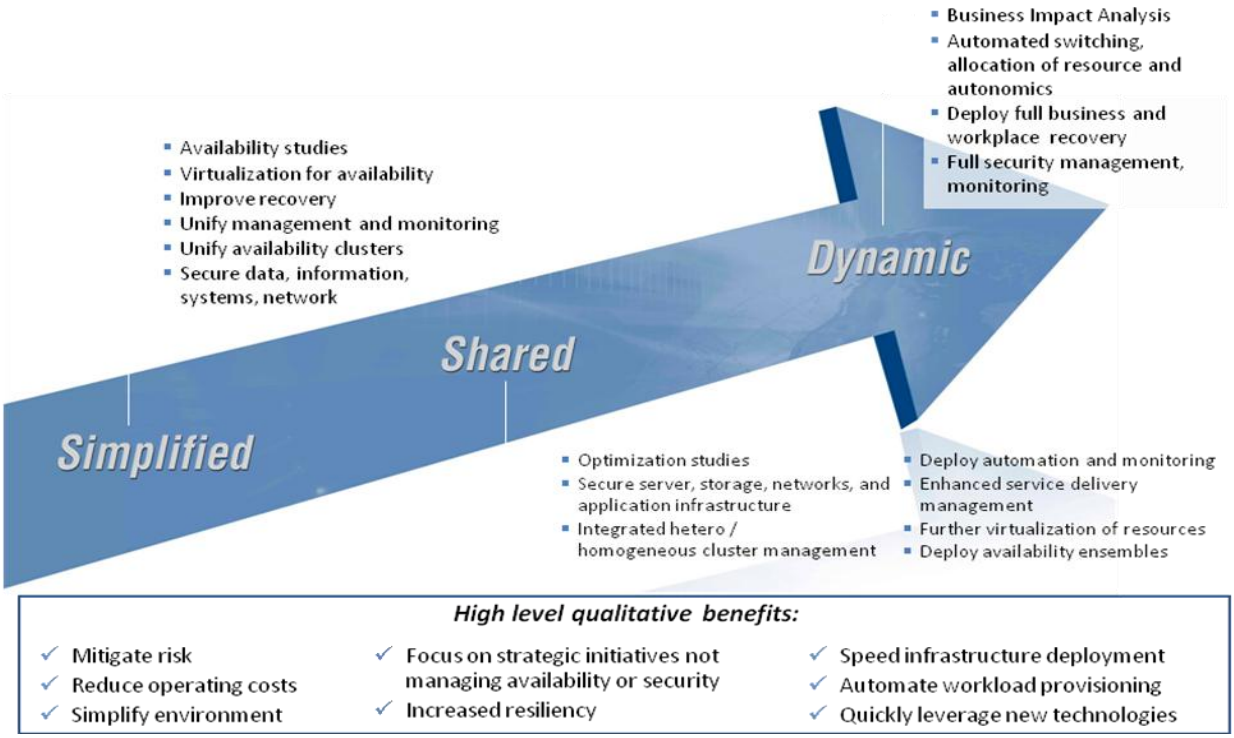

Source: IBM Corporation — July, 2008

PART 3: Step-by-Step Recommendations for Building a New Enterprise Data Center

This section examines the "big picture" when it comes to business resiliency planning — and then drills down into a specific, five point action plan for deploying resilient business systems.

Figure 5 illustrates how business resiliency spans across the entire enterprise, reaching into organizational needs, process support, application/data design, technology deployment and facilities management.

Figure 5: Cross-organizational, Multi-dimensional Business Resiliency Planning

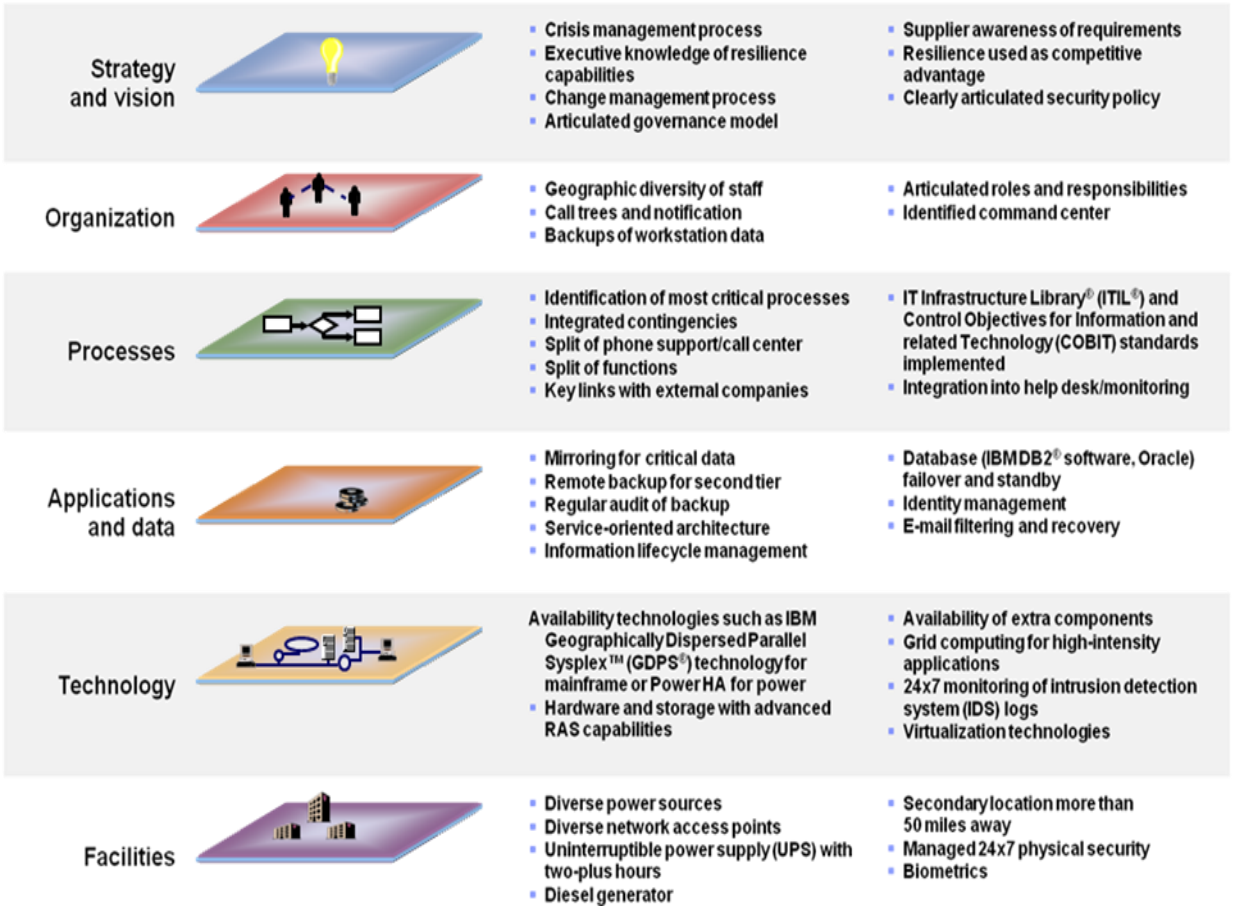

Source: IBM — July, 2008

Some of the key examples of capabilities and best practices in Figure 5 include:

- When building a business resiliency strategic plan, elements such as crisis management, change management, and security — as well as high-level executive sponsorship — should be considered. Further, note that resiliency is often implemented to help enable compliance (compliance with uptime or recovery time requirements, access to online data requirements);
- From an organizational perspective, geographical location of staff (geographical diversity) needs to be considered — and, should a business failure occur, roles and responsibilities to respond to that business failure need to be clearly articulated.
- Process flow is also important when building a business resiliency plan. Processes need to be assessed — and a list of which processes are the most critical is vital should an enterprise need to respond to an outage. Further, contingency plans need to be developed to ensure process flow.
- It is important to back-up critical data, as well as to ensure that critical applications remain running in the event of a failure;

- From a technology perspective, decisions need to be made on which applications should be mirrored (and which applications should receive high-availability support or fault tolerant support). Further, login, authentication, and security policies need to be established for mirrored systems; and,
- From a facilities perspective, diverse power sources should be made available — as well as diverse network connection options. Uninterruptable power supplies should be put in place to ensure information systems operations — and secondary operations need to be established at least 50 miles away from a main computing site to weather catastrophes.

Step-by-Step Business Resiliency Design Considerations

Based upon its experience to date in building new enterprise data centers, IBM suggests that IT systems designers/strategic planners focus on four activities as they seek to build fully resilient business systems. Those activities involve are:

1. Improving recovery (faster recover from unplanned outages or catastrophic failure);
2. Improving uptime/availability (increased uptime and availability for servers and applications);
3. Introduce enterprise availability management (improve operational efficiency and reduce downtime); and then
4. Implementing a continuously operational business environment (implement business continuity by building a highly automated information system/service management environment that yields no downtime, offers full business recovery, and is fully secure).

The remainder of this section examines these points in greater detail.

Improving Recovery: Building a Disaster Recovery Plan

First and foremost in building business resiliency is the implementation of a comprehensive disaster recovery plan. This plan should be developed with the following 5 steps in mind:

1. Design a plan aligned to your vulnerabilities, compliance requirements, business needs and risk-tolerance levels.
2. Implement designs to reduce the risk of downtime and associated revenue losses during a disaster.
3. Extend disaster recovery to new components, systems and locations.
4. Incorporate site, workplace and full-business recovery. And,
5. Test and update plans regularly.

Improving Uptime and Availability: Clustering, High Availability, and Fault Tolerance

A clustering, high availability, and/or fault tolerance plan needs to assess back-up/recovery and availability needs by analyzing application, system, and business unit requirements — and aligning availability requirements with specific business needs. The recommended approach for doing this is to:

1. Assess existing implementation against requirements by application, system, business unit.
2. Strengthen existing implementations (clusters, virtualization, automatic restart, fault-tolerance, ensembles).

3. Implement redundancy for business or mission critical applications.
4. Accommodate new business requirements. And,
5. Align availability with business needs.

Introduce Enterprise Availability Management

Enterprise availability management not only means monitor/control activities to ensure systems up-time, it also includes facilities that allow for demand fluctuation, system health assessments, and so on. When implementing an enterprise availability management environment, consider the following steps:

1. Assess existing implementations for improvement.
2. Design and implement highly available, resilient solutions.
3. Be prepared for demand fluctuations to allow for seamless adding of capacity.
4. Visualize business performance and operational health of services and process through dashboards. And,
5. Leverage this information to prioritize response and investments according to business impact.

Implement a Business Continuity Environment

IT designers need to start this phase of business resiliency implementation by reviewing the level of business resiliency that is required for applications and information systems within their respective environments — and meshing those requirements with existing disaster recovery and availability plans. Attention also needs to be focused on the scalability characteristics of information systems hardware as well as scalability requirements of the enterprise's various application software solutions. In this phase, IT designers should also consider using solutions that integrate storage with information systems in order to simplify management while increasing storage utilization rates. Plans also need to be developed for automating the management of information systems environments, allowing resiliency and continuity features to "kick-in" automatically should a failure occur.

The five steps in this phase, therefore, include:

1. Review resiliency requirements against existing implementations for disaster recovery and availability.
2. Design & implement required resiliency enhancements.
3. Use integrated hardware and software solutions that are scalable and adapt to your business model.
4. Use integrated storage solutions for data recovery and availability. And,
5. Implement automated system/site failover with fast server, storage, network and application recovery.

Improving recovery capabilities, implementing various levels of availability, and implementing an enterprise availability management environment all contribute to the end game — a fully automated, secure, resilient, service-oriented information systems environment. To reach this goal, all of the above recommendations should be implemented.

Summary Observations

The "new enterprise data center" is model that focuses on helping enterprises dramatically reduce IT operating costs, while realigning information systems to support service flows that are closely aligned with enterprise goals and objectives. This model relies heavily on consistent and reliable service delivery from the underlying computing "utility". Accordingly, systems hardware needs to be made reliable, available, and secure.

More specifically, the new data center model calls for a flexible IT infrastructure that is capable of handling dynamically changing workloads. To ensure service delivery, redundant components may be required (in high-availability/fault tolerant configurations). A service management environment also needs to be put in place to ensure service level delivery. And multiple-level security across the entire enterprise is a necessity. Ultimately, workplace/data center/full business recovery services may be required.

To ensure that businesses can recover from unplanned outages, IBM emphasizes that a comprehensive business recovery plan be put in place. IBM can assist in building such a comprehensive plan. The company's business resiliency background includes:

- 100% record for recovery of disasters declared;
- Vast experience in design and deployment of secure and resilient business systems — as well as disciplined, field-tested methodologies and technical support;
- Reliable, available, and secure (RAS) systems and storage environments that provide unmatched security and business uptime while reducing risk (IBM systems and storage delivers advanced availability and security features, information and data protection, extensive global technical support, and more);
- Industry leading software, servers, storage, and operating environments based on 40 years of providing high availability solutions;
- Award-winning software, technology leadership in hardware, industry-leading services and fault-tolerant systems;
- Award winning security and privacy software and services. (IBM holds more S&P copyrights than any other company in the world); and,
- The availability of custom designed, business resiliency and security solutions for enterprises in specific industries in specific geographies.

As for competitive offerings, some vendors can provide some of these products and services — but *only one vendor (IBM)* can provide a one-stop shop for all of the integrated hardware, software, and strategic planning, deployment and testing services needed to help reduce operational complexity, improve security, and ensure business resiliency across the entire new enterprise data center.

Clabby Analytics
http://www.clabbyanalytics.com
Telephone: 001 (207) 846-0498