Gartner

Research

Publication Date: 8 September 2010 **ID Number: G00206424**

Hype Cycle for Business Continuity Management and IT Disaster Recovery Management, 2010

Roberta J. Witty, John P Morency

Business continuity management and IT disaster recovery management are maturing practices as a result of supply chain pressures and the realization that a disaster can happen to any organization.

TABLE OF CONTENTS

Analysis 4
What You Need to Know 4
The Hype Cycle 5
Added Profiles 7
Renamed Profiles 7
Forward-Moving Profiles of Particular Note 7
Obsolete Before Plateau Profiles 8
Early Mainstream Profiles That Are Still Early in the Hype Cycle 8
The Priority Matrix 10
Off the Hype Cycle 11
On the Rise 12
Recovery Exercising 12
Cloud-Based Recovery Services 13
Disaster Recovery Service-Level Management 16
IT Service Failover Automation 17
Long-Distance Live Migration 19
U.S. PL 110-53, Title IX 20
Data Dependency Mapping Technology 22
Continuous Availability Architectures 24
At the Peak 25
Cloud Storage 25
Virtual Machine Recovery 26
Risk Assessment for BCM 28
Mobile Service-Level Management Software 29
Sliding Into the Trough 30
DR Insourcing 30
IT Service Dependency Mapping 32
Appliance-Based Replication 33
Lights-Out Recovery Operations Management 34
Data Deduplication 36
Hosted Virtual Desktops 37
Humanitarian Disaster Relief 39
Crisis/Incident Management 40
Business Continuity Management Planning Software 43
Emergency/Mass Notification Software 46
Workforce Continuity 49
BCM Methodologies, Standards and Frameworks 52
WAN Optimization Services 55
Climbing the Slope 56
Bare-Metal Restore 56
Business Impact Analysis 57
Distributed Virtual Tape 59
Pandemic Preparedness Planning 61
Work Area Recovery 64
Outage Management Systems 65
Print/Mail Business Continuity and Disaster Recovery 67
Entering the Plateau 68
WAN Optimization Controllers 68
E-Mail Continuity 69
Distributed Tape Backup 70

Gartner

Appendixes 72
Hype Cycle Phases, Benefit Ratings and Maturity Levels 74
Recommended Reading 75

LIST OF TABLES

Table 1. Hype Cycle Phases 74
Table 2. Benefit Ratings 74
Table 3. Maturity Levels 75

LIST OF FIGURES

Figure 1. Hype Cycle for Business Continuity Management and IT Disaster Recovery Management, 2010 9
Figure 2. Priority Matrix for Business Continuity Management and IT Disaster Recovery Management, 2010 11
Figure 3. Hype Cycle for Business Continuity Management, 2009 72

What You Need to Know

The main factors (listed below) driving growth in business continuity management (BCM) and IT disaster recovery management (IT-DRM) programs have not changed between 2009 and 2010:

- The 24/7 business delivery model
- Increasing operational risks
- Globalization of business operations

As a result of these drivers and their impact on business operations, compliance to standards and regulations is increasing globally. The U.S. Department of Homeland Security (DHS) and the U.S. Federal Emergency Management Agency (FEMA), as well as the governments of Australia and New Zealand, continue to deliver resources for all organizations to improve their business continuity readiness. However, the lingering impact of the 2008 recession has resulted in a large number of BCM professionals being out of work, so some BCM programs are not maturing as quickly as needed.

If done well, then a BCM program will create a single place in the organization where all business and technology processes are documented and maintained. The information collected and managed by BCM processes is a valuable resource that management can leverage for strategic and tactical nondisaster planning activities. However, the full potential of BCM has yet to be realized.

BCM processes are mature, but their implementation is not consistent within individual organizations and across all industries. For many organizations, some components (e.g., crisis/incident management, emergency response, IT-DRM, business recovery, contingency planning and pandemic planning) are more mature than others — for example IT-DRM is a fairly mature practice in large enterprises, but much less so in small and midsize businesses. A second example is business recovery — even if you have the strongest IT-DRM program, a disaster at a non-data-center facility could mean the demise of the organization because recovery procedures for business operations were not in place. In addition, BCM program responsibility must be positioned higher up in the organization for its related benefits to be more fully realized.

BCM tools have improved during the past three years to help document and manage information, but some are not being used to their full potential — e.g., BCM planning (BCMP) and crisis/incident management planning tools in particular. We moved the Crisis/Incident Management technology profile from prepeak to the Trough of Disillusionment because we included the full scope of crisis/management tools in the 2010 Hype Cycle profile — i.e., environmental, health and safety (EH&S), command and control, as well as BCM/operational. Emergency/Mass Notification Software (EMNS) adoption is growing quite well. Therefore, it will likely take another three to five years for many organizations to change over to a state in which BCM is used to support strategic and tactical operation resiliency across the organization.

One of the key reasons for this three- to five-year forecast is that IT-DRM, an essential foundation for effective BCM, remains labor-intensive for many organizations, especially in the area of recovery plan exercising. This labor intensity will become a significant barrier to scaling the IT-DRM program as more in-scope business processes, applications and data are added. Improved recovery services and management automation are critical to overcoming this barrier. However, many key enablers — such as more-formalized Disaster Recovery Service-Level Management, IT Service Dependency Mapping, Data Dependency Mapping Technology, Virtual Machine

Recovery and automated Lights-Out Recovery Operations Management — are still at relatively early stages in the Hype Cycle.

This contrasts with the increasing availability and viability of recovery facility alternatives from traditional colocation and hosting providers as well as an evolving community of "recovery in the cloud" providers. Because of the continued slow progress of recovery management automation maturity, more IT organizations are beginning to seriously evaluate the extent to which they should instead focus their time and resources toward building out a continuously available IT infrastructure. In addition, continuously available IT operations are a key technology prerequisite for sustainable business resiliency. However, assessing the required implementation time and cost is only at the due diligence stage. For many organizations, concrete realization is at least three to five years away.

The Hype Cycle

This Hype Cycle (see Figure 1) will aid BCM and IT-DRM leaders in identifying important processes and technologies, in assessing their level of adoption, and in modifying their management strategies to close key technology, process and knowledge gaps as these leaders chart a course toward a mature BCM program. The title of this year's Hype Cycle has been changed from the 2009 title, "Hype Cycle for Business Continuity Management, 2009," to include IT-DRM, and to make it explicit that IT-DRM processes and technologies are covered under BCM.

In this Hype Cycle — unlike most other Hype Cycles — technologies and methodologies whose business benefits were rated "High" were far more essential for maintaining existing business operations than for generating significant cost savings or enabling new competitive opportunities. For example, in most organizations, e-mail is considered an essential application to keep the business running; therefore, e-mail recovery has a high benefit to the organization. Additionally, the implementation of BCM processes, such as crisis/incident management, ensures that an organization can more effectively manage its public reputation and image in case a disruptive disaster event should occur.

The 2009 H1N1 pandemic was a catalyst for governments and organizations worldwide to enact strong pandemic preparedness planning processes and programs. In the Gartner 2010 Risk and Security Survey, 49% of responding participants said that they were doing so, as opposed to 37% in 2008. The concern is that, because H1N1 proved to be of limited impact, pandemic preparedness planning can easily fall back into the Trough of Disillusionment. However, the stage has been set for all organizations to monitor the threat and take action as needed. One complicating factor is the issue of delayed vaccine production. The other is antibiotic resistance diseases, such as the recent outbreak of a new gene (NDM-1) that alters bacteria, making it resistant to nearly all known antibiotics. Medical experts are warning that the booming medical tourism industries in India and Pakistan could fuel a surge in infections as patients import an increasing number of drug-resistant diseases into their home countries.

As BCM programs mature and expand their scope to include the use of work-at-home options, Work Area Recovery will make the transition from organization-provided facilities only to multiple options being supported in the overall BCM program. In addition, Hosted Virtual Desktop (HVD) technology may enable enterprise IT organizations to become economically viable providers for work-at-home services.

The growing visibility of BCM in boardrooms worldwide is focusing considerable attention on the development of a best-practice model for BCM methodologies, procedures, terminology, and so on. Still, for most enterprises, no single regulation, standard or framework exists that defines the BCM requirements that organizations should meet. However, the U.S. Private Sector

Preparedness (PS-Prep) program's selection of standards for organization certification — i.e., ASIS SPC.1-2009, British Standard (BS) 25999 and National Fire Protection Association (NFPA) 1600 — will move us closer to a "global standard" for BCM methodology. Supply chain pressure will advance the adoption and implementation of a global standard or set thereof. Hence, we moved the BCM Methodologies, Standards and Frameworks technology profile to post-trough.

Because IT-DRM is not a management category that has directly supporting products (unlike other management categories, such as configuration, change, availability and performance management), software product support for IT-DRM is, at best, fragmented and nonintegrated. The result is that IT-DRM recovery exercising and execution continue to be labor-intensive, with the burdens of execution and management software integration being left to the user. BCMP tools provide automation to structure the exercising process and should be considered for use.

Relatively new technologies — such as Virtual Machine Recovery, Bare-Metal Restore, IT Service Dependency Mapping, Data Dependency Mapping Technology and IT Service Failover Automation — have the potential to improve the degree of IT-DRM automation and reduce associated operating costs.

New generations of cloud-based infrastructure-as-a-service (IaaS) offerings, including recovery as a service (RaaS) and managed backup storage clouds, have the potential to significantly impact IT-DRM service delivery economics. However, Gartner believes that both service types are still at very early delivery stages. Therefore, given their relative newness and relatively small sets of supporting providers, clients should not assume that the use of cloud services to support DRM will largely subsume the use of traditional disaster recovery service providers, or hosting and colocation vendors, for at least the next three years.

In addition, the increasing number of privacy management and data breach notification regulations is causing many organizations to more closely scrutinize the breadth and maturity of service providers' operations management controls. The key challenge is in ensuring that these services can be securely, reliably and economically used to complement or entirely supplant more traditional shared-recovery services.

Finally, the issue of liability is a growing concern for organizations. Recent outages of third-party service providers have resulted in multiday business delays. The issue of how organizations are monitoring their third parties — and then who pays and at what level if there is an outage — is not well-known. Currently, the standard practice is to reimburse the customer for service fees. However, lost business compensation is not covered. The service provider cannot move the liability to the customer's business interruption insurance policy if it is a service provider problem. However, the total liability for all the service provider's customers can be more than the total value of the service provider. No insurance company would insure the service provider in these cases. Discussions are centering on the startup of a derivatives market for third-party risk trading — an option that, if not done transparently, could have a bad result.

Because operations recovery times need to become more predictable, the importance of IT-DRM service levels is continuing to grow. This has resulted in some providers — especially recovery-in-the-cloud providers — beginning to offer specific recovery time objective (RTO) and recovery point objective (RPO) service-level guarantees.

Because of the significant drop in required application and data recovery times, a more dedicated recovery infrastructure that is capable of supporting continuous disk-to-disk data backup (as an alternative to off-shift, tape-based backup) is causing many organizations to rethink their data center sourcing and management strategies. Increasing numbers of organizations are insourcing more IT-DRM program management responsibilities for reduced cost and operations recovery risk reasons. This transition is just beginning for many organizations and will require at least three to five years to complete.

IT organizations investing in technologies (such as the use of wireless e-mail) to mitigate the impact of a business outage should evaluate the availability service levels required by mobile users, and should identify potential risk factors for partial or complete interruptions of service.

The following specific changes to the 2009 BCM Hype Cycle were made to ensure that the major technology, service and methodology changes that occurred during the past year were properly documented and positioned relative to the BCM Hype Cycle curve.

Added Profiles

- Recovery Exercising, which is in a state of process and supporting technology transition
- Long-Distance Live Migration, the transition of virtual machine operations from one data center location to another
- U.S. PL 110-53, Title IX certification for private enterprises by the U.S. Department of Homeland Security
- Cloud Storage services for production data backup and recovery
- Appliance-Based Replication services
- WAN Optimization Services
- Data Deduplication
- Distributed Virtual Tape
- WAN Optimization Controllers
- Distributed Tape Backup
- Print/Mail Business Continuity and Disaster Recovery

Renamed Profiles

- 2009 BCM Regulations, Standards and Frameworks renamed BCM Methodologies, Standards and Frameworks in 2010
- 2009 Emergency Notification/Mass Notification Software renamed Emergency/Mass Notification Software in 2010
- 2009 Business Continuity Management Planning Tools renamed Business Continuity Management Planning Software in 2010

Forward-Moving Profiles of Particular Note

- The Crisis/Incident Management profile was moved from prepeak to the Trough of Disillusionment because we included the full scope of crisis/management tools in the 2010 Hype Cycle profile — EH&S, command and control, as well as BCM/operational.
- The Workforce Continuity profile was moved from prepeak to post-trough. In prior BCM Hype Cycle reports, our Hype Cycle rating for Workforce Continuity was based on the maturity of the technical solutions for workforce continuity — none of which are being used to any great extent at this point, and which have been de-emphasized in most recovery vendor offerings. For 2010, our rating is based on the overall process of

ensuring that workforce continuity preparedness needs are addressed. Therefore, it moved significantly to a post-trough position.

- The Virtual Machine Recovery profile was moved from mid-prepeak to peak because the rate of server virtualization technology adoption is increasing, and because the workloads being virtualized are becoming more mission-critical. Gartner estimates that, by the end of 2010, approximately 25% of enterprise workloads will reside on virtual machines.

- The Risk Assessment for BCM profile position was changed from prepeak to postpeak because BCMP has often been conducted at a very superficial level of risk assessment, or even none at all. Although it has been well-understood that risk assessments are a necessary component of BCMP, the line of business often considers them to be time-consuming and too resource-intensive. This opinion has been justified, given the general lack of effective risk assessment methods and tools, and has often been exacerbated by the inappropriate use of such tools and methods. However, expectations of better levels of practice are increasing, encouraged to some extent by standards such as ITIL, Control Objectives for Information and Related Technology (CobiT), International Organization for Standardization (ISO)/International Electrotechnical Commission (IEC) 27001, BS-25999, ASIS SPC.1-2009 and NFPA 1600.

Obsolete Before Plateau Profiles

- Outage Management Systems (OMSs) are predicted to become obsolete before maturity, because the new breed of distribution management systems (DMSs) will eventually incorporate the OMS functionality as we know it. DMSs will include OMSs within real-time advanced distribution supervisory control and data acquisition (SCADA), which also will include automated restoration and self-healing "smart grid" functionality.

- Although Virtual Machine Recovery tools showed increased adoption in 2009, these point solutions are likely to give way to more-generalized tools. These solutions could be extensions of current recovery applications and tools.

Early Mainstream Profiles That Are Still Early in the Hype Cycle

- A number of profiles are ranked as early mainstream adoption, but still early on the Hype Cycle — e.g., Cloud-Based Recovery Services, Long-Distance Live Migration, Humanitarian Disaster Relief, Crisis/Incident Management and so forth. The reason for this placement is that recovery technology adoption lags the production-ready technology usage. The recovery technology is more mature because it has to follow the production-ready technology offering, but market adoption is low due to recovery efforts lagging overall in focus and investment in the majority of organizations. Specifically for cloud-based recovery solutions, privacy and security are outstanding issues that are yet to be resolved; therefore, adoption is low for this reason as well.

Figure 1. Hype Cycle for Business Continuity Management and IT Disaster Recovery Management, 2010

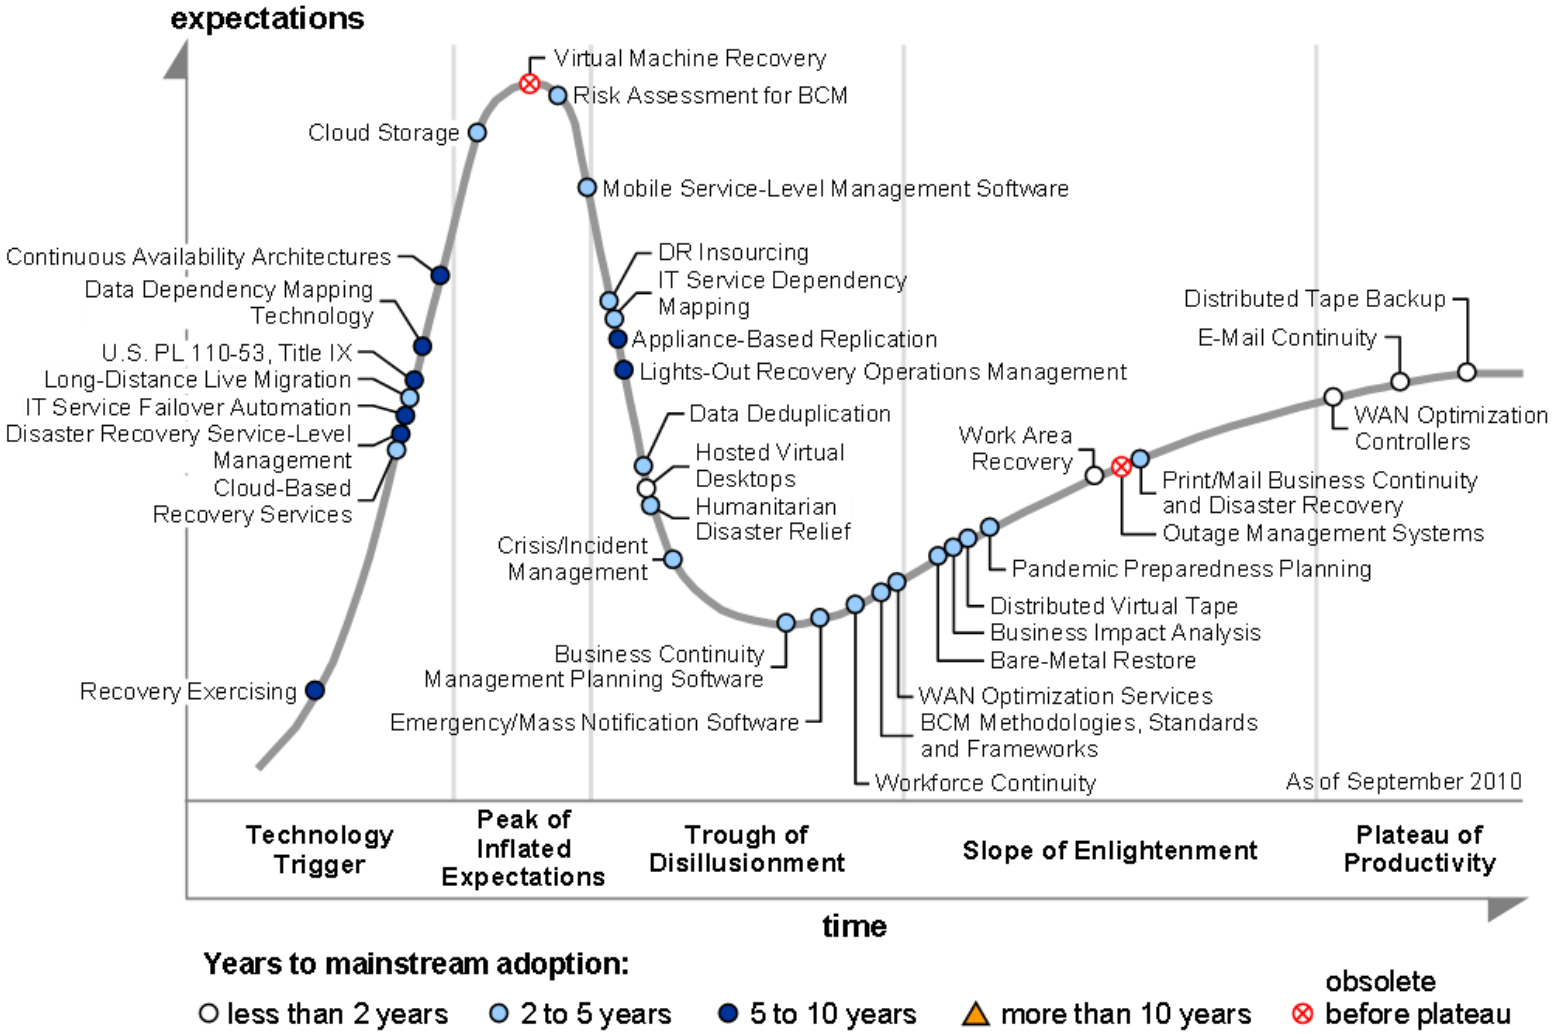

Source: Gartner (September 2010)

The Priority Matrix

The BCM Priority Matrix (see Figure 2) shows the valuation of the 35 continuity and recovery profiles in the 2010 Hype Cycle. The Priority Matrix maps the benefit rating of a process or technology against the length of time that Gartner expects it will take to reach the beginning of mainstream adoption. This mapping is displayed in an easy-to-read grid format that answers these questions:

- How much value will an enterprise get from a process or technology in its BCM program? When will the process or technology be mature enough to provide this value?
- In the case of a process, when will most enterprises surpass the obstacle that inhibits their ability to achieve mature BCM programs?

This alternative perspective helps users determine how to prioritize their BCM investments. In general, companies should begin in the upper-left quadrant of the chart, where the processes and technologies have the most dramatic impact on ensuring a strong ability to recover and restore business and IT operations after a business disruption (and these processes and technologies are available now or will be in the near term). Organizations should continue to evaluate alternatives that are high-impact, but further out on the time scale, as well as those that have less impact, but are closer in time.

Many profiles that have a "High" ranking are process-oriented. Therefore, the most important piece of advice that Gartner can provide is to look at the BCM methodology being used in the BCM program so that consistency of program implementation is achieved across all lines of business.

No profiles have a "Low" ranking because *all* recovery activities are critical when disaster strikes and the business needs to be recovered.

No profile has an adoption rate of more than 10 years, because BCM and IT-DRM processes and technologies follow the business process and technical architecture of the organization. As a technology is developed, the recovery for that technology is developed.

The Data Deduplication technology profile has a ranking of "Transformational" because it can reduce disk storage costs by a factor of 15 to 25 times over nondeduplication recovery solutions.

Figure 2. Priority Matrix for Business Continuity Management and IT Disaster Recovery Management, 2010

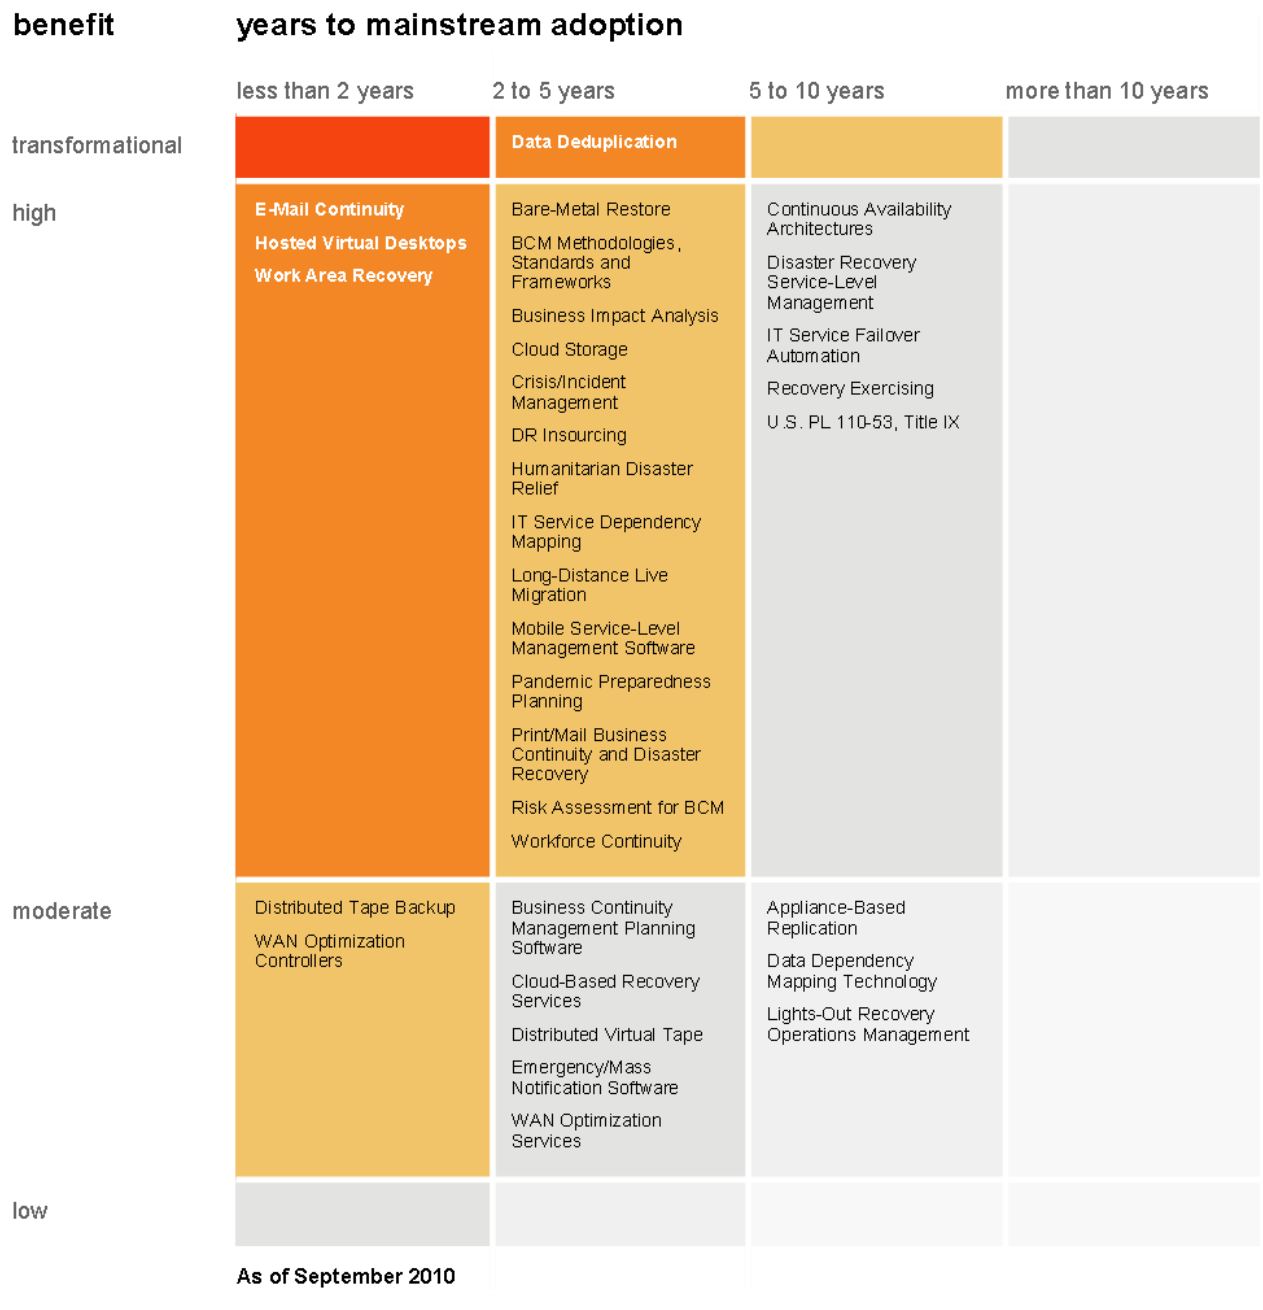

Source: Gartner (September 2010)

Off the Hype Cycle

We did not include the following BCM technology in the 2010 Hype Cycle:

- Internet Data Center Colocation, which was positioned on the Slope of Enlightenment in the 2009 BCM Hype Cycle and is considered highly mature at this point.

On the Rise

Recovery Exercising

Analysis By: John Morency

Definition: Recovery exercising (also known as disaster recovery [DR] testing) is the set of sequenced testing tasks typically performed at a recovery data center facility and which are focused on ensuring that the access and usage of a production application (or group of production applications) can be restarted within a specified time period (the recovery time objective [RTO]) with the required level of data consistency and an acceptable level of data loss (the recovery point objective [RPO]). As the recovery scope of mission-critical business processes, applications and data increases, however, sustaining the quality and consistency of recovery exercises can be a daunting technical and logistical challenge, especially as the frequency with which recovery exercises are held increases, in addition to increased change frequency.

Regardless of the frequency with which recovery exercises are held, however, an inescapable fact is that the consistency between the current state of the production data center infrastructure, applications and data and their state at the time of the last recovery test erodes on a daily basis as a direct side effect of the changes that are applied to support new business requirements. Given that the number of monthly changes processed in many production data centers is on the order of hundreds, the variance between the current change state of production applications and data and the state that existed at the most recent recovery exercise can potentially be on the order of many hundreds, if not many thousands, of changes. In addition, travel and entertainment (T&E) costs associated with performing recovery exercising at a remote data center facility have come under increased executive scrutiny, making the challenge of maintaining the current testing frequency even more daunting. Finally, recovery testing is still either a partially or totally manual exercise, making exercise scalability more difficult as new in-scope applications and data are brought into production.

Hardware, software and service advances are beginning to improve how recovery exercising is performed. Secondary development and test configurations housed in dedicated data center space have become an increasingly attractive testing alternative to DR provider-managed space and equipment for large enterprises. To reduce test time, especially for mission-critical applications, some organizations are rearchitecting their most-critical applications for active-active processing, meaning that the application runs "live" in two or more data centers and that testing is done every day by means of the production implementation across two sites. In addition, for all applications, we recommend the use of dependency mapping tools and/or configuration compliance tools so that application configurations can be compared across the two data centers, to find misconfigurations (for example, changes applied to one data center but not the other) in a proactive way, and not waiting for a test (or worse, a disaster) to uncover a recovery risk.

A new generation of recovery-in-the-cloud offerings has the potential to both improve the frequency with which provider customers can conduct live testing, as well as eliminate recovery configuration server and storage capital cost. To remain competitive, DR providers for their part are continuously improving remote access into their data centers, which enables in-house recovery management teams to remotely orchestrate live exercises without having to travel. Supporting this, however, requires the provider to ensure that proper authentication, access and (if needed) data encryption controls are in place.

Unfortunately, the number of software products that support recovery exercise task workflow automation are very few. Some (such as VMware Site Recovery Manager [SRM] or Ermestel Reliable DR) are specific to a single hypervisor vendor while others (such as Sanovi DRM and

Symantec VCS FireDrill) are vendor-independent. This means that most organizations must still manage the arduous tasks of developing, testing and maintaining custom recovery run book scripts that are typically manually executed with results being logged and later analyzed for future improvement opportunities.

Position and Adoption Speed Justification: Despite the increasing availability of supporting products during the past few years, today's reality is that the majority of organizations must still manage very labor-intensive recovery exercises. Given that these same organizations will likely be adding new mission-critical applications and data to the production workload to remain competitive, exercising time and cost challenges will become increasingly exacerbated. This will inevitably increase both interest in and application of exercise management products during the next few years. For this reason, the initial positioning of recovery exercising management technology on the Hype Cycle curve is at the post-trigger 10% point.

User Advice: As a near-term alternative, one approach to consider is the consolidation of what were previously separate preproduction QA testing and DR test teams. The multipurposing of server and storage assets for supporting both day-to-day development and testing activities, as well as recovery exercising (and application failover operations, if needed), combined with the needs for tighter change synchronization between configurations and automation of more-frequent test exercises, have been some of the key drivers. It is important to note, however, that an organizational and tools consolidation, by itself, may not necessarily be sufficient. The ideal scenario is one in which a separate test environment can be configured for exclusive use by the merged organization. However, due to budget-related constraints, this may not always be immediately doable. In addition, an important success factor is the availability of automated application test management software, combined with customized run book automation scripts, to manage the orchestration of the infrastructure and application failover sequences that may be required.

The benefits that have been realized by some of the early adopters of this approach include increasingly reliable and more-effective test exercises, combined with a more thorough testing of representative production inquiries and transactions against the recovery configuration. The latter benefit improves the likelihood that recovery operations can be initiated within required RTO and RPO targets, as well as ensuring more-stable recovery operations. If this is not already being evaluated, Gartner recommends that this approach be considered as one possible alternative to continued maintenance of the status quo.

Business Impact: The ability to automate recovery exercise tasks in an automated, repeatable and timely manner is becoming increasingly important for many organizations. As Web applications support an increasing number of business-critical process activities, effective recovery exercise management will become an important foundation for the realization of improved business resiliency.

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Sample Vendors: Sanovi; VMware

Cloud-Based Recovery Services

Analysis By: John Morency

Definition: Disaster recovery (DR) services delivered by public cloud providers today are primarily infrastructure-as-a-service (IaaS)-class offerings. These include recovery-in-the-cloud

and cloud-based storage services. This profile specifically focuses upon recovery-in-the-cloud services. Cloud-based storage services are discussed in the Cloud Storage profile of this Hype Cycle.

Recovery-in-the-cloud services typically support a combination of server image and production data backup to the service provider's data center. When access to the replicated server images and production data is required by the customer for plan exercising or to support live recovery operations, the server images are dynamically restored to available hardware and reactivated. All recovery-in-the-cloud providers support the restoration of VMware .vmdk image files and some (e.g., Amazon Elastic Compute Cloud (Amazon EC2) services supplemented by the use of Double-Take Software's Cloud) can support a combination of virtual machines and physical servers through the use of bare-metal restore technology.

Once server image restoration is complete, network access to the relevant applications and production data can begin. A software-based agent or provider-specific appliance installed at the customer site supports the initiation and orchestration of image and data backups. In addition, storage area network (SAN)-to-SAN replication, supported through compatible SANs (e.g., NetApp-SAN-to-NetApp-SAN, EMC SAN-to-EMC-SAN) on the customer and provider sides, is a third backup option.

The mean time for a server image to become operational inside the provider's cloud will vary, based on a number of factors. Some of these factors include the amount of prerequisite system configuration, customization, patch updating and data replication needed**.** In addition, Web-based application startup time will vary, depending on the number and ordering of server images required to support application startup. Cloud-based storage can be used for several different applications, including nightly data backups, archival data storage or ad hoc DR test storage that can be provisioned on demand.

The recovery-in-the-cloud value proposition is twofold. First, because server restoration on demand does not require the preallocation of specific computing equipment or floor space, provider customers have the opportunity to exercise their recovery plans more frequently. This contrasts with the 10 to 14 weeks advance notice that must be given to providers of equipment-subscription-based recovery services prior to beginning a test exercise. In addition, because server images are restored to providers' server hardware when needed, and production data has already been stored inside the provider cloud, the need for either shared-subscription or dedicated server and storage equipment can be significantly reduced, if not totally eliminated. Typically, the server hardware maintained inside the provider's cloud primarily supports either Windows- or Linux-based applications.

Representative service providers include Amazon (through its EC2 service), BlueLock, Databarracks, Big Little Fish, Doyenz, GCloud3, Geminare, i365, iland, NaviSite, Rackspace, SunGard Availability Services and Terremark Worldwide. Service subscription pricing per server image ranges from $100 to $300 per month, depending on the provider. In addition, data storage pricing ranges from as little as $0.25/month/gigabyte to as much as $3.00/month/gigabyte, depending on the provider and the degree of additional provider-supplied storage management that may be required.

Position and Adoption Speed Justification: These are key criteria for evaluating a recovery-in-the-cloud provider:

- Provider reputation
- Recovery data center availability and accessibility

- Formalized service level management (either based on recovery time objectives [RTO] or recovery point objectives [RPOs])
- Lead time required for server provisioning, configuration (if needed) and reactivation
- Network connectivity management (including circuit failover management)
- Frequency with which the provider client can update server images
- Frequency with which the provider client can update production data
- Additional recovery exercise orchestration management services
- Breadth and depth of production management controls, especially as they relate to data privacy and integrity management
- Provider policy on geographical hosting of server images and production data

The natural service strengths of DR-in-the-cloud providers align reasonably well with the first two criteria, assuming the customer can strongly influence data center location. In addition, the third criterion can also be addressed for all network connectivity — with the exception of the last mile. The criteria for which there is far less alignment include the availability and provisioning of computing equipment that is not Windows- or Linux-based (for example, z/OS-based mainframes, IBM iSeries systems, HP-UX-based servers, IBM AIX-based servers or Oracle Solaris servers). Another shortfall is the degree of choice for data replication options, as well as the extent of customer data reformatting that may be required to map the customer data into the format(s) that can be transferred and/or stored by the provider. In addition, the customer is still generally responsible for managing recovery exercising and testing.

Regardless of the specific provider under consideration, recovery-in-the-cloud services are still at a very early delivery stage. Some providers may not have the operations management maturity and process rigor that one would find in a more established DR service provider, established hosting vendor or IT operations outsourcer. However, given the evolving maturity of both server image replication and cloud storage technologies (see the Virtual Machine Recovery and Cloud Storage profiles in this Hype Cycle), as well as the increasing number of more-established providers (e.g., Amazon, NaviSite, Rackspace SunGard Availability Services and Terremark) offering recovery-in-the-cloud services, Gartner has raised the Hype Cycle position for recovery-in-the cloud services to Trigger-Peak Midpoint.

User Advice: Given the relative newness of recovery-in-the-cloud services, do not assume that cloud-based services will subsume the use of traditional DR service providers or hosting and colocation vendors in the near term or midterm. DR and business continuity requires the management of technology and operations management disciplines, not all of which can be readily addressed by cloud-based DR services. Examples include work area recovery, incident management, crisis communications and data archival. Therefore, it is important to look at cloud-based services as just one possible means by which server image backup and data replication requirements can be addressed.

Consider cloud infrastructure when you need DR capabilities for applications whose primary residence is already outside your data center or that are Windows- or Linux-based, or when you need low-cost recovery for browser-based applications. Because cloud services are still nascent and few providers currently offer DR-specific service levels for recovery time or recovery point objectives, carefully weigh the cost-benefits against the service management risks as an integral part of the decision-making process for DR sourcing.

Business Impact: The business impact is moderate. The actual benefits will vary, depending on the diversity of computing platforms that require recovery support and the extent to which the customer can orchestrate (and ideally automate) the recurring recovery testing tasks that need to be performed, as well as the extent to which the customer can transparently and efficiently utilize same-provider cloud storage for ongoing data backup, replication and archival. The key challenge is in ensuring that these services can be securely, reliably and economically utilized to complement or entirely supplant the use of more-traditional equipment-subscription-based services.

Benefit Rating: Moderate

Market Penetration: Less than 1% of target audience

Maturity: Emerging

Sample Vendors: Amazon; BlueLock; i365; NaviSite; Rackspace; SunGard Availability Services; Terremark International

Recommended Reading: "Predicts 2010: New IT Disaster Recovery Technologies Are Emerging, but Most Are in the Early Stage"

"Disaster Recovery Sourcing: The Time to Make More Informed Decisions Has Come"

"IT Disaster Recovery Sourcing Considerations"

"Web Hosting and Cloud Infrastructure Prices, North America, 2010"

"Dataquest Insight: The Changing Colocation and Data Center Market"

Disaster Recovery Service-Level Management

Analysis By: John Morency

Definition: Disaster recovery service levels are defined by recovery time objective (RTO) and recovery point objective (RPO), which are typically defined in units of minutes, hours or days. Disaster recovery service-level management refers to the support procedures and technology needed to ensure that committed RTO and RPO service levels are met during recovery plan exercising or following an actual disaster declaration.

External service providers offer two types of disaster recovery service levels. The first is an application-specific RTO- and/or RPO-based service level, which software-as-a-service (SaaS) providers often offer. One example of application-specific service levels is Oracle's Enterprise Disaster Recovery Option for its CRM On Demand service offering (supporting details are provided in "Oracle CRM On Demand Release 16 Lifts Custom Object Limit"). As announced by Oracle, this option supports RTO service-level guarantees of 24 hours, and has a one-hour RPO guarantee.

The second type of disaster recovery service level is an application-independent service level that is supported by a combination of server virtualization and virtual machine failover to a provider's managed facility. Virtual Server Replication Service (offered by SunGard Availability Services) is one example of provider-managed virtual machine replication and failover that supports a worst-case RTO of six hours. In general, application-independent recovery service levels require longer recovery times, due to the more-complex orchestration steps that are needed to ensure that all (versus one) in-scope systems and applications have been recovered and reactivated in a time frame consistent with the service-level targets. In addition, because RPO targets are very application-specific, application-independent service levels that are managed by external service providers are just defined in terms of the RTO.

In addition to external providers, IT recovery teams also support formal IT Disaster Recovery Management (IT-DRM) service-level targets. However, this is more typical in organizations that have fairly high IT-DRM maturity, and is not yet considered a mainstream best practice.

Position and Adoption Speed Justification: Not only is the market need for more predictable operations recovery increasing, but also the required recovery times for the most important mission-critical applications continue to be measured on the order of hours versus days. One potentially beneficial side effect of an increase in the number and quality of managed recovery service-level options could be a decrease in the time and effort needed for recovery testing, especially if the scope and effectiveness of the underlying application and data failover automation continue to improve.

This improvement will not happen in the near term; however, the result will be in disaster recovery service-level management being positioned at a relatively early stage in the implementation life cycle. Because of its dependence on base technologies, such as virtual machine recovery, data dependency mapping and continuous availability architecture, disaster recovery service-level management cannot realistically be positioned on the Hype Cycle at a point later than its technological prerequisites. For this reason, and due to the very low growth in provider-supported RTO- and/or RPO-based service-level agreements during the past year, disaster recovery service-level management remains at the trigger-peak midpoint position on the Hype Cycle curve in 2010.

User Advice: Over time, recovery, hosting, application and storage cloud providers may offer a more robust protection alternative to what in-house IT organizations can deliver. Because this is a nascent, but fast-growing, provider service differentiator, it is important to continually re-evaluate the recovery sourcing strategy to ensure that IT operations recovery continues to remain predictable, sustainable and cost-effective, regardless of who is responsible for delivering it at any point in time. Because service-level excellence is so critical to long-term provider viability, expect the growth of managed recovery service-level-based offerings for the near future.

Business Impact: The ability to manage recovery service levels in an automated, repeatable and timely manner is becoming increasingly important for many organizations. As Web applications support an increasing number of business-critical process activities, managed recovery service levels will become an important foundation for the realization of improved business resiliency.

Benefit Rating: High

Market Penetration: Less than 1% of target audience

Maturity: Embryonic

Sample Vendors: Oracle (Siebel CRM); SunGard Availability Services

Recommended Reading: "Toolkit Best Practice: Disaster Recovery Service Levels: What Makes Them Different and Why They Are Important"

IT Service Failover Automation

Analysis By: Donna Scott

Definition: IT service failover automation provides end-to-end IT service startup, shut-down and failover operations (local and/or remote for disaster recovery [DR]). It establishes ordering and dependency rules, as well as failover policies that are implemented in the infrastructure architecture. Most enterprises have implemented scripts for end-to-end IT service startup, shut-down and failover for DR and high-availability clustering across heterogeneous physical and virtual computing environments. These scripts are increasingly complex to maintain, with

changing requirements for more-granular high-availability (HA)/DR architectures. For example, most organizations implement simple DR failover policies, rather than more-fine-grained policies, which would enable partial failover functionality to exist because of the complexities of recovery script implementation and maintenance.

Some enterprises have taken this approach to the next step and have implemented these capabilities in their job-scheduling systems, because they have the concept of ordering and contain dependencies. However, the capabilities are more relevant for batch jobs (for example, if this file exists, then start this job), rather than for day-to-day startup, shut-down and failover operations of interactive systems. Removing and externalizing these automation procedures and policies from scripts into a more visual and service-oriented framework are necessary steps to reduce manual errors, enable ease of maintenance and foster the policy-based automation required for real-time infrastructure architectures. In addition, these prebuilt automation procedures enable IT operation centers to take control and restart failover applications with the "push of a button," including any required component ordering and dependencies.

Position and Adoption Speed Justification: Emerging tools in this arena come from the clustering heritage of Symantec (Veritas Cluster Server) and IBM (Tivoli System Automation, which has a heritage in clustering from High Availability Cluster Multiprocessing [HACMP]), which have an inherent understanding of resources and groups, and in which IT service dependencies and policies are built. Today, the vendors' definition of IT services is manually built, rather than derived from a configuration management database (CMDB). Both products have few customers in the distributed space, although IBM has a significant installed base on z/OS, and has since brought the functionality to Linux, AIX, Solaris and Windows.

Virtual server platform tools such as VMware vCenter Site Recovery Manager (SRM) have also entered this space, but only for IT services running on virtual-machine platforms. Run book automation (RBA) tools have also started to emerge with products for use in this area, with clients using the tools to aid in developing visualization around recovery automation (startup, ordering, dependencies), as well as a framework in which to organize their scripts. However, unlike clustering tools, RBA tools do not have an inherent understanding of the underlying physical infrastructure (but could be notified by clustering tools). As a result, RBA tools tend to be used to enable simpler maintenance around startup and shut-down routines, but not generally for granular failover policy automation (as clustering-based tools would be).

User Advice: Emerging IT service failover automation tools have the potential to broaden availability to the end-to-end IT service, and to enable granular rules for failover automation (for example, if there is a loss of the application server tier and a desire to fail it over to another data center, should the database tier failover as well?). While these tools add complexity to the environment, they have the potential to reduce it through visualization and external policy creation and maintenance. Moreover, IT service startup, shut-down and migration/failover rules could be established, maintained and tested all in one place for operations and HA/DR purposes, thus increasing the likelihood that HA/DR strategies will work when failures occur.

We recommend that enterprises watch this space and assess emerging IT service failover automation technologies to replace fragile, script-based recovery routines that are easier to maintain, as well as gain more granularity, greater consistency and efficiencies. In addition, because emerging tools in this space tend to be more loosely coupled, rather than tightly coupled, like that of traditional clustering, enterprises will be more likely to reduce the "spare" (that is, available) infrastructures required for HA, and thus reduce the overall cost of providing HA. Moreover, as more virtualized environments are deployed into production, these tools will be able to make use of the underlying virtual platform for HA, as well as virtual server mobility.

Business Impact: The potential business impact of this emerging technology is high, reducing the amount of spare infrastructure needed to ensure HA, as well as helping to ensure that

recovery policies work when failures occur. Moreover, efficiencies are increased because the maintenance of fragile recovery scripts is replaced with external policies that are easier to visualize, maintain and test. Furthermore, with IT service startup and shut-down ordering and dependencies being used both operationally and for recovery, external policies are more likely to be kept up to date, thus increasing consistency, service quality and recoverability.

Benefit Rating: High

Market Penetration: Less than 1% of target audience

Maturity: Emerging

Sample Vendors: IBM; Symantec; VMware

Long-Distance Live Migration

Analysis By: Donna Scott; Neil MacDonald

Definition: Long-distance live migration enables the movement of a running virtual machine (containing, for example, an application or a component of an application) from one data center location to another, over WANs with no loss of service to the application users. It is similar in concept to local virtual server live migration, but is performed over longer distances, for example among geographically disperse data centers.

Position and Adoption Speed Justification: Long-distance workload mobility is being developed, advanced and tested by hardware, software and virtualization infrastructure vendors, which see this as a logical evolution of local virtual server live migration for high availability, and as an enabling technology for hybrid cloud environments. It also is a great example of how virtualization can result in flexibility and agility (in this case, enabling an application to run in any data center). Because of significant research being done in this area and its appeal to IT organizations, we believe the technology will advance fairly rapidly — in two to five years, rather than in five to 10 years, which is more typical of new technologies.

There are many problems to be solved, however, including reliable session state transfer, potential latency issues — if an application or service is split across two data centers (as opposed to the entire application or service moved to the alternative location) — real-time data availability in both data centers, and movement of IT service traffic and processing from one data center to another. Because of complex application-specific dependencies, we do not foresee this technology fully replacing clusters or recovery technology; enterprises still require planning and testing to recover applications and infrastructures in the event of outages and disaster scenarios.

User Advice: The current technology is embryonic, but is advancing fairly quickly, due to significant vendor R&D efforts. Some application architectures and workloads (such as Web applications) may be more tolerant of long-distance migrations (versus others that may experience application time-out issues, or have long-running transactions that may interfere with session movement), so organizations should test early technologies and begin planning for use cases where it can be helpful in the future, such as for data center migrations and consolidations, as well as in anticipation of pending and known disasters. Enterprises should be sure to test the end-to-end movement and not just a live migration of a single workload, especially for use cases such as data center migrations (for example, due to consolidation or in anticipation of disaster). Moreover, enterprises should be sure that data replication provides the right degree of data availability in both data centers — the primary location and the backup site. In today's architectures, data is typically replicated and available in active/passive mode, meaning that there is a 15- to 60-minute startup time required to get an application online at an alternative data center. For long-distance live migration to meet its goals of application migration across data

centers, the data will have to be available in an active/active mode without compromising the integrity of the data itself.

Business Impact: Long-distance live migration would enable an IT organization to move workloads in anticipation of a disaster (such as a pending hurricane or flood) without the downtime that is typically required when performing a failover from one data center to another. It could also be used for data center moves, migrations and planned maintenance, or for rebalancing capacity across data centers, while reducing or eliminating the downtime associated with these initiatives or projects. Moreover, long-distance live migration could enable workload migration across internal and external service provider/cloud data centers in hybrid private/public cloud architectures. The technology has significant potential business benefits, from a runtime flexibility perspective, especially if the technology evolves to be enabled through the infrastructure (e.g., virtual server, storage and networking), without complex application design and engineering requirements. However, it will be critical to enable not just specific workload live migrations, but also the entire end-to-end IT service.

Benefit Rating: High

Market Penetration: Less than 1% of target audience

Maturity: Embryonic

Sample Vendors: Cisco; EMC; F5; VMware

U.S. PL 110-53, Title IX

Analysis By: Roberta Witty

Definition: On 3 August 2007, the U.S. passed legislation entitled Implementing Recommendations of the 9/11 Commission Act of 2007 (Public Law 110-53), which tasked the Department of Homeland Security (DHS) to develop a voluntary accreditation and certification program (see www.fema.gov/privatesector/preparedness) concerning private enterprises' emergency preparedness and business resiliency (Public Law 110-53, Title IX, Section 524). The program is commonly referred to as PS-Prep (Private-Sector Preparedness). The purpose of this law is to improve the preparedness of the private sector in three areas: (1) disaster management; (2) emergency management; and (3) business continuity. The law sets out the following program requirements: (1) Select preparedness standards that meet the certification requirements established by DHS; (2) establish the accreditation and certification program and the associated accreditation process; (3) obtain stakeholder and public comment on the viability of the proposed standards; (4) make provisions for the hardship placed on small businesses that want to or need to obtain certification; and (5) account for sector-specific emergency management and business continuity needs, especially if members of those sectors have already obtained certification under a different program.

Position and Adoption Speed Justification: The PS-Prep certification is for U.S.-based private enterprises and is completely voluntary. U.S. federal government agencies are mandated through continuity of operations to have a recovery program in place. There is no indication from DHS that PS-Prep certification will become part of a future government contract for securing services from vendors. Non-U.S. private enterprises may be asked by their supply chain partners to obtain PS-Prep certification if doing business in the U.S. or with a U.S.-headquartered enterprise.

PL 110-53, Title IX does indicate that credit may be granted toward PS-Prep certification for enterprises that have fulfilled industry-level availability requirements, such as financial services. However, these enterprises will have to wait until DHS issues its mapping of PS-Prep requirements to industry-specific availability requirements to see how much, if any, of their

compliance solutions industry requirements will apply to PS-Prep certification. DHS is actively working on this mapping at the time of this writing.

We position PL 110-53, Title IX, Section 524, at the Trigger-Peak midpoint for a number of reasons: (1) DHS got off to a slow start in rolling out the PS-Prep program — it assigned the ANSI-ASQ National Accreditation Board (ANAB) as the accreditation body after the 210-day deadline outlined in the law; and (2) it was only in October 2009 it announced the three proposed standards: NFPA 1600 2007, ASIS SPC.1-2009 and BS 25999-2. The public comment period has just ended, and it will take DHS months to review the comments and make its decision about the final standards that will be part of PS-Prep. (3) It will take ANAB many months to roll out the accreditation program for the auditors who will actually perform the PS-Prep certification audits. (4) The uptake by private enterprise will be slow because most organizations' business continuity management (BCM) programs aren't mature enough to warrant certification. There will be much added focus on BCM programs though, because the supply chain might pressure firms to receive certification. This market pressure would be more palatable to private enterprise than government regulation if this certification program was not voluntary.

User Advice: Go slowly in adopting PS-Prep for these reasons:

- Few organizations have evolved their recovery programs to a point where they could consider a publicly acknowledged level of maturity.
- There is no gradation of passing the certification — it is pass/fail.
- No case law yet exists on the legal implications of an organization obtaining certification and then failing to recover from a disaster, or not maintaining the certification.
- Follow the PS-Prep program development, even if you have no plans to be certified in BCM at this point in time. You may be asked to be certified in the future by your supply chain.
- Assess your BCM program against ASIS SPC.1-2009, BS 25999-2 and NFPA 1600 2007.
- Decide if organizational certification is appropriate for your organization. Review requests received from customers and trading partners during the past three years to see if the need is sufficient to make the investment.
- Determine what and where within your organization would require organizational certification — for example, which products, services and locations will need to be certified.
- Prioritize the organization certification schedule, based on the largest revenue provider or regulatory requirement.
- If your organization has already been reviewed for disaster preparedness and that review has been used successfully by a customer or trading partner to assess your organization's ability to meet its recovery requirements, then ask future customers, trading partners and prospects to use the same review. There is no reason to waste money in getting a review every time you are asked for one if you have already proven your capabilities. In the PS-Prep process, this reuse process is referred to as an "attestation."
- Small and midsize businesses (SMBs): Ask your supply chain partner how it can support your organization's ability to meet the recovery requirements (PS-Prep or other) you are being asked to comply with. This support can come in several forms: direct recovery

financial support or angel funding to company operations as a whole (which is likely available only for SMBs that are providing a unique product or service to the partner); expert guidance through the transfer of knowledge between the partner's recovery staff and the SMB; or a combination of both.

- Nonregulated U.S.-based organizations: Follow the work being done in relation to U.S. Title IX to understand how it might influence your models.
- Use the Gartner BCM Activity Cycle and maturity self-assessment tool to help manage your own BCM programs, as well as assess the current level of recovery or preparedness maturity. Build a road map to improve maturity during the next five years.

Business Impact: Anytime is a good time for an organization to demonstrate its ability to recover from a disaster. The worst position is to not be prepared. Having more U.S. organizations certified by their emergency preparedness and business continuity programs means that the nation as a whole is better able to respond, recover and restore business operations after a natural or man-made disaster; and that lives and livelihoods are protected.

However, a number of costs are associated with obtaining certification. (1) Costs associated with technology, recovery facilities and consulting services to ensure your recovery strategies and solutions meet business availability requirements: If you don't meet the current business availability requirements, do not even attempt certification. (2) Costs for third-party reviews of your BCM program: Organizations interested in obtaining certification should have an independent, third-party review of their program before starting down the certification path. (3) Costs associated with the certification itself, including fees for hiring an accredited auditor to perform the certification audit, and program management investments needed to obtain and maintain the certification.

Benefit Rating: High

Market Penetration: Less than 1% of target audience

Maturity: Embryonic

Data Dependency Mapping Technology

Analysis By: John Morency; David Russell

Definition: Data dependency mapping products are software products that determine and report on the likelihood of achieving specified recovery targets, based on analyzing and correlating data from applications, databases, clusters, OSs, virtual systems, networking and storage replication mechanisms. These products operate on direct-attached storage (DAS), storage-area-network (SAN)-connected storage and network-attached storage (NAS) at the primary production and secondary recovery data centers.

Position and Adoption Speed Justification: Without these solutions, there were only two ways of determining whether a specific recovery time objective (RTO) could be achieved. This was done through data restoration testing or through operations failovers that were conducted during a live recovery test exercise. A frequent outcome of the live test was the discovery of missing, nonsynchronized or corrupted data that was not detected during normal backup, asynchronous replication or synchronous mirroring, resulting in unplanned losses of data that could potentially cause disruption in one or more business processes.

Because of the high cost incurred to discover and remediate these problems, as well as the costs incurred from re-exercising the test, a new generation of data assurance technologies has been developed to create more-granular knowledge of application-specific data dependencies, as well

as the identification of possible content inconsistencies that result from application software bugs or misapplied changes (both largely attributable to human error and the complexity and dynamic nature of the IT environment).

One technology approach that's being taken by a number of vendors is the use of well-defined storage management problem "signatures" supported by industry-standard storage and data management software, in combination with the passive traffic monitoring of local and remote storage traffic (through software-resident agents). This traffic monitoring is used to detect control and user data anomalies and inconsistencies in a timelier way, notifying storage operations staff of the occurrence of the issue, and also to project the negative RTO effect through onboard analytics. The automation of the verification process and an attempt to quantify the impact on the business are the key deliverables of these solutions.

In 2010, Gartner adjusted the "Time to Plateau" forecast to "five to 10 years" from the "two to five years" that had been forecast in the 2009 Hype Cycle report. The main reason why is because, based on the number of direct client inquiries, the market does not appear to be evolving as quickly as Gartner had originally expected. Another contributing factor is that data dependency mapping products are still primarily offered on a stand-alone basis, as opposed to being bundled as part of larger storage management or backup solutions. This is also contributing to the slower-than-expected adoption rates.

User Advice: Preproduction pilots are a viable option that may be worth pursuing. In some cases, vendors (such as 21st Century Software and Continuity Software) offer very-low-cost pilot projects in which the vendor software auto-discovers and reports potential problems that may previously have been unknown to the storage and server operations staff. The validation of the solution in the actual environment lowers the risk that the solution will not meet the organization's needs, and the outcome of the pilot can be used to justify the purchase of the solution.

Some vendor products (such as InMage's DR-Scout) go one step further through the use of proprietary multivendor data checkpoint and replication technology. As a result, a much richer degree of data integrity and consistency assurance can be supported across primary production and secondary recovery data centers. In addition, data checkpoint, replication and validation support for application-platform-specific software, such as Microsoft Exchange, BlackBerry Enterprise Server and Microsoft SharePoint, constitute additional product alternatives for users whose immediate recovery needs are far more specific.

While some vendors (for example, 21st Century Software) have been offering production quality products for over six years, the overall maturity of this technology is still at a very early stage. Nonetheless, vendors such as EMC and NetApp have acquired companies that offer these capabilities because of their long-term potential to further automate storage management. Other large vendors, such as Symantec, have rebranded this software (in Symantec's case, from Continuity Software), which has led Gartner to adjust the data dependency mapping Hype Cycle curve positioning to reflect technology adoption and use by large management (Symantec) and storage vendors (EMC and NetApp), as well as by at least one prominent service provider (SunGard).

Business Impact: The primary benefit attributable to this technology is improved predictability and efficiency for reaching critical RTO targets, especially for mission-critical Tier 1 and Tier 2 applications. This is especially important for high-volume transaction applications, or for low to medium transaction applications for which the average revenue per transaction is high. In recent years, recovery exercising has become much more limited, making such tools even more valuable in avoiding data loss.

Benefit Rating: Moderate

Market Penetration: Less than 1% of target audience

Maturity: Emerging

Sample Vendors: 21st Century Software; Continuity Software; EMC; InMage; NetApp; Symantec

Continuous Availability Architectures

Analysis By: Donna Scott

Definition: Continuous availability or resiliency architectures enable 24/7 access to IT-enabled business functions, processes and applications, despite unplanned outages or planned maintenance. Continuous availability involves two strategies: high availability (minimizing unplanned downtime) and continuous operations (minimizing planned downtime, such as maintenance and upgrades). Sometimes, continuous availability architectures embody disaster recovery strategies, which means that the IT service requires continuous processing, despite disaster events. Gartner refers to this as "multisite continuous availability."

The inclusion or exclusion of disaster events in the definition depends directly on the overall business continuity strategy (and business process recovery and resumption plans). Continuous availability is measured from the user's perspective, meaning that users perceive 100% availability, even though IT service components may fail or have maintenance performed on them. The IT service can be hosted internally or at an external service provider (e.g., outsourcer, software-as-a-service [SaaS] provider or cloud-computing provider) location, or a combination of both. However, the IT organization is typically responsible for ensuring the overall IT service levels required, including the coordination of architecture, service levels, integration and processes across the multisourced environment. Continuous availability is usually implemented via active/active architectures, where multiple sites participate in IT service processing (so that processing can continue even if a site or its components is down), or where processing occurs primarily in one site and the other site takes over quickly and with minimal or no user impact during an outage or disaster.

Designing for no planned downtime requires extra planning, so that upgrades and maintenance can be performed while the IT service is operational. In addition to business and IT service architecture (and the level of resilience to outages and maintenance activities), the IT operations management architecture plays a significant role in achieving high levels of availability. Organizations need mature processes (for example, in incident, problem, change, release and configuration management) to achieve high levels of availability.

Position and Adoption Speed Justification: The trend toward automating business processes to achieve competitive advantage suggests that businesses and customers expect IT services to be available 24/7. Although overall service quality has improved — in part, based on IT management process maturity — service availability still isn't "good enough" for most enterprises, because the cost of downtime continues to escalate. The pressure is on for business and IT services to have 100% availability, not some number of nines, even during disasters. Despite businesses' desire for 100% availability and "availability as a utility," IT is not a utility. It lacks standards, the architectures are complex and interdependent on many components, and the level of people and processes in IT service delivery increases the risk of downtime versus technology risks.

Most enterprises approach availability in an opportunistic way, after they have put an IT service in production. However, achieving 100% or near-100% availability requires a top-down approach to engineering it into the application, infrastructure and operating architectures. Moreover, the greater the availability requirements, the more customization and integration are required across multivendor components in the solution. Thus, achieving continuous availability is an expensive

proposition reserved for the most critical business and IT services, such as those affecting life and safety, significant revenue generation, customer service and critical infrastructure protection (for example, payment processing/clearing systems).

This is the justification for continuous availability's location in the early part of the Hype Cycle. Less-critical applications rely on lower service levels at lower costs (and, often, are implemented for mainstream IT services). We estimate that fewer than 1% of IT services will be designed for continuous availability through 2014, because of cost and complexity. Continuously available IT services that exist today are constrained to approximately 5% of large enterprises.

Every hardware and software vendor offers some type of feature for high availability at a minimum (e.g., redundancy and failover), but not all offer features to enable continuous operations. The challenge for enterprises is to take high-availability components designed independently and integrate them seamlessly across the IT service. This is no simple task. Moreover, designing for continuous operations often requires it to be engineered into the product (including the supporting infrastructure), meaning that workarounds may not be otherwise available. Thus, continuous operations may not be possible without rearchitecture.

User Advice: Determine availability requirements at the design phase for new business and IT services, and update them annually through the business impact assessment process. Justify higher levels of availability (and costs) based on the business impact that outages will have on your enterprise. Not all IT services require the same level of availability. Enterprises should classify their applications based on availability, and should plan to invest more for availability that must occur for higher-level requirements.

Business Impact: The business impact of downtime can be significant for some business processes, such as those affecting revenue, regulatory compliance, customer loyalty, health and safety. Where the business impact is significant, enterprises should invest in continuous availability architectures.

Benefit Rating: High

Market Penetration: 1% to 5% of target audience

Maturity: Early mainstream

Recommended Reading: "How to Calculate the Cost of Continuously Available IT Services"

"How to Assess Your IT Service Availability Levels"

"How to Design Software for Continuous Operations"

At the Peak

Cloud Storage

Analysis By: Adam Couture; Stanley Zaffos

Definition: Cloud storage is a storage utility offering that is defined by the following characteristics: pay-per-use model, software-agnostic, reservationless provisioning and provider-owned; it is also frequently geographically separated from the servers that are using it. For the purposes of this Hype Cycle, we focus on the enterprise, rather than the consumer community, and exclude software as a service (SaaS) because the storage associated with it is unavailable for other applications.

Position and Adoption Speed Justification: Cloud storage currently has several iterations available on the market. Its evolution is being driven primarily by market demand for low-cost

storage alternatives. We expect this segment to peak in about five years as the larger established vendors that have recently entered the market capture early and mainstream adopters of new technologies and as emerging companies reach sustainable revenue growth rates. Full-scale adoption is not likely to occur earlier due to as-yet unanswered security issues, which could lead to potential legal exposure. Unpredictable monthly costs due to usage variability and the current lack of sufficiently differentiated storage services will also affect the absorption rate of the market.

User Advice: Evaluate cloud storage as a low-cost option for certain applications such as archiving and backup because they generally are not classified as mission-critical workloads, and they tolerate relatively long latencies better than transactional workloads; additionally, security requirements may dictate the cost and management expense of data encryption, both in flight and at rest. Due diligence should also include an evaluation of the organization's key management strategy and potential vendors' service and support capabilities, including monitoring and resolving issues and customer satisfaction. Considerable investments in time and money will often be required to integrate cloud storage options into current applications and environments.

Business Impact: The cost and agility expectations set by the public cloud storage vendors are forcing in-house IT operations to change their storage infrastructure management procedures and storage infrastructure strategy. User demands for lower costs, more agility and operations that are more autonomic are also influencing vendor R&D investments and cloud service offerings. Those services that have already been influenced by user demands include: backup, versioning, encryption and secure erasure. And in response to cost concerns, vendors are offering a variety of different pricing models that allow end users to align their storage costs with their usage rates, with the goal of lowering costs in the short and long term.

Benefit Rating: High

Market Penetration: 1% to 5% of target audience

Maturity: Adolescent

Sample Vendors: Amazon; EMC; Iron Mountain; Mezeo Software; Nasuni; Nirvanix

Recommended Reading: "The Storage Utility: From Outsourcing to the Cloud"

"SNIA Launches First Cloud Storage Standard"

"Cloud Storage: Benefits, Risks and Cost Considerations"

Virtual Machine Recovery

Analysis By: David Russell

Definition: Virtual machine (VM) recovery focuses on protecting and recovering data from VMs, as opposed to the physical server that the VMs run on. Server virtualization for the x86 platform from vendors such as VMware, Citrix and Microsoft is gaining considerable attention, and the deployment of x86 VMs is roughly doubling every year (the market is dominated by VMware). VM recovery solutions help recover from problems including user or application administrator error (such as the accidental deletion or overwrite of a file), logical errors (such as viruses), physical errors (such as disk failures) and disaster recovery (such as site loss). They can offer improved granular recovery of data/files in a VM environment.

Position and Adoption Speed Justification: The rate of server virtualization technology adoption is increasing, and the workloads being virtualized are becoming more mission-critical. As more production data is housed in or generated by VM environments, the need to protect data

in these environments is increasing. Recoverability of the virtual infrastructure is a new and significant component of an organization's overall data availability and disaster recovery plan.

The adoption, speed and time to plateau will vary, depending on the server platform that is virtualized. The x86 platform is experiencing the most growth and has the largest ecosystem of supporting vendors with value-added products and services.

Although VM-specific recovery tools showed increased adoption in 2009, these point solutions are likely to give way to more-generalized tools. These solutions could be extensions of current recovery applications and tools.

User Advice: Many products and methods are available for VM data recovery. Most traditional backup applications can install their agents in VMs, which may be acceptable in a small deployment. As the number and mobility of VMs increases, more-advanced backup — such as VMware Virtual Consolidated Backup, the recently announced vSphere application programming interfaces (APIs) or block-level incremental capabilities —should be considered, as well as snapshot, replication and data reduction (including data deduplication) techniques, and deeper integration with the server virtualization provider. With hundreds to thousands of VMs deployed in the enterprise, and with 10 or more mission-critical VMs on a physical server, improved data capture, bandwidth utilization, and monitoring and reporting capabilities will be required to provide improved protection, without complex scripting and administrative overhead.

This is a nascent, but fast-growing, market. Continually re-evaluate your options, especially if you decide to invest in a point solution that is different from the rest of the recovery tools that are deployed for the physical environment. Server virtualization vendors may eventually provide more-robust protection capabilities, and the traditional physical server protection vendors will expand their support for VM recovery within their standard products. In addition, expect nontraditional storage management companies that are creating virtualization management tools to expand their capabilities to include data protection and disaster recovery for blended virtual/physical environments.

Business Impact: The ability to protect and recover VMs in an automated, repeatable and timely manner is important for many organizations. As server-virtualized environments become pervasive and are used for more business-critical activities, VM recovery will become necessary to ensure timely access to data and continuation of business operations. Gartner estimates that, by the end of 2010, approximately 25% of enterprise workloads will reside on VMs.

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Sample Vendors: Acronis; Asigra; Atempo; BakBone Software; CA; CommVault; Double-Take Software; EMC; FalconStor Software; HP; i365; IBM; InMage; Microsoft; NetApp; Novell; PHD Virtual; Quest Software; Symantec; Syncsort; Veeam

Recommended Reading: "Backup and Recovery in a Server-Virtualized World"

"Enterprise Backup/Recovery Market Update: Change Driven by Virtualization and Data Reduction"

"Competitive Landscape: Enterprise Distributed Backup/Recovery Software Growth Driven by Virtualization and Data Reduction"

Risk Assessment for BCM

Analysis By: Tom Scholtz

Definition: Risk assessment in the business continuity management (BCM) context is the process of identifying and treating risks to business process availability and the continuity of operations. It is an essential first step (along with the business impact analysis) in the overall BCM process, and is necessary to reduce the frequency and effect of business interruptions, addressing risks related to technology, location, geopolitics, regulations, industry, as well as the business and IT supply chains. Whereas the business impact assessment (BIA) is the process of identifying and evaluating the impact of given events, risk assessment is the process of evaluating the likelihood of the event occurring in the first place. While a BIA typically includes a mini-risk assessment (e.g., a location's generic risks), a fully fledged, detailed assessment of the likelihood of, for example, a fire, pandemic, earthquake, hurricane or security breach happening to the organization often is not conducted.

Position and Adoption Speed Justification: Unfortunately, BCM planning is often conducted at a very superficial level of risk assessment, or even none at all. Although it has been well-understood that risk assessments are a necessary component of BCM planning, the line of business often considers them to be time-consuming and too resource-intensive. This opinion has been justified, given the general lack of effective risk assessment methods and tools, and often exacerbated by the inappropriate use of such tools and methods. However, expectations of better levels of practice are increasing, encouraged to some extent by standards such as Information Technology Infrastructure Library (ITIL), Control Objectives for Information and Related Technology (CobiT), International Organization for Standardization (ISO)/International Electrotechnical Commission (IEC) 27001, British Standard (BS) 25999 and PS-Prep.

Today, risk assessments are recommended in almost every BCM framework, and risk assessment tools are being included as integrated or stand-alone modules in BCM toolsets. Using these tools still requires specific BCM skills and time, which often are not available, but this situation is improving. Increasing emphasis on the importance and value of risk assessment in all spheres of business management is driving increased adoption of the discipline as a key component of BCM. However, unrealistic expectations about risk assessment being a panacea for ensuring business involvement in the BCM process, coupled with the inappropriate use of risk assessment tools (e.g., using very algorithmic, mathematical models with a business audience that manage risk in a more-intuitive manner), will result in some disillusionment and a lack of business unit buy-in.

User Advice: Make formal risk assessments that identify key control weaknesses and single points of failure mandatory components of your BCM program. Define the extent to which risk assessments will be performed based on BCM project scope, resources and time availability. If existing processes are not effective, then change them. Consider replacing complex mathematical tools with more-intuitive assessment methods (e.g., scenario planning, Delphic brainstorming), if it will better suit the cultural approach to risk management. Improve efficiency and reduce the time demands on business managers by leveraging risk assessments performed by operational or IT risk teams. Work with those teams to ensure that their data is sufficiently granular to meet BCM needs. As you become more mature at BCM risk assessment, make the transition to a continuous improvement process that accommodates BCM, IT and security risks. This will ensure that BCM team members — business and IT — are included and kept apprised of new or changing threats. Use standard terminology and processes to ensure consistency in assessment and risk prioritization. Investigate the use of software tools. They will not eliminate the need for an experienced risk assessor, but they can simplify the risk assessment process. Additionally, they provide an important repository for risk information, tracking assessments and treatment activities, as well as documentation for auditors and aid to program improvement. BCM

planning (BCMP) tools, which often provide integrated risk assessment functionality, are increasingly being used as hosted/software-as-a-service solutions. This potentially allows the business continuity manager to realize value at a lower price entry point.

Business Impact: Implementing BCM plans can be expensive and disruptive. Risk assessments are essential for pre-emptive action to reduce threat occurrences and constrain the effect of any disaster. Risk assessments ("What are the chances of a disaster happening?") also provide critical information for effective BIAs ("What will the impact be if a disaster becomes reality?").

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Sample Vendors: Archer; part of RSA; The Security Division of EMC; CPACS; eBRP Solutions; Fusion Risk Management; Linus Information Security Solutions; RiskWatch; Siemens Enterprise Communications; Strategic BCP; SunGard Availability Services

Mobile Service-Level Management Software

Analysis By: Monica Basso

Definition: Mobile service-level management products enable enterprises to increase the resilience of their mobility deployments and provide appropriate service levels for mobile users, while implementing more cost-effective operations. Wireless e-mail is the area with the highest demand, because e-mail is a mission-critical application, and users do not tolerate service faults and outages.

Position and Adoption Speed Justification: Organizations with midsize to large wireless deployments face the challenges of service resilience and business continuity. Scaling up deployments while maintaining the appropriate service quality and performance is an additional challenge. Software products, such as those sold by BoxTone, Zenprise and others, support a range of techniques to deal with these challenges, including:

- Mobile-user real-time monitoring and management
- Mobile-user self-service
- Troubleshooting
- Capacity management/load balancing
- High availability
- Disaster recovery
- Failover/failback

Growing size and complexity of mobility deployments force organizations to acquire appropriate management capabilities. As operational costs represent about 60% of mobility costs, IT organizations will increasingly invest to acquire tools for containing these costs. Organizations begin to invest in tools for optimizing wireless e-mail and application deployments, increasing service resilience and quality for IT users more cost-effectively.

We expect demand to grow significantly for tools that support better service levels to users, such as real-time monitoring, user self-service and support, and cost optimization tools, as well as for

techniques and tools implementing data center disciplines, such as high availability, fault tolerance, disaster recovery, load balancing and capacity management.

Some wireless e-mail platforms already offer some of these capabilities natively. Research In Motion's BlackBerry Enterprise Server offers some monitoring, reporting, alerting and troubleshooting capabilities, as well as high availability. Microsoft Exchange 2007 offers high availability as well. However, some wireless e-mail platforms are limited because they tend to focus on single nodes or capabilities instead of taking a holistic approach (for example, in the BlackBerry support, there is no help to deal with network operations center [NOC] failures and consequent service outages).

Mobile service-level management products tend to provide capabilities that are complementary to those provided by mobile device management and telecom expense management tools. A few vendors, such as Mobile Iron and Good Technology, have offerings that span across these areas.

Much fragmentation exists in this market, and tools tend to focus on a subset of capabilities (with either a user focus or a data center focus). Some vendors are partnering with mobile operators to offer services to the enterprise market (for example, Neverfail Group and Vodafone in the U.K.).

We expect to see consolidation among these vendors. More competition will come from vendors in related markets, such as enterprise wireless e-mail platforms (for example, Research In Motion), IT outsourcing services (for example, HP, IBM and CSC), and mobile operators active in enterprise markets. Mobile device management and PC management life cycle vendors (for example, Microsoft, Sybase, LANDesk, BigFix, HP and Capricode) might also expand in this area, possibly through acquisition of one of these vendors. However, these point solution vendors will drive innovation in this area during the next 18 months.

User Advice: IT organizations investing in mobility, particularly wireless e-mail, should evaluate service levels required by mobile users, identify potential factors of risk to interrupt or downgrade those services, choose mobile service-level management tools that allow them to cope with the risk and minimize the impact of rapid scaling deployments, faults and outages on end users.

Business Impact: This technology can be used to improve user support (for example, through self-service portals or by troubleshooting problems rapidly), to control costs and to optimize usage of server resources in the data center.

Benefit Rating: High

Market Penetration: 1% to 5% of target audience

Maturity: Adolescent

Sample Vendors: BoxTone; Conceivium Business Solutions; Devinto; Good Technology; InterNoded; Mobile Iron; Neverfail Group; Zenprise

Recommended Reading: "The Five Phases of the Mobile Device Management Life Cycle"

"Cool Vendors in IT Operations Management, 2007"

"Latest RIM Outage Shows That Customers Need a Backup Plan"

Sliding Into the Trough

DR Insourcing

Analysis By: John Morency

Definition: Disaster recovery (DR) insourcing is the use of company personnel and resources working with or without DR provider services to recover from technology and site failures.

Position and Adoption Speed Justification: Due to the large numbers of technology, risk and cost management trade-offs that change fairly frequently, DR needs to be managed more as a continuous life cycle process, rather than as a once- or twice-a-year, stand-alone test ritual. The management of key life cycle deliverables (such as business impact assessments, DR plans, recovery infrastructures, and application and data recovery processes) and their associated interdependencies (including technology enablement, execution management and process definitions) has resulted in the creation of a more-complex management challenge. There are three options for managing the creation of the key deliverables, as well as the resolution of their related technology, management and process interdependencies:

- The IT organization can manage deliverables creation and interdependency resolution on its own, if that is the most-cost-effective option.
- All deliverables creation and interdependency resolution can be entirely outsourced to a third party, including cloud service providers.
- Multisourcing (i.e., supplementing the efforts of the in-house recovery management team with services from one or more external providers) can be used to partition management responsibilities. This approach is increasingly being used for the recovery management of mainframe and open systems. Because of cost and/or logistics reasons, some clients are electing to continue their relationships with external service providers for mainframe recovery, while choosing to insource the recovery management of open systems (Windows and Unix/Linux platforms), leveraging colocation and hosting services, as well as (in some cases) building out an internal facility.

In terms of recovery-facilities-specific insourcing, the technologies that can most impact its adoption are public application and storage cloud services, because they represent a service delivery paradigm that is completely different from the traditional shared DR service model. However, these technologies are still at an early stage. Hence, Gartner is positioning the overall state of DR insourcing at the same point on the Hype Cycle curve as in 2009.

User Advice: A decision to insource a DR program is typically driven by one or more of the following factors: dissatisfaction with the quality and pricing of shared recovery services provided by an incumbent DR vendor; relatively inflexible long-term recovery service contracts; and the need for reduced recovery risk that drives the need for internal program management and a more-disciplined approach to exercising the recovery plan, rather than the simple execution of a once-a-year recovery plan exercising event.

If one or more of these factors are highly relevant to your recovery strategy and program, then you should evaluate the possibility of insourcing some or all of the following:

- DR program management
- Recovery data center management
- Recovery infrastructure management
- Recovery plan exercising

Insourcing any one of these items does not necessarily mean that the same approach should be taken for all four items. Successful insourcing decisions are based on a combination of cost, risk avoidance and operations management considerations.

Business Impact: The business impact is high, for business recoverability reasons.

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Early mainstream

IT Service Dependency Mapping

Analysis By: Ronni Colville; Patricia Adams

Definition: IT service dependency mapping tools enable IT organizations to discover, document and track relationships by mapping dependencies among the infrastructure components, such as servers, networks, storage and applications, that constitute an IT service. The primary use of these tools is for applications, servers and databases; some also discover network devices (such as switches and routers), storage devices, mainframe-unique attributes and virtual infrastructures, thereby presenting a more-complete service map. New enhancements to these tools include support for virtualized environments, enabling tracking of virtual servers and their relationships to physical and other virtual systems. Some tools offer the capability to track configuration change activity for compliance.

These tools use agent-based or agentless capability to discover networked assets, and build a map view of the interrelationships of the various components. Vendors in this category provide sets of blueprints or templates for the discovery of various packaged applications and infrastructure components. IT organizations can use a software development kit (SDK) to create desired-state blueprints or templates of internally developed or custom applications, which can then be discovered with the IT service dependency mapping tools.

Position and Adoption Speed Justification: Prior to the emergence of IT service dependency mapping tools, enterprises struggled to maintain an accurate and up-to-date view of the dependencies across IT infrastructure components, relying on data manually entered into Visio diagrams and spreadsheets, or not at all. There was no (near) real-time view of the infrastructure components that made up an IT service or how these components interrelated. Traditional discovery tools provide insight about the individual components and basic peer-to-peer information, but they did not provide the necessary parent/child hierarchical relationship information.

During the past five years, the marketplace has undergone consolidation, and there are no longer any stand-alone IT service dependency mapping vendors. The last was acquired in late 2009. The dependency mapping vendors have typically been bought by vendors that offer configuration management database (CMDB) tools as a means of jump-starting the data population of the CMDB with the IT service or application service models. However, for many enterprises, these tools still fall short in the area of homegrown or custom applications. Although the tools provide an SDK to develop the blueprints, this task remains labor-intensive, which will slow the enterprisewide adoption of the tools beyond its primary use of discovery. For the first several years since their acquisitions, IT service dependency mapping tools did not maintain development investment to keep pace with new technology requirements (for example, virtualization), but as IT organizations have matured in their use of CMDBs and the requirements for dependency mapping, the vendors have prioritized investment in this area. IT service dependency mapping tools still require expanded functionality for breadth and depth of discovery (such as broad range of storage devices and mainframe); however, they do offer dramatic improvements, compared with the prior manual methods. The adoption of these tools has seen an increase in the last 12 months, because new stakeholders (for example, disaster recovery

planners) and business drivers (for example, data center consolidation and migration projects) have emerged. Therefore, market awareness and sales traction have improved.

User Advice: Evaluate IT service dependency mapping tools to address requirements for configuration discovery for servers and applications. The tools should also be considered as a precursor to CMDB initiatives. If the primary focus is for an IT service view, then be aware that if you select one tool, the vendor is likely to try to "thrust" its companion CMDB technology on you. If the IT service dependency mapping tool you select is different from the CMDB, then ensure that the IT service dependency mapping vendor has an adapter to integrate and federate to the desired or purchased CMDB.

These tools can also be used for other initiatives, such as business service management, configuration management, business continuity, application management and other tasks that benefit from a near-real-time view of the relationships across a data center infrastructure. Although most of these tools aren't capable of action-oriented configuration modification, the discovery of the relationships can be used for a variety of high-profile projects in which a near-real-time view of the relationships in a data center is required, including compliance, audit, disaster recovery and data center moves (consolidation and migration). IT service dependency mapping tools can document what is installed and where, and provide an audit trail of configuration changes to a server and application.

Business Impact: These tools will have an effect on high-profile initiatives, such as CMDB, by establishing a baseline configuration and helping to populate the CMDB. IT service dependency mapping tools will also have a less glamorous, but significant, effect on the day-to-day requirements to improve configuration change control by enabling near-real-time change impact analysis and providing missing relationship data that's critical to disaster recovery initiatives.

The overall value of IT service dependency mapping tools will be to improve quality of service by providing a mechanism for understanding and analyzing the effect of change to one component and its related components within a service. These tools provide a mechanism that enables a near-real-time view of relationships that previously would have been maintained manually with extensive time delays for updates. The value is in the real-time view of the infrastructure so that the effect of a change can be easily understood prior to release. Using dependency mapping tools in conjunction with tools that can do configuration-level changes, companies have experienced labor efficiencies that have enabled them to manage their environments more effectively and improved stability of the IT services.

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Early mainstream

Sample Vendors: BMC-Tideway; CA Technologies; HP; IBM; VMware-EMC

Recommended Reading: "Selection Criteria for IT Service Dependency Mapping Vendors"

Appliance-Based Replication

Analysis By: Stanley Zaffos

Definition: Replication appliances provide network-based storage-vendor-neutral replication services that can include local clones, snapshots or continuous data protection, and synchronous or asynchronous block-level remote copy of protected volumes. Integration into the storage infrastructure can be via software agents, storage area network director or switch APIs, or direct storage system support.

Position and Adoption Speed Justification: Replication appliances provide common replication services across dissimilar storage systems. Offloading replication service overhead into the appliance preserves the native performance of storage systems and reduces the effort needed to switch storage vendors. Key impediments to appliance market share growth include the reluctance of end users to add another device to the input/output path, competition from storage virtualization appliances, storage vendor lock-ins, and channel limitations caused by the bias of storage vendors to protect their storage-based software revenue and the reluctance of indirect channels to directly participate in this market.

User Advice: Users should consider the use of replication appliances when there is: a need to provide common replication services across different tiers of storage; a need to create a constant timeline across multiple homogeneous or heterogeneous storage systems; a problem with the usability, performance or cost of the existing replication solution; a need to preserve investments in existing storage systems; or a desire to pursue a dual-vendor strategy.

Business Impact: Appliance-based replication services can:

- Provide the benefits of storage-based replication solutions, without the lock-ins that storage-system-based replication solutions create
- Delay storage system upgrades by offloading replication overhead from the storage system
- Provide heterogeneous replication targets to allow lower-cost solutions

Benefit Rating: Moderate

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Sample Vendors: BakBone Software; CA; Cisco; DataCore Software; EMC; FalconStor Software; Hitachi Data Systems; IBM; InMage; Oracle

Lights-Out Recovery Operations Management

Analysis By: John Morency

Definition: A "lights-out" operation refers to the management of a remote (and largely unmanned) recovery data center through the use of remote management software. Active management of the recovery data center may be to support recovery plan exercising, or to orchestrate a post-disaster event operations recovery. A key requirement for a lights-out operation is that remote management support must be uniformly supported across the heterogeneous computing, storage and network infrastructure components located at the remote recovery facility.

Position and Adoption Speed Justification: Currently, the principle enablers for lights-out operations are the use of keyboard, video and mouse (KVM) switches and hardware service processors that support the level of equipment console access required to support basic equipment configuration, event monitoring and change management. Service processors are hardware-vendor-specific and typically do not support all the models of any one server product vendor's product line. Given this limitation, recovery managers may be better-served by the use of remote KVM switches similar to those used to support server and storage equipment.

In the case of Solaris, Unix or Linux servers, the use of KVM switches can be supplemented with the privileged use of utilities, such as rlogin, rsh, Secure Shell (SSH) and many others. Physical or virtual remote 3270 access to mainframe applications and transaction subsystem management

commands has been a staple of the mainframe environment for many years. In addition, most of the remote configuration, change, incident, problem and performance management functionality for server, storage and network equipment (as well as business applications and data) can be supported by products from the "Big Four" system management vendors (that is, BMC Software, CA Technologies, HP and IBM), as well as management products from the major storage and network product vendors.

The services that can most impact the extent to which distributed operations can be managed with little to no in-house support staff are managed hosting and software as a service (SaaS). Given the fact that the supporting management software for these services continues to evolve at a fairly rapid rate, Gartner has moved the Hype Cycle curve position for lights-out recovery operations management to peak-trough midpoint.

For these reasons, the gap between user expectations and the reality of what can actually be supported by the vendors has lessened considerably over the past few years. Nonetheless, the need for remote staffing presence to manage equipment installation, configuration and basic problem triage still exists. This may not be an issue during recovery exercising, but it would be a real consideration during an actual recovery, especially one lasting several days or longer.

User Advice: Given that equipment malfunctions will always occur, independent of whether the equipment is locally or remotely located, it is generally necessary to have at least some on-site support presence (such as a facilities person) who is authorized to enter the data center; turn the equipment off and on; support requested moves, adds and changes; and escort equipment vendors who come to the remote data center to add, remove or repair equipment. In addition, it is generally in the IT organization’s best interest to have a technician with Level 1 (or maybe even Level 2) support skills to provide basic incident management, and to coordinate problem triage with the central IT operations team.

Lights-out operations in remote recovery centers are increasingly being used to reduce the costs associated with flying to a disaster recovery (DR) service provider location in another state, another country or even a remote portion of the same country. Additionally, it can reduce the time required to commence live recovery operations, because key support staff can perform recovery management either from their homes or from a recovery location that is much closer to the primary production data center location. In addition, DR service providers are improving the remote system and data access that they provide to a customer's recovery team in order to facilitate the management of recovery testing from the customer site, thereby reducing travel and entertainment (T&E) expense.

Because of the potential time and cost benefits, Gartner recommends detailed due diligence of the application of remote management technologies and processes in order to ensure that they fully support the recovery requirements of the business. In addition, if your existing DR service contract is coming up for renewal, ensure that your provider due diligence process includes the assessment of provider data center access to support remote recovery exercising.

Business Impact: The business impact is medium. The actual benefits will vary depending on the frequency with which live recovery exercising is performed each year.

Benefit Rating: Moderate

Market Penetration: 1% to 5% of target audience

Maturity: Adolescent

Sample Vendors: Avocent; Dell; HP; IBM; IBM BCRS; Minicom; Oracle; Raritan; SunGard Availability Services

Data Deduplication

Analysis By: David Russell

Definition: Data deduplication, also known as "data de-duplication" or "data dedupe," is a form of compression that eliminates redundant data on a subfile level, hence improving storage utilization. In this process, only one copy of the data is stored; all the other redundant data will be eliminated, leaving only a pointer to the previous copy of the data. Deduplication can significantly reduce the required disk space, since only the unique data is stored. It can offer greater savings than a single-instance store (SIS) method, because SIS can only eliminate a copy of that entire redundant file, while deduplication is more granular.

While deduplication is most commonly used in backup activities, this can be applied to many other use cases, such as long-term archiving and even primary storage, with file storage of unstructured data, in particular, most frequently considered.

Solutions vary in terms of where and when the deduplication takes place, which significantly can affect performance and ease of installation. When used with backup, deduplication that occurs on a protected machine is referred to as "client-side" or "source" deduplication, whereas deduplication that takes place *after* the backup server is called "target side" deduplication. There is also a distinction that is made between solutions that deduplicate the data as it is processed (which is called "in-line" deduplication) and products that have the data land on disk, as it would without deduplication, and then process it later, which is called "post processing" or "deferred" deduplication. Deduplication solutions also vary in how granular they are, but 4KB to 128KB chunks, or segments, of data are typical, and some deduplication algorithms are content-aware, meaning that they apply special logic for further processing, depending on the type of application and data that is being stored.

Position and Adoption Speed Justification: This technology reduces the amount of physical storage required, significantly improving the economics of disk-based solutions for backup, archiving and primary storage. Gartner clients using deduplication for backups typically report seven times to 25 times the reductions (7:1 to 25:1) in the size of their data, and sometimes higher than 100:1 for file system data or server virtualized images when data deduplication is used. To achieve the highest levels of reduction, backup workloads during a period of three to four months are typically needed, and a traditional model of nightly incremental backups and weekly full backups is also required; however, even short-term backups can achieve two times to three times the reductions.

Archiving deduplication ratios are often in the 3:1 to 10:1 range, and primary file data commonly yields 3:1 to 5:1.

Deduplication is a common topic in Gartner end-user inquiries, and the adoption of the technology for production deployments has been high for backup workloads, with increasing interest since 2009 for archive and primary data as well.

User Advice: Purpose-built deduplication offerings are now available as host-based client software, target-based software, and/or bundled software and a hardware appliance. In addition, this capability could also be part of an operating system, a network device (such as a WAN optimization controller), network-attached storage (NAS) filer, backup/restore software, archiving solution, disk subsystem or a virtual tape library (VTL).

Product implementations can vary, with some solutions supporting a limited matrix of operating systems, applications and backup software. In contrast, other implementations are completely independent of such things. Today, one vendor (CommVault) markets the ability to write deduplicated data to tape seamlessly, but others are expected to follow suit. Deduplication

solutions can also differ in their performance characteristics and scalability, so it's important to ensure that considered solutions meet current and anticipated requirements.

For backup workloads, carefully consider the architectures and design points mentioned above. Understand that client-side deduplication assumes that the backup software is switched to the deduplication software, which organizations are sometimes surprised to learn.

If deduplication is used for primary storage, then ensure that the workload matches the performance characteristics of the deduplication approach, because performance is sometimes negatively affected; however, not all data types require the highest levels of performance. With so many types of solutions, it's important to investigate several vendors and implementation architectures. Also, because this is an early, but fast-growing, market, expect additional vendors to enter the market, potentially offering several solutions. This means that architectures, road maps and pricing across the providers could vary widely.

Business Impact: The effects of deduplication primarily involve the improved cost structure of disk-based solutions (as less disks need to be purchased, deployed, powered and cooled). As a result, businesses may be able to use disks for more of their storage requirements, and retain data on disks for longer periods of time, thus enabling recovery or read access from disks versus retrieval from slower media (such as magnetic tape).

Backup-to and restore-from disks can improve performance, compared with tape-based approaches. The effective cost of replication can also be reduced if the data has previously been deduplicated, because potentially less bandwidth would be required to move the same amount of nondeduplicated data. Deduplication can also improve the cost structure for disk-based archives and primary storage, because fewer resources are utilized.

Benefit Rating: Transformational

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Sample Vendors: Asigra; CommVault; Data Domain; EMC; ExaGrid Systems; FalconStor Software; Hitachi Data Systems; i365; IBM; NetApp; Nexsan; NEC; Ocarina Networks; Oracle; Permabit; Quantum; Sepaton; Spectra Logic; Symantec

Recommended Reading: "EMC Acquires Data Domain, Becomes Deduplication Leader, Signals Deduplication as a 'Must Have' Capability"

"How EMC's Data Domain Buy Will Affect the Industry"

"Use Data Deduplication to Improve Availability and Lower Cost"

"Emerging Technology Analysis: Primary Data Deduplication, Storage Software and Hardware Technologies"

"Data Deduplication Is Poised to Transform Backup and Recovery"

"New Storage Solutions Can Modernize Data Life Cycle Management"

Hosted Virtual Desktops

Analysis By: Brian Gammage; Mark Margevicius; Ronni Colville

Definition: A hosted virtual desktop (HVD) is a full, thick-client user environment, which is run as a virtual machine (VM) on a server and accessed remotely. HVD implementations comprise server virtualization software to host desktop software (as a server workload), brokering/session

management software to connect users to their desktop environment, and tools for managing the provisioning and maintenance (e.g., updates and patches) of the virtual desktop software stack.

Position and Adoption Speed Justification: An HVD involves the use of server virtualization to support the disaggregation of a thick-client desktop stack that can be accessed remotely by its user. By combining server virtualization software with a brokering/session manager that connects users to their desktop instances (that is, the operating system, applications and data), enterprises can centralize user data and applications, and manage personalized desktop instances centrally. Because only the presentation layer is sent to the accessing device, a thin-client terminal can be used. For most early adopters, the appeal of HVDs has been the ability to "thin" the accessing device without significant re-engineering at the application level (as is usually required for server-based computing).

Early adoption was hindered by licensing compliance issues for the Windows client operating system, but that has been resolved through Microsoft's Windows Virtual Enterprise Centralized Desktop program. Beginning in mid-2010, Microsoft will reduce the license cost premium of the Windows operating system for HVD installations and expand roaming rights, enabling an HVD image to be accessed from multiple devices for a single license fee. However, other technical issues must still be resolved before mainstream viability is reached. Improvements in the complexity of brokering software and remote-access protocols will continue to occur through 2011, extending the range of desktop applications and users that HVDs can address.

Through 2011, broader manageability of HVD VMs will improve, as techniques to reduce HVD storage volumes (introduced in late 2008) lead to new mechanisms for provisioning and managing HVD images by segmenting them into more-isolated components (including operating systems, applications, persistent personalization and data). These subsequent manageability improvements will extend the viability of HVD deployments significantly beyond the structured task worker community, first to desk-based knowledge workers and then later to mobile/offline users.

Since late 2007, HVD deployments have grown steadily, reaching around 1.5 million at the end of 2009. Because of the constraints previously discussed, broad applicability of HVDs has been limited to specific scenarios, primarily structured task workers in call centers, and kiosks, trading floors and secure remote access; about 50 million endpoints is the current target population. These deployments will continue to expand through year-end 2010. General deployments should begin shortly thereafter, driven by the expansion of implementations to the broader user population. Inhibitors to general adoption involve the cost of the data center infrastructure that is required to host the desktop images (servers and storage, in particular), network constraints, availability of the skills necessary to manage virtual desktops and the limited potential for operational cost savings (the users to whom HVDs are currently most applicable are already those who cost least to manage).

HVDs promise to deliver diminishing marginal per-user costs, due to the high level of standardization and automation required for successful implementation, but this is currently only achievable for persistent users where images remain intact. As other virtualization technologies mature (e.g., brokers and persistent personalization), this restraint will be reduced. This creates a business case for organizations that adopt HVDs to expand their deployments, as soon as the technology permits more users to be viably addressed. However, the rate at which HVD technologies and products are improving also imposes cost penalties for early adopters. Enterprises that adopt HVDs aggressively will see later adopters achieve superior results for less cost, but will also need to migrate to new broker and complementary management software as products mature and standards emerge. This phenomenon is set to push HVDs into the Trough of Disillusionment in late 2010.

User Advice: Unless your organization has an urgent requirement to deploy HVDs immediately for securing your environment or centralizing data management, wait until late 2011 before initiating deployments for mainstream desktop user scenarios. Through 2011, all organizations should carefully assess the user types for which this technology is best suited. You will need to balance the benefits of centralized management with the additional overhead of the infrastructure and resource costs. Customers should recognize that HVDs may resolve some management issues, but they will not become panaceas for unmanaged desktops. In most cases, promised reductions in total cost of ownership will not be significant and will require initial capital expenditures to achieve. The best-case scenario for HVDs continues to be for securing and centralizing data management or for structured task users.

Organizations must optimize desktop processes, IT staff responsibilities and best practices to fit HVDs, just as organizations did with traditional PCs. Leverage desktop management processes for lessons learned. The range of users and applications that can be viably addressed through HVDs will grow steadily through 2011. Although the user population is narrow, it will eventually include mobile/offline users as well. Organizations that deploy HVDs should plan for growing viability across their user populations, but they should be wary of rolling out deployments too quickly. Diligence should be employed in testing to ensure a good fit of HVD capabilities with management infrastructure and processes. Visibility into future product road maps from suppliers is essential.

Business Impact: HVDs provide mechanisms for centralizing a thick-client desktop PC without re-engineering each application for centralized execution. This appeals to enterprises on the basis of manageability and data security.

Benefit Rating: High

Market Penetration: 1% to 5% of target audience

Maturity: Adolescent

Sample Vendors: Citrix; NEC; Parallels; Red Hat; VMware

Humanitarian Disaster Relief

Analysis By: Roberta Witty

Definition: Humanitarian disaster management is a complex process for the authority in control of the event: Organizations with specific disaster relief expertise from many locations need to be managed, and emergency responders and volunteers often arrive without official interface to the authorities and are not given specific deployment instructions, making their response chaotic. Free software and open-source software, as well as social media, have been key to delivering disaster relief for major disasters, such as the 2009 earthquake in Haiti.

Position and Adoption Speed Justification: The Hype Cycle position of pre-trough 30% was selected as the initial position for this profile because, even though the tools have been used successfully in many disasters, every time there is a disaster, coordination among all the participating humanitarian disaster relief suppliers requires an extensive effort. To further benefit the process of humanitarian disaster relief, the suppliers need to add more integration into their tools so that each implementation of the mashup tools is easier and faster. The emergency management organization in the country experiencing the disaster is typically the point of coordination, and it is at this point — in each country — where the process needs to be optimized.

User Advice:

- National emergency and disaster management personnel should investigate the use of open-source software for managing humanitarian relief aid and collaboration during a disaster.

- Local emergency management offices that run normal operations — for example, high altitude, missing person and shelter management — should integrate humanitarian disaster relief into their operations before an event occurs. Doing so will reduce the chaos of unmanaged search and rescue efforts.

- Humanitarian aid organizations will find humanitarian disaster relief useful for managing the intake and distribution of their own resources — for example, relief camp management and food distribution management.

- Business continuity management (BCM) professionals will find humanitarian disaster relief useful as a way to integrate all BCM tools and platforms for private enterprise disaster management, as well as to coordinate disaster management activities across multiple parties (private and public) during a large-scale (for example, regional) event.

- Academic institutions seeking to participate in open-source projects will find the open-source humanitarian disaster relief software market an opportunity for rich open-source interactions.

Business Impact: Managing humanitarian disaster relief is a challenge and can benefit from the use of automation so that emergency resources are managed in a manner that brings support to victims faster and mitigates further damage due to the recovery effort itself.

Benefit Rating: High

Market Penetration: 1% to 5% of target audience

Maturity: Adolescent

Sample Vendors: CrisisCommons; Facebook; FortiusOne; Google; InSTEDD; International Network of Crisis Mappers; MySpace; Open Solutions Group; Sahana Software Foundation; Star-Tides; Twitter; Ushahidi

Recommended Reading: "Sahana: Humanitarian Disaster Management and Collaboration System"

Crisis/Incident Management

Analysis By: Roberta Witty; Daniel Miklovic; Jeff Vining

Definition: Crisis/incident management is the process of managing multiple groups and workflow responses to a particular crisis/incident (earthquakes, fires, floods, collapsing bridges, severe weather conditions, terrorist attacks, chemical spills or accidental discharges), with a consistent and quick approach so as to return to normal as soon as possible. The goal of crisis/incident management is communications that minimize damage to individuals, localities, businesses and public agencies. Damage can be done to an organization's reputation, operations and revenue streams, as well as a government's ability to reduce any adverse impact on public safety.

In recent years, we've seen the commercialization of specialized crisis/incident management software tools designed for government, utility and private-enterprise use. These tools are used for the following purposes:

- To manage relationships with all organization stakeholders (internal and external)

- To manage crisis/incident/situation response, recovery and restoration actions
- To communicate information internally and externally
- Provide reports for postmortem reviews of the crisis/incident for regulatory reporting purposes and business continuity management (BCM) process improvement efforts

Solutions may be:

- Specialized to the operations of one industry — for example, electric utilities, transportation, or oil and gas.
- Generalized for management of any type of crisis/incident normally found in a BCM plan.
- Part of a larger solution, such as an environmental, health and safety (EH&S) application.
- Part of a case management tool. Many of these products are evolving into centralized "systems of record" and general risk-management tools.

Crisis/incident management software functionality should include:

- Emergency response planning
- A continuity plan repository
- Plan training/exercising capability
- Crisis/incident response- and action-tracking capabilities, including workforce, expenses and response actions
- Internal and community interaction support
- Interfaces to external data resources, such as a geographic information system, emergency notification and situational emergency message alerting systems
- Formal crisis/incident command reporting and processes

Position and Adoption Speed Justification: Government agencies and private enterprises in industries such as electric utilities, transportation, and oil and gas have embraced crisis/incident management processes and their supporting technologies to protect public safety and business operations, improve the efficiency of crisis/incident command and related emergency responses, and continually communicate and assess progress when responding to a disaster (for example, natural disasters, power failures, transportation accidents, a pipeline break or biohazard accident) that interrupts the delivery of goods and services. The St. Louis Area Regional Response System and New York State's Emergency Management Office NY-Alert and NY-Delivers are examples of government efforts.

An interesting dichotomy in managing the process of crisis management is that, in most organizations, alignment has not been made between managing an operational disruption, such as an electric utility responding to a power outage, and an administrative disruption normally covered under a BCM plan, such as a fire at the data center. Most recently, and due to events such as pandemics and large-scale disasters, such as hurricanes, organizations are reviewing how a supply chain interruption affects their abilities to recover and continue business services. Therefore, as institutional crisis/incident management maturity grows, Gartner projects that once-disparate crisis/incident management solutions will begin to show convergence of features, form

and functionality. Regional and national-scope disasters increasingly will require enterprise-based crisis/incident management for the critical infrastructure sectors to interact, at least at the level of status reporting and communicating with each other and government agencies. In addition, legal and regulatory agency requirements, such as the U.S. Occupational Safety and Health Administration (OSHA) and National Incident Management System/Incident Command System (NIMS/ICS), are driving more organizations to move to automation.

We've moved the position of crisis/incident management software to just past the midpoint of Peak of Inflated Expectations and the Trough of Disillusionment because firms are starting to use to these tools, but are finding them to be complicated to use or fit for purpose to only one standard, for example NIMS/ICS. More-flexible tools are required for market adoption to rise.

User Advice: Develop and document crisis/incident command process and procedures.

Use social networking technologies, such as Facebook, blogs, wikis and Twitter, to augment, but not replace, crisis/incident management software functionality.

Match the type of crisis/incident management software solution deployed to the most likely and critical types of crisis/incidents that pose the greatest operational risk to a company, based on a formal, board-approved risk assessment. A financial services company might opt for a solution that provides functionality aligned with an IT outage, a natural disaster or a pandemic, while a heavy industry manufacturing entity might choose one with functionality tailored for response to EH&S-related crisis/incidents.

Ensure that the chosen software solution adheres to public-sector crisis/incident protocols relevant in the geographic regions in which the solution is deployed. For example, in the U.S., any solution targeted to respond to physical crisis/incidents, such as environmental mishaps, safety issues, or natural disasters affecting health and safety, should adhere to the ICS, as mandated by the U.S. Department of Homeland Security. This will ensure interoperability with public-sector response agencies.

Manufacturers with exposure to EH&S issues as a result of disruptions caused by natural disasters should adopt solutions that are interoperable with regional public-service protocols to ensure timely and efficient responses to minimize brand damage.

Consult with corporate counsel for jurisdictional issues relating to privacy and rules of evidence.

Educate all staff on basic preparedness for themselves and their families, and develop communication plans to communicate with various government and nongovernment entities.

Business Impact: Crisis/incident management processes and software solutions help organizations manage all the actions taken in response to a disaster. Therefore, they do the following:

- Improve the organization's ability to protect public safety and restore business services as quickly as possible.
- Ensure recovery of expenses incurred during the disaster from business interruption insurance policies.
- Protect the reputation of the organization in the eyes of all stakeholders — employees, customers, citizens, partners/suppliers, auditors and regulators.

Using a system that imposes a standardized best-practice or leading-practice model extends uniform managerial controls across the organization, cuts staff training time, and ensures better integration with the broader internal and external community involved in recovering from a disaster.

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Sample Vendors: Archer Technologies (EMC/RSA); Cintellate; Coop Systems; Crisis Commander; Crisis Management Software; Dell; eBRP Solutions; Enablon; Enviance; ERMS; ESi; ESS (IHS); Global Alertlink; Intelex Technologies; Intergraph; IntraPoint; Ixtrom Group; MissionMode Solutions; National Center for Crisis and Continuity Coordination; Pier Systems; Preparis; Previstar; RecoveryPlanner.com; Send Word Now; Strategic BCP; SunGard Availability Services; Syntex Management Systems

Recommended Reading: "Turning EH&S Challenges Into Benefits"

"Toolkit: Requirements for Crisis Command and Emergency Operations Centers"

"Research Roundup: Business Continuity Management and IT Disaster Recovery Management, 3Q09"

"How to Understand and Select Business Continuity Management Software"

"Case Study: PetroChina Goes for World-Class EH&S Performance"

Business Continuity Management Planning Software

Analysis By: Roberta Witty; Tom Scholtz

Definition: Gartner includes the following three markets as part of the business continuity management (BCM) software marketplace (see "How to Understand and Select Business Continuity Management Software"):

- Business continuity management planning (BCMP) tools that provide risk assessment, business impact analysis (BIA) and plan management functionality (growth is moderate)
- Emergency or mass notification service (EMNS) tools that automate the call tree (growth is aggressive — see "MarketScope for Emergency and Mass Notification Services")
- Crisis/incident management (CIM) tools that let the organization manage the event once it's occurred (growth is embryonic)

BCMP products have been in the market for more than 20 years — growing from word-processing templates to sophisticated, interactive decision support tools. The increased need for usable recovery plans of all types (crisis management to damage assessment, emergency response, emergency notification, external communications, insurance support, travel support, procurement/vendor management, customer/partner support, shelter in place, IT disaster recovery, business recovery, business resumption, restoration and stand-down), as well as a consistent and repeatable plan development process, has resulted in increased sophistication in the products. In addition, they are more likely than in the past to integrate with other BCM tools, such as EMNS and CIM.

A BCMP tool will typically include the following components:

- Risk assessment for BCM
- BIA
- Business process and IT dependency mapping

- Information libraries of common business and IT equipment, processes, personnel and so forth
- Plan development and management
- Plan exercising
- Workflow management for plan development and maintenance actions

Some products include other features, such as:

- A modeling capability that lets the organization assess the impact on the business as a result of an outage — whether it be an application, location, business process and so forth.
- An incident management function that can be used when exercising and executing recovery plans (this function should not be confused with pure-play CIM tools in the market that provide a much richer set of functions than those described above).

Position and Adoption Speed Justification: The BCMP market has fairly low adoption — the 2009 revenue estimate is $76 million, with average revenue growth of 10% and 13% for 2008 and 2009, respectively. Large or regulated enterprises as well as government agencies typically use the tools, while small and midsize firms have not (but are increasingly looking to use them). The financial services market leads the pack in implementations and vendor marketing efforts. In our 2010 MarketScope, we rated this market as Promising.

The small BCMP market size is reflected in an average, and rather flat, 31% adoption rate of BCMP tools from our risk and security surveys during the past five years, 2005 through 2010.

We are quite surprised at the size of this market (very small) in relation to the importance of the use of these tools if an organization has a disaster from which it needs to recover. Having current and effective recovery plans is the key to success during a disaster, and these tools are "ground zero" for effective crisis and business recovery. We anticipate adoption to grow in the next five years, given the increased focus from government agencies — federal, state and local — as well as private-sector preparedness initiatives.

There are two different views of what a BCMP tool should provide. Some organizations want only a document management system that helps them manage their various recovery plans. These are typically firms without a strong commitment to BCM, those that have developed their own customized tools for their operations, those that do not have complex organizational structures and those that do not want to be forced into a particular BCM methodology. Others want a tool that allows them to enter and manage BCM data for the entire enterprise in a dynamic manner that allows for analysis of the data, the creation of the most-current recovery plans and integration into the day-to-day business operations — often with a strong commitment to enterprise risk management. Additional considerations regarding the use of BCMP tools include how many resources organizations have to apply to BCM, and the maturity of the BCM program. However, typically, a mature BCM program with dedicated staff is using BCMP automation because they have obtained long-term program commitment from management, and they have more-complex business operations — multiple operating locations, often on more than one continent. Therefore, coordinating, analyzing and managing large amounts of availability information become almost impossible to do without a tool.

Because there is a low level of market penetration, a large addressable market needs lots of convincing as to the need for these tools. Therefore, we moved back the Hype Cycle position from 10% post-trough to in the trough.

User Advice: Review Gartner's "MarketScope for Business Continuity Management Planning Software" to understand the BCMP market before you start your own implementation.

Consider a pure-play BCMP tool when:

- You are starting a new BCMP and want to follow standard practices throughout the organization.
- You are updating your current BCMP and processes.
- You need to integrate plans and partial plans from a number of departments and business units into one consistent, accessible and easily updated plan.
- A merger or acquisition has presented you with the need to create a BCM program reflecting all the elements of your organization.
- You want to conduct the research and planning process in-house, with minimum assistance from outside consultants.

Consider the BCM module from a governance, risk and compliance (GRC) product when:

- You already have purchased a GRC tool.
- You are starting an integrated, enterprisewide approach to operational risk management.

In all cases, choose a tool that matches your organization's complexity — do not overbuy. Focus on:

- Ease of use in the hands of business users (not IT users only)
- Ease of customization to your organization's branding and culture
- Ease of configuration for reporting, import/export capability to/from HR systems, student management systems, IT asset management systems, IT configuration management systems, IT service dependency mapping tools and so forth
- Integration with other BCM software that your organization may already have purchased (emergency notification or crisis/incident management)

Like all policies and procedures, even the best recovery plan can rapidly become obsolete. Consider the recovery plan a living document, and put in place a continuous process improvement process for regular (annually at a minimum) and event-triggered plan reviews (such as changes in operational risk profiles, business or IT processes, and applicable regulations, as well as exercise results showing a gap in plan actions versus current recovery needs). Besides an inaccurate 10K or 10Q, this is the one organization document most likely to result in lost revenue or damaged reputation if it is not current, or worse, not developed.

Business Impact: BCMP tools will benefit any organization that needs to perform a comprehensive analysis of its preparedness to cope with business or IT interruptions, and to have in place an up-to-date, accessible plan to facilitate response, recovery and restoration actions. If used to its fullest potential, a BCMP tool can be used to enhance business resilience in areas such as HR management, business re-engineering, mergers and acquisitions, and so forth.

Benefit Rating: Moderate

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Sample Vendors: Archer (EMC/RSA); Bold Planning Solutions; Business Protection Systems International; Coop Systems; CPACS; CPOtracker; eBRP Solutions; EverGreen Data Continuity; Factonomy; Linus Information Security Solutions; LiveProcess; netEOP; Paradigm Solutions International; RecoveryPlanner.com; Strategic BCP; SunGard Availability Services; Tamp Systems; Virtual Corporation

Recommended Reading: "MarketScope for Business Continuity Management Planning Software"

"Research Roundup: Business Continuity Management and IT Disaster Recovery Management, 3Q09"

"Best Practices for Conducting a Business Impact Analysis"

"How to Understand and Select Business Continuity Management Software"

Emergency/Mass Notification Software

Analysis By: Roberta Witty; John Girard; Jeff Vining

Definition: An emergency or mass notification service (EMNS) is the automated call-out to notify groups or individuals — disaster recovery teams, employees, citizens, residents, students/parents, customers, suppliers or government officials — and is critical for managing a crisis. In other words, these tools automate manual call processes (aka call trees). Not everyone is on duty when the incident occurs, but they must be notified to take action. EMNS offerings are tools focused on the electronic activation and management of notification messages, thus streamlining an organization's mass communications capability. The software can be used to organize contacts into an unlimited number of groups or subgroups; to send emergency messages (for example, announcing fires, power outages, natural disasters, severe weather conditions, volcanic ash events, terrorist attacks, hostage crises, bridge collapses, child abductions or criminal activity); and then to track receipts or responses for message delivery confirmation.

Activation can be accomplished by logging onto a Web portal; accessing the system by a telephone or by calling the vendor's call center; and then securely sending a custom or previously crafted voice or text message to multiple endpoint devices, such as phones, PDAs, desktops, e-mail systems, fax machines, physical security systems or public announcement systems. EMNS software can send thousands of messages to endpoint devices simultaneously. However, there is no guarantee that the person to whom the endpoint device belongs actually receives the message due to the telecommunications infrastructure being used, as well as recipient issues — for example, not being properly trained to use the service, information overload during an event, not wanting to participate in the service and so forth.

Position and Adoption Speed Justification: Critical incidents today range from localized events such as a fire or power outage to regional and catastrophic disasters such as earthquakes (Haiti and Chile), hurricanes/tsunamis and terrorist attacks. They don't have to cause major physical damage in order to have a major business interruption — for example, the 2010 Iceland volcanic ash event and the 2009 and 2010 H1N1 virus. As a result, organizations are increasingly implementing EMNS, thereby building a stronger crisis management program. The EMNS market is growing fast: 2009 revenue is estimated at $570 million, with median revenue growth of 37% and 29% for 2008 and 2009, respectively. Many vendors exist, and barriers to entry are few. In fact, many vendors expressed their concern about the lack of barriers to entry and the competitive pricing tactics that follow. Our first EMNS MarketScope focuses on enterprise-level offerings and has an overall rating of Positive.

Organizations are recognizing that they can use EMNS for more than emergency/mass notification purposes, with the following being the six main use cases:

1. Emergency/crisis events that require stakeholder notification (workforce, customers, partners and so forth)
2. Business operations notifications, such as workforce management roll call or mustering, call-outs to parents for absentee students, upcoming and special event announcements, important meeting reminders, and so forth
3. Business context-based alerting that gets triggered from another business process (for example, checking account overdraft, late payment, flight delays, work availability options by locale — "Send Work Now," grade delivery, incoming injured patient and so forth)
4. IT service alerting
5. Reverse and enhanced public emergency call numbers (for example, 911 and E-911 in the U.S.)
6. Public safety (for example, student tracking on a college campus)

Organizations will expand their use of these offerings, thereby driving down the cost, making these offerings feasible for the smallest of organizations.

No vendor has an offering that supports all use cases. The EMNS market addresses the first, second and third messaging use cases, with emergency/crisis event alerting being the primary reason for the use of these tools. Given the business evolution of a few of the EMNS vendors, some have their use cases and associated message volume just the reverse — business operations and context-based alerting are the primary uses of their tool. At present, there is some vendor overlap between the EMNS and communications-enabled business process (CEBP) markets (see "Hype Cycle for Enterprise Communication Applications, 2010") through an EMNS product API for integration to a triggering business application.

Gartner forecasts a growing relationship between the EMNS and CEBP markets within the next five years as alerting and notification of all kinds become routine. Many firms use multiple products to address all their alerting/notification needs. In some cases, these firms are looking to consolidate their vendor portfolio; however, when it comes to alerting that is triggered from a business application, it is not always possible to consolidate because of the complexity of the integration between the business application and the EMNS tool.

User Advice: Review Gartner's "MarketScope for Emergency and Mass Notification Services" to understand the EMNS market before you start your own implementation. Use Gartner's EMNS RFP template ("Toolkit: Emergency/Mass Notification RFP Template") to develop your own EMNS RFP.

Vendors focus on the following main markets: higher education (K-12 vendors specific to this market are not covered in the MarketScope), healthcare, government and private enterprise — regulated and not. Choosing a vendor that has experience in your market will result in a more-aligned offering to your business operations.

Customers prefer a subscription-based or hosted solution, which means that the software and hardware necessary to operate the EMNS system are located off-site and accessed via a Web portal, desktop API or handheld device. Decide which approach is best for your own operations and security/privacy needs.

EMNS pricing is competitive, but pricing models vary. Most are based on the number of contacts in the contact database, plus additional charges for message volumes per endpoint. E-mail messages tend to be unlimited; phone messages are usually restricted to a certain volume and price point; and proprietary SMS messages are priced like a cell phone call.

Review and examine the types of use cases for which you will be using EMNS. Knowing all of the usage can help in product selection, as well as understanding pricing quotes among the vendors.

Some EMNS vendors use resellers, such as telecommunications companies, to resell their products. For customers looking to expand their use cases beyond emergency notification, these resellers might be of interest, because some have an EMNS offering with unlimited messaging for all endpoints. However, the price may be beyond what some firms are willing to pay.

Government organizations should not, for the sake of redundancy, opt to use multiple EMNS vendor technologies, because, when activated in unison, they have the potential to overload servers. Some of these systems can be linked to a geographic information system map interface to develop a more-targeted approach, such as a certain postal code or a neighborhood within a certain radius of a chemical spill.

Carefully plan your enrollment procedure to ensure that all people needing to be contacted are included in the service and that their contact information is current and complete.

Carefully plan the types, number and content of notification messages because:

- Recipients of notification messages may ignore notices if too many are sent about the same event.
- Carrier-based character restrictions on text messaging make the formation of a meaningful message a challenge.
- During a regional disaster, don't overload the telecommunications infrastructure with needless messages.

No vendor can ensure or guarantee message delivery — all they can prove is that message volume levels are leaving their system. EMNS prospects and vendors need to set realistic expectations regarding message volume for each endpoint — phone call, e-mail, SMS and so forth — and in the aggregate.

If you want 24/7 availability of a service — and if the vendor has a business interruption such as a disaster, scheduled maintenance that runs over time, unscheduled maintenance and so forth, and a documented service-level agreement to go along with it — then you must validate your needs against the EMNS vendor's capability and delivery of that capability. At times, you might have to contract for it in addition to what the vendor provides in its base offering. EMNS users should apply a holistic analysis to SLAs and look for potentially unrecognized factors.

Business Impact: The interest in and need for EMNS tools continue to grow among governments, private enterprises (regulated or not), educational institutions and operators of critical infrastructures. The use of EMNS reduces overall costs by consolidating functions and improves the capability to deliver and update uniform message delivery to targeted and mass groups. The business benefits of using an EMNS tool include:

- Many key personnel can be notified in minutes.
- Management can focus on critical decision making and exception handling, instead of message delivery.

- Human error, misinformation, rumors, emotion and distraction, so often found during a crisis, are eliminated from automated EMNS communications.
- A documented notification audit log can be provided for real-time and post-event management.

Benefit Rating: Moderate

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Sample Vendors: Amcom Software; Amtelco; AtHoc; Benbria; Blackboard; Cooper Industries; Dell; Emergin; Enera; Everbridge; Federal Signal; FirstCall; Global AlertLink; MIR3; Omnilert; PlantCML; Rave Mobile Safety; ReadyAlert Services; Send Word Now; SpectraRep; SunGard Availability Services; Transformyx; Twenty First Century Communications; Varolii

Recommended Reading: "MarketScope for Emergency and Mass Notification Services"

"Toolkit: Emergency/Mass Notification RFP Template"

"Research Roundup: Business Continuity Management and IT Disaster Recovery Management, 3Q09"

"Toolkit: Requirements for Crisis Command and Emergency Operations Centers"

"New York Projects Show Critical Need for Unified Emergency Management"

"Q&A: How Universities Can Notify Students of a Crisis"

"Case Study: City of Chicago and ChicagoFIRST Public-Private Partnership"

Workforce Continuity

Analysis By: Roberta Witty

Definition: Workforce continuity identifies the technical and human resources needed to run the business in the recovery process. There are three goals:

- Enhance the organization's ability to recover from a disaster.
- Preserve the reputation of the organization.
- Maintain the social networks normally found at work.

A workforce continuity program delivers solutions for these critical recovery areas:

- Life and safety (the most important objective of any recovery plan)
- Availability to immediate resources, such as cash, food and shelter
- Workforce personal preparedness, so that recovery team members are comfortable leaving their families and homes to come to the aid of the organization
- Workforce event preparedness, so that they know how to respond to highly stressful events and don't make the situation worse than it is
- Workforce communications to ensure that rumors are controlled and the reputation of the organization is maintained throughout the event

- Recovery staffing to ensure that you have the correct teams, structure and members
- Work space — where workforce personnel will work from when their primary workplaces are no longer available
- Application access, which includes desktop availability, as well as other technical delivery mechanisms, to the applications and data that recovery teams need

Actions included in a workforce continuity plan fall into the following three categories: recovery operations, human resources and workforce communications.

Position and Adoption Speed Justification: Organizations that have any type of recovery plan have addressed their workforce continuity needs to some extent. For example, the IT disaster recovery plan must, by definition, address the IT recovery teams needed to recover the data center; work-at-home solutions are a key part of many continuity strategies.

Pandemic planning has also jump-started organizations addressing workforce continuity planning. However, unless the organization has made the transition from IT disaster recovery to business continuity management (BCM), it typically has not addressed all aspects of its workforce continuity needs. For example, workforce collaboration tools, smartphones and social-networking software are being implemented with greater frequency so that workforce members can respond to a crisis from any location, thereby enhancing crisis/incident management, as well as recovery needs.

Therefore, Gartner predicts that, as organizations make this transition, workforce continuity needs and solutions will mature. In prior BCM Hype Cycle reports, our Hype Cycle rating for workforce continuity was based on the maturity of these technical solutions. For 2010, our rating is based on the overall process of ensuring that workforce continuity preparedness needs are addressed. Therefore, it moved significantly to a post-trough position.

User Advice: Consider these areas when developing a workforce continuity plan:

- Does the workforce know how to evacuate the building, where to receive immediate medical care and trauma counseling, and whom to contact during an emergency?
- Is your workforce prepared at home so that it can leave loved ones to come to the aid of the organization? Have you implemented a personal preparedness training program for all workers?
- Is your workforce prepared with expert training in hazardous materials handling, emotional/psychological procedures, emergency medical care and so forth so that they know how to respond to highly stressful events and don't make the situation worse than it is?
- Do you understand their willingness to come to work during a disaster, especially a regional disaster, and have you built mitigation and incentive controls into the formation of your workforce recovery teams? For example, many workers will need to stay home to care for children and elders — if they are primary/mission-critical workers, you need to have a backup plan for them if they cannot come to work. Also, some workers are concerned about taking on tasks that they are not 100% trained for and are concerned about legal repercussions if they underperform or make mistakes when executing.
- Do you have access to immediate disaster resources, such as cash, housing arrangements and transportation, for at least two weeks?

- During a crisis, communication is key. Communicate often with your workforce, but don't overdo it, or the messages will be diluted and ignored. Recovery teams must be in constant and immediate touch with management and colleagues during a crisis. Do you provide collaboration tools, such as instant messaging, to the workforce? Does your command center have a way to share information across the virtual team and with others?

- Do you know what staff is needed for each business, as well as the decision-making process at each location and when they are needed? Some people will be more critical (for example, customer-facing staff) and others less critical (for example, programmers). Some workers will be able to work from home more easily. By taking a segmented strategy to the formation of your workforce recovery teams, you can identify high-risk, high-priority groups and make different preparations accordingly. Do you have secondary support for mission-critical business processes? Can you procure a content expert with a phone call?

- Are you building in shift rotation to alleviate fatigued workers?

- When planning for regional events, many enterprises look at work area recovery in a 150-mile radius from the production site. Hotels (mostly prewired or wireless for Internet access) are becoming a common approach for large groups to collocate during a crisis.

- Does your workforce have remote access to applications to do their jobs at an alternate facility? Do you know who can work from home, and who is required to be on-site at the recovery/alternate location?

- What is the organization's culture and management style? Is it centralized/decentralized or formal/informal? Each dictates a workforce recovery approach. How does information flow — bottom-up and fast, or top-down and slow? Information must flow to the workforce regularly; management styles may need to loosen to accommodate a workforce in crisis.

Business Impact: Having a workforce continuity plan is critical to ensuring that the organization can recover from a local or regional disaster. Without your workforce, there is no business.

Benefit Rating: High

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Sample Vendors: Apple; Computer Alternative Processing Sites; Facebook; HP; IBM BCRS; Preparis; Research In Motion; SunGard Availability Services; Twitter; XBRM

Recommended Reading: "Findings: The U.S. GSA Is a Telework Success Story"

"Workforce Continuity: Best Practices for Workforce Management"

"Workforce Continuity Defined"

"Personal Preparedness Enhances Corporate Recovery"

"London Bombings Confirm Need for People-Based Continuity Plans"

"Safeguarding the Workforce in Uncertain Times"

"Protecting People, Knowledge, Work: Are You Prepared?"

BCM Methodologies, Standards and Frameworks

Analysis By: Roberta Witty

Definition: The growing visibility of business continuity management (BCM) in boardrooms around the world is putting considerable attention on the development of a best-practice model for BCM methodologies, best practices, terminology and so forth.

There are many initiatives:

- Industry-based, such as the Federal Financial Institutions Examination Council (FFIEC); New York Stock Exchange (NYSE); National Association of Securities Dealers (NASD); North American Electric Reliability Corporation (NERC); Securities Industry and Financial Markets Association (SIFMA); the Health Insurance Portability and Accountability Act (HIPAA); Hong Kong Monetary Authority; Australian Prudential Regulatory Authority (APRA); Ontario Securities Commission (OSC); central banks in India, Indonesia, Russia and so forth; The Center for Financial Industry Information Systems (FISC — Japan); Monetary Authority of Singapore (MAS); Financial Services Commission (Korea); and Basel Joint Forum.
- Industry-neutral, such as the Information Technology Infrastructure Library (ITIL) and the Supply Chain Risk Leadership Council.
- Country-specific, such as the National Fire Protection Association (NFPA) from the U.S., the British Standards Institution (BSI) in the U.K., Standards Australia, Standards New Zealand, Associacao Brasileira de Normas Tecnicas (ABNT) and Korean Industrial Standards.
- Many BCM vendors, disaster recovery vendors and service providers that promote an existing model, or their own proprietary model, and influence what organizations use.
- Personal certification processes, such as DRI International (DRII) and Business Continuity Institute (BCI).
- Organizational certification — As of June 2010, three standards support organizational certification (as opposed to personal certification through DRII and BCI): NFPA 1600 (2007 and 2010 version), BS 25999-2:2007 and ASIS SPC.1-2009. These three standards were selected as part of the U.S. voluntary accreditation and certification program for emergency preparedness and business resilience called PS-Prep, "Implementing Recommendations of the 9/11 Commission Act of 2007" (aka Public Law 110-53, Title IX, Section 524). This initiative will not create a new standard; its directive from the U.S. Congress is to adopt existing work and to expand the list of acceptable models over time.

The following are examples of BCM regulations, standards and frameworks:

- ASIS International Business Continuity Guidelines (international)
- ASIS SPC.1-2009
- Basel Capital Accord
- BCI — The BCI Good Practice Guidelines (international)
- BITS Shared Assessments Program
- BSI — Business Continuity Management Specification (BS 25999 and BS 25777, U.K.)

- BSI Publicly Available Specification (PAS) 77:2006 IT service continuity management (U.K.)
- Canadian Standards Association — Z1600 Standard on Emergency Management and Business Continuity Programs
- DRII — Generally Accepted Practices for Business Continuity Practitioners (international)
- Expedited Funds Availability Act
- Federal Emergency Management Agency (FEMA) — National Incident Management System (NIMS)/Incident Command System (ICS)
- FFIEC Business Continuity Planning IT Examination Handbook: March 2008 (U.S.)
- Gramm-Leach-Bliley Financial Services Modernization Act
- Standards Australia/Standards New Zealand — HB 221:2004
- Standards Australia — HB 292-2006
- HB 292-2006 A Practitioner's Guide to Business Continuity Management (Australia)
- HIPAA/Joint Commission — healthcare in the U.S.
- ITIL v.3 (international) — "IT Service Continuity Management" is part of the "Service Design" book in ITIL v.3
- International Organization for Standardization (ISO) —27001:2005, 27002:2005, PAS 22399:2007, and 24762:2008
- Monetary Authority of Singapore — Business Continuity Management Guidelines
- NFPA 1600 — Standard on Disaster/Emergency Management and Business Continuity Programs (2007 and 2010 Editions, U.S.)
- NASD 3510
- National Institute of Standards and Technology (NIST) SP 800-34 (U.S.)
- NBR1 5999-1, 5999-2 and NBR ISO/IEC 24762 — direct translations of BS 25999-1, BS 25999-2 and ISO/IEC 24764
- NERC Critical Infrastructure Protection Cyber Security Requirement 9
- NYSE 446
- Applied Prudential Standard (APS) 232 — Business Continuity Management (Australia) for financial institutions
- Singapore Standard SS507:2004 business continuity/disaster recovery service providers
- Singapore Technical Reference TR19:2005 Business Continuity Management

Position and Adoption Speed Justification: For most enterprises, no single regulation, standard or framework exists that defines the BCM requirements they should meet, although enterprises may be subject to service-level agreements and contractual requirements that contain

availability requirements. The exceptions are financial services (driven by regulations to identify and protect against anticipated or reasonably foreseeable internal and external threats, or hazards that could result in unauthorized disclosure, misuse, alteration, or destruction of customer information and customer information systems; substantial harm; or inconvenience to any customer), healthcare, nuclear industry and electric utilities.

A proliferation of regulations, standards and frameworks has occurred over the past 10 years. Because of PS-Prep, the three standards selected for that program will likely be considered as representing a set of BCM best practices for nonregulated industries. The overall uptake of PS-Prep will take time, so Gartner does not predict that one will predominate worldwide within the next 12 months. There will be pockets of acceptance by organizations, depending on industry, geographic location and maturity of the existing BCM program. ISO is working on a business continuity management standard — ISO 22301 — and that will likely be the worldwide adopted standard when it is published. It is also working on an IT disaster recovery standard currently named ISO/IEC 27031 "Information technology — Security techniques — Specification for ICT Readiness for Business Continuity." Because no single model has been agreed on, there is no single set of audit standards that can be used for business continuity in the same way that specific auditing standards have been defined for regulations such as Sarbanes-Oxley 404, HIPAA and Payment Card Industry. ISACA Doc G32 provides guidance to IT auditors for assessing BCM plans.

As organizations recognize their need to work more closely together with multiple service providers or business partners, there is pressure to use common frameworks and standards to make it easier to integrate processes across organizational boundaries. As a result, there is increasing focus on compatibility across the supply chain, such as enforcing the use of ITIL as a common framework for IT service management in a multisourced environment. In the same way, the use of a common BCM model would facilitate BCM planning, testing and auditing of business activities across the organization's product/service delivery ecosystem. Enterprises are putting more pressure on their service providers to meet the enterprise's recovery requirements; this is evidenced in the growing number of vendor/supply chain availability risk management programs in a number of large enterprises. Therefore, organization certification will be the best avenue for the service provider to demonstrate to multiple customers that it can, in fact, meet recovery requirements.

The position on the 2010 Hype Cycle for business continuity management has been changed to slightly post-trough because of the PS-Prep program's decision on which standards to include.

User Advice:

- Organizations should review a number of existing models, and select or develop their own BCM model based on appropriate industry and country regulations and standards. Using multiple references will provide a broader view of BCM to assist in developing a program applicable to an organization's business needs.
- Organizations should find out which models their service providers, trading partners, customers and external auditors are using for their audit work. In the event of an actual interruption, most businesses today are heavily dependent on their interactions and interdependence with other stakeholders in their business ecosystems.
- Nonregulated U.S.-based organizations should follow the work being done in relation to U.S. PS-Prep to understand how it might influence their models and organizational certification initiatives.
- Due to the lack of a single BCM model and supporting framework, the only means by which organizations can assess the effectiveness of recovery and continuity controls is

through the use of live testing, or experiencing a real disaster and executing a recovery plan in real time. Even organizational certification does not provide a guarantee that the organization will effectively recover from an actual disaster — it only confirms the maturity of a BCM program process and management thereof.

- Know your organization's culture and management style in regard to identifying gaps, single points of failure, or mistakes in current business and IT processes. Organizations that shy away from, or refuse to identify, root causes of failure will respond less effectively than an organization that takes a proactive approach to risk identification and mitigation.

Business Impact: One benefit of a BCM model is that there will be more consistency and completeness (by using a third-party-vetted process) across the enterprise in the execution of the BCM program. A second benefit is that responding to BCM program validation requests from auditors, customers and trading partners will be easier due to that consistency. However, it is one thing to have a solid BCM model; it's another thing to be "process mature" and performing well in each of the categories listed in that model. Finally, having a BCM model is a goal that will benefit organizations in many ways, but it cannot guarantee that organizations will be able to recover from a disaster.

Because of the proliferation of regulations, standards and frameworks, it is easier than ever to develop a good-quality BCM model. In addition to assessing the quality of an enterprise's BCM model, determine:

- How mature the organization is in terms of embedding BCM-style thinking and acting into business operations management (process maturity)

- How effectively the organization performs in each of the domain areas — that is, having the right plan in place to ensure recovery during continuous exercising, practices and demonstrations

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Early mainstream

WAN Optimization Services

Analysis By: Joe Skorupa

Definition: WAN optimization services provide enterprises with improved application performance across the WAN through the use of protocol spoofing, route control, compression and caching in the network cloud, or as provisioned and managed WAN optimization controller (WOC) services. In time, more WOC functions will be built into the carrier infrastructure.

Position and Adoption Speed Justification: As WOCs mature and gain deeper penetration in the business market, service providers will respond with more-integrated services to retain and gain customers. To date, the global carriers have taken a cautious stance with these services, basically offering to manage WAN optimization devices. The notable exception is Akamai's network-based acceleration service. However, during the next 12 to 18 months, we expect a more aggressive global rollout of WAN optimization services, with the eventual integration of services into the carrier network itself. While protocol spoofing, route control, compression and application-specific accelerations can be delivered effectively via network-based services, last-mile bandwidth limitations will force on-premises solutions for data-intensive applications.

Cloud-based application providers may choose to offer network-based acceleration as a options for their HTTP/HTML applications. In many cases, this approach will boost performance without incurring the cost and complexity of an on-premises equipment-based approach.

User Advice: Organizations that have customer-facing Internet applications, or those that want to purchase WAN optimization as a service, should consider specialist service providers or global carriers for a variety of offerings. Managed services that include software-based WOC (SoftWOC) capabilities can be valuable for supporting mobile workers, as well as for supporting critical employees during disruptions, including pandemics and natural disasters.

Business Impact: WAN optimization services can reduce the cost of WAN bandwidth, while delivering significant gains in application performance.

Benefit Rating: Moderate

Market Penetration: Less than 1% of target audience

Maturity: Emerging

Sample Vendors: Akamai; AT&T; BT; Orange Business Services; Verizon Business; Virtela

Recommended Reading: "How Application Networking Services Are Evolving"

Climbing the Slope

Bare-Metal Restore

Analysis By: David Russell

Definition: Bare-metal restore (BMR) products provide a way to recover (or redeploy) the system, applications and data to a PC or server that is "bare metal" — i.e., it has no previously installed or corrupted software and/or operating system. These products deploy a sector-by-sector, disk-imaging approach to making a copy of the contents of a hard disk. Most let you boot a repaired computer from a CD-ROM or external USB drive and restore your disk drive from CD-ROM, DVD, a second disk drive connected to the computer, a disk on the network or, in some cases, tape. Products provide backup and restoration of servers, networked workstations or PCs at the level of discrete files or the entire disk volume. The ability to restore to hardware that is not the same as the original system is now a requirement. In addition, some vendors enable users to restore different operating systems to the same hardware (for example, a Linux system image to "Wintel" hardware).

Although solutions are available for many operating systems (including Linux, Unix and Windows), the Windows platform is the predominant use case for BMR solutions.

Position and Adoption Speed Justification: BMR products have long been used to redeploy or repair PCs, and most organizations have a product for that purpose. Previously, server BMR has been less widely used, because older Microsoft DOS-based solutions required the server to be taken offline, and required a reboot to do the backup. Today's Windows- and Linux-based products have removed that requirement, raising interest and increasing usability, because the process is less invasive. Many solutions enable incremental images to be taken after an initial base backup, decreasing the time required to back up the system. BMR is particularly needed in organizations with large deployments of Windows servers, where security or other patch deployments may require quick deployment and a fast recovery of a system or files if problems arise.

User Advice: Although backup vendors have continued to improve the system recovery features of their traditional backup solutions, and service solutions have added system recovery capabilities, stand-alone solutions should be considered if ease of use and advanced features are a consideration. Choosing solutions that offer dissimilar hardware restoration is important, because PC and server hardware configurations change frequently, and the ability to restore to different equipment can be valuable in a disaster recovery scenario. For smaller single servers or individual desktops and laptops, these tools can serve as backup products for file and application data, as well as provide protection for the operating system.

Business Impact: The need for rapid system recovery is more important than ever, because an entire business model can hinge on a company's servers functioning properly. In the event of a hardware failure, a traditional recovery can take many hours or several days, especially if the new systems are not identical to the ones that went down. BMR dramatically reduces recovery times for servers and can get PCs up and running rapidly.

Benefit Rating: High

Market Penetration: 20% to 50% of target audience

Maturity: Mature mainstream

Sample Vendors: Acronis; CA; Cristie; EMC; IBM; Novell; StorageCraft; Symantec; UltraBac Software; Unitrends

Recommended Reading: "Enterprise Backup/Recovery Market Update: Change Driven by Virtualization and Data Reduction"

"Poll Shows Disk-Based Backup on the Rise, With a Few Surprises"

Business Impact Analysis

Analysis By: Roberta Witty

Definition: A business impact analysis (BIA) is a process that is used to identify and evaluate the potential effects (financial, life/safety, regulatory, legal/contractual, reputation and so forth) of natural and man-made events on business operations. The BIA is considered to be the starting point of a business continuity management (BCM) program — after the program formation, of course. The BIA establishes:

- The period of time within which the workforce, technology (systems, applications and IT services), vital records and equipment must be recovered — the recovery time objective (RTO)
- The level of acceptable losses of data, or recovery point objective (RPO)
- Operational performance levels that must be achieved after a disruption
- The order/priority in which normal levels of operation need to be resumed
- Operational resource requirements over time (hours, days, weeks, months) after a disruption
- Critical interdependencies between each business process and other internal and external third parties, including suppliers
- The cost of downtime — an analysis of how the organization will be impacted in finances, reputation, legal/contractual issues, penalties and so forth, so that an

appropriate investment level and funding allocation for recovery solutions and technologies can be established

- Enhanced senior management awareness and understanding of the business risks associated with business disruptions
- One of the most complete views anywhere in the enterprise of business operations that can be used in business decision making for activities such as rightsizing and mergers and acquisitions

A BIA cannot be conducted without the business and the IT organization's input. Focusing on only one side is a waste of resources and can lead to a false sense of security in the organization's recovery readiness. It is not unusual to have the IT department more active in the BCM process early in the project. However, sometimes, the IT organization's view on what is important doesn't align with the business view. Therefore, both sides need to work together to develop an enterprise view of the BIA.

There are a variety of proprietary software products available to conduct business impact analyses that may be useful, but they are not essential. The key benefits of using a software tool include ease of distributing questionnaires, collating responses, analyzing results, storage of information and reporting of the results. Some vendors have advanced analytics that produce results, such as projected RTO and RPO targets, given a set of cost of downtime inputs — an important feature that makes the RTO and RPO determination a bit less subjective when doing it manually.

Position and Adoption Speed Justification: The BIA is the foundation on which recovery strategy and solutions are built, and it provides the basis for funding recovery solutions. Although it is generally accepted that comprehensive, detailed business impact analyses are essential for IT disaster recovery management (IT DRM) and BCM programs, they are often not repeated on an annual basis — the recommended cycle — due to the amount of work involved, the complexity of the process and the shortage of skilled practitioners within the enterprise. It is not uncommon for organizations to engage BCM consultants to develop the BIA process, execute it and refresh it on a regular basis. BIAs have been applied superficially in many organizations, usually as a component of IT DRM planning, as risk assessments, or during project and/or change management. IT DRM teams have been inclined to use simple ad hoc methods to determine application and system criticality, rather than to determine business impact.

Due to the growing number of disasters experienced worldwide, and the pressure from the supply chain and auditors, many firms are revisiting this part of the BCM program life cycle and refreshing the BIA — which might have last been updated five years ago and, therefore, has little resemblance to current business operations. Therefore, we have changed the 2010 position for the BIA profile to acknowledge this maturing of the overall BCM program.

User Advice: As BCM processes mature, more-precise BIA practices will be required and adopted. Therefore:

- Select and implement a formal BIA methodology that meets BCM, IT DRM and risk management requirements so that a standard approach is used to gather and assess information, and establish a common repository for all risk-related information.
- Select a method that is aligned with a generally accepted standard or framework, such as those published by BCM frameworks and standards bodies (see BCM Methodologies, Standards and Frameworks).

- Consider implementing a BIA software product because it will simplify the process. However, be aware that it will not eliminate the need for an experienced BIA practitioner, nor will it eliminate interviews with or involvement of individuals knowledgeable in the activity being analyzed. When an organization selects a BCM tool, stringent BIA capabilities must be specified as key selection criteria.

Business Impact: Globally, there is a growing awareness of availability risk in organizations in all market sectors. One effort called the Supply Chain Risk Leadership Council is working on an ISO standard for supply chain risk management that cuts across all market sectors. The demand for an accurate assessment of business impacts due to man-made and natural threats will rise as management boards are required to formally report on risk management measures. A standard BIA approach will enable organizations to gather information from across the organization for making critical risk-related decisions and to formulate rational risk responses. The BIA can be used as a powerful catalyst for a management team to start investing time and money in BCM by bringing attention to potential losses and mismatched recovery and continuity capabilities.

Benefit Rating: High

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Sample Vendors: Archer Technologies; Avalution; Bold Planning Solutions; Business Protection Systems International; COOP Systems; CPO; eBRP Solutions; EverGreen Data Continuity; Factonomy; Linus; Paradigm Solutions International; RecoveryPlanner.com; Strategic BCP; SunGard; Tamp Systems; Virtual Corporation

Recommended Reading: "MarketScope for Business Continuity Management Planning Software"

"How to Understand and Select Business Continuity Management Software"

"Best Practices for Conducting a Business Impact Analysis"

Distributed Virtual Tape

Analysis By: David Russell

Definition: Virtual tape for distributed system technologies uses software and disks to emulate physical tape automated libraries, tape drives and tape volumes. This enables traditional tape backup/recovery software products to achieve the benefits of disk-based backup and recovery speeds, with little or no change in the backup software or process. In addition to the potential benefits of increased automation and backup and restore performance improvements, compared with physical tape, some customers use virtual tape libraries (VTLs) to constrain the costs associated with having to add additional tape drives or tape capacity by placing a VTL in front of their physical tape infrastructure, thereby offsetting additional investments in physical tape and/or preserving legacy investments.

In addition, ease of management and potential performance benefits over conventional disk — coupled with advanced features such as compression, deduplication, encryption, data shredding and replication — are value-added services that further extend the value proposition of VTLs beyond simply emulating physical tape. (Virtual tape technology is also deployed on mainframe platforms, providing similar benefits; however, this section is focused on the distributed platform implementations.)

Position and Adoption Speed Justification: Appliances built using virtual tape software have gained traction in the market since 2004 as fast and easy methods for incorporating disks into recovery environments. All major storage vendors offer virtual tape products, and many backup/recovery applications enable users to create virtual tape areas on disks using the backup software. Since late 2009, there has been a rising preference toward advanced disk solutions that do not present a VTL interface. Instead, these non-VTL disk-based solutions use file interfaces, such as CIFS, NFS or Symantec's OST, which do not require the additional management of setting up VTLs, tape drives and tape volumes.

User Advice: Users should consider VTLs because they improve the speed of backup and recovery operations, compared with traditional tape; offer an integrated hardware/software solution; and can be easily integrated into the backup infrastructure, with little or no modification required. However, other bottlenecks may become more apparent after the VTL has been installed, such as the speed of older 1GB/second host bus adapters on the protected server not being able to transfer input/output as fast as the VTL system receives it.

Users may want to evaluate VTLs for their ability to scale to higher levels of capacity and performance than is easily available with conventional disks, as well as for advanced features, such as their ability to replicate to another site and their data deduplication capabilities. In addition, redundant array of independent disks (RAID) storage offers another level of protection that tape doesn't provide. Users also have the potential to reduce the number of physical tape drives, cartridges or libraries required in the future.

VTL solutions typically cost more than physical tape or even conventional disk approaches; however, their ease of implementation, manageability, increased performance and advanced features may offset this. With the advent of disk cost-saving technologies (such as the Serial Advanced Technology Attachment [SATA] disk, compression and deduplication), the economics of a VTL with these characteristics, along with a likely lower overall maintenance bill, could represent an overall reduction in the total cost of ownership (TCO), compared with some physical tape environments over a multiyear period.

For many organizations, VTLs will represent a staging area or cache for more-recent data, with older data spilling to physical tape. However, small data sizes and deduplication can affect if and when the data spills to tape.

Going forward, for new investment decisions, Gartner advises that a VTL interface may not be the most appropriate choice for a user seeking disk-based solutions, as the advent of 10Gb Ethernet and the maturity of deduplication appliances may offer just as fast, or even faster, performance that does not require the overhead and management of tape operations. However, for users who are using physical tape export as the final stage for their data, a VTL's tape simulation feature is still important, making a VTL potentially more appropriate than the current non-VTL disk appliances.

Business Impact: Distributed VTLs improve backup and recovery times, increase backup success rates, provide operational efficiencies with minimal impact on the backup/recovery infrastructure, and increasingly incorporate advanced services (such as encryption, replication, compression, data shredding and data deduplication) into the VTL platform.

Benefit Rating: Moderate

Market Penetration: More than 50% of target audience

Maturity: Mature mainstream

Sample Vendors: Crossroads Systems; Cybernetics; Data Domain; EMC; FalconStor Software; Fujitsu; Hitachi Data Systems; HP; IBM Storage; Oracle; Overland Storage; Quantum; Sepaton; Spectra Logic

Recommended Reading: "Enterprise Backup/Recovery Market Update: Change Driven by Virtualization and Data Reduction"

"Poll Shows Disk-Based Backup on the Rise, With a Few Surprises"

"Competitive Landscape: Enterprise Distributed Backup/Recovery Software Growth Driven by Virtualization and Data Reduction"

"New Storage Solutions Can Modernize Data Life Cycle Management"

"The Half-Life of a Virtual Tape Library"

Pandemic Preparedness Planning

Analysis By: Roberta Witty; Steve Bittinger

Definition: The world has experienced pandemics in the past. It just recently successfully managed the 2005 SARS and avian flu, as well as the 2009 H1N1 virus (swine flu) — which was officially declared complete on Tuesday, 10 August 2010. It will certainly suffer pandemics in the future. The rising magnitude of global travel means that pandemics are more likely than ever to spread rapidly. Pandemics differ from normal disasters. They are at least regional and most likely global in scope, are of indefinite duration, and will drive staff absenteeism that may exceed 40% for extended periods. The accuracy of pandemic predictions is less important than the guidance that a particular prediction can have in improving decision making. Organization-level pandemic preparedness planning was greatly enhanced and matured through the 2009 H1N1 virus; SARS and the avian flu threat in 2005 were wake-up calls and demonstrated the potential for harm that pandemics can do.

Whatever direction events take — and predictions are that the next pandemic will be new to humans or a more deadly adaptation of a common virus (source: Joseph Fair, director of global field operations for the Global Viral Forecasting Initiative), enterprises can plan their pandemic strategies by developing their responses using scenario planning. A pandemic won't immediately affect IT systems, but during the long term, outages will be encountered, because providers and vendors will be affected too. As real-time processing becomes the norm in business operations, Gartner predicts that IT system outages will start around the second week into a regional pandemic that affects large numbers of people and organizations.

Position and Adoption Speed Justification: A Gartner 2009 cross-industry business continuity management (BCM) survey revealed that 49% of survey participants started their pandemic preparedness planning, and 23% planned to start in the next 12 months. That is a 32% and 44% increase over 2009's reported numbers of 37% and 16%, respectively. However, these survey results do not indicate the validity of such preparedness planning. Full-scale simulations in the U.S. financial services industry in 2007 revealed that existing BCM plans for pandemics are insufficient and require immediate upgrades. Existing BCM plans, normally based on limitation of access for short periods to individual locations or systems, are unlikely to provide adequate protection. Driven by strict regulations and examinations, financial institutions in the U.S. are considered by many to be the "model" industry for BCM; however, less than 12% of those participating in this nationwide pandemic simulation found their BCM plans to be "very effective" at maintaining normal operations. It was unlikely that adequate testing, planning and remediation of existing plans could take less than 18 months. This test was not redone in 2008 nor 2009.

Overall, there is not a lot more development to be done by organizations in terms of learning how to plan and prepare for a pandemic; the work is in putting into practice the advice already provided by Gartner, the World Health Organization, government entities and many other organizations worldwide. H1N1 response gave us — for the first time — a real-time update of the spread of the virus around the world. Much automation was implemented to track the virus's spread. Planning areas that continue to need more focus include:

- Joint planning with external stakeholders within the broader business ecosystem and potentially building relationships with "complementary" organizations. As an example, in one pandemic preparedness planning discussion with a grocery firm, it had made arrangements with a local bus company so that, if the grocery firm's truck drivers were sick, they might be able to call on the services of other bus drivers as backup.
- Medical care capacity. We were lucky that the 2009 H1N1 was not a major disaster; the next pandemic may not be as kind to the world.
- Vaccine development. Because each year brings new viruses, infections and diseases, the challenge is for the pharmaceutical companies to develop an appropriate and timely vaccine for that year's particular infection.
- Supply chain management. Organizations need to plan with their partners how each will respond if impacted by a large workforce outage. One organization was able to develop an "impact profile" of their suppliers in Mexico and the impact on revenue by product and customer using risk management automation that they had put in place as part of their supply chain risk management program.

The apathy that we generally saw before the emergence of H1N1 suggested that pandemic preparedness planning was in the Trough of Disillusionment (2008). However, H1N1 forced organizations to take the threat seriously, and in 2009, we positioned it at post-trough 10%. A year later, we place pandemic preparedness planning at the post-trough 30% position, indicating that many governments — Australia, the U.S., the U.K. and New Zealand in particular — and organizations made progress in preparedness planning. However, because of the challenge of vaccine development, we will never reach the Plateau of Productivity.

User Advice: The experience of the 2009 H1N1 pandemic should reaffirm the commitment to continually refine pandemic preparedness planning. IT can only address a few of the changes businesses need to make to be better prepared. Most changes will come in changing the business processes. Therefore, the following advice is provided to jump-start your pandemic preparedness planning initiative:

- Use scenario planning (low, medium or high impact) that includes both the spread and severity of a pandemic to build your plans now. Don't wait for an outbreak, because it will be too late. Service providers will not be able to address your organization's needs during times of crisis.
- Do not build plans addressing the response to a pandemic on competitive instincts. Rather, ensure that the plans are collaborative and supportive in nature.
- Start honest, rigorous, inventive, ongoing and documented testing immediately to isolate and remediate problem areas.
- Ensure that IT and business managers worldwide identify critical operation skill shortages and initiate staff cross-training, testing and certification. This requires the longest lead time and is the most disruptive of the improvements.

- Determine realistic business operation sustainability and the likely downtime for IT staff. Plan for absentee rates of 40% or more for 90 days, with various combinations of leaders and skilled staff. Do not expect heroic devotion — some workers will be home taking care of children or elders, because schools and long-term care facilities have closed; others may be "able" to go to work, but will be unwilling to do so.
- Work with your legal department and outside counsel to understand the legal ramifications of policy and business decisions you are making in regard to pandemic preparedness planning.
- Educate your workforce on personal hygiene issues.
- Work with public health and other government agencies, especially around loss of government services, vaccines and antiviral medication.
- Plan for the loss of travel and transportation, and reduced food supplies.
- Implement more work-at-home options (for example, virtual private network access and videoconferencing).
- Look to external parties in advance to provide skills where appropriate.
- Stagger the workforce work schedule into shifts.
- Work with customers and partners to minimize any disruption by developing coordinated crisis response capabilities. Implement supply chain risk management automation to enhance your planning and response capabilities in this area.
- Prepare for privacy protection reduction (for example, tracking sick people and travelers).
- Review where your business might have spare capacity as a result of the pandemic impact, and see where else it can be used.
- Review business continuity and disaster recovery plans to determine where and when they can be used.
- Stockpile critical supplies and raw materials — just-in-time inventories won't work.
- Use more online systems (for example, order taking, training, FAQs and knowledge bases).
- Implement WAN optimization controller software to reduce the network bandwidth of applications — reducing the potential for bottlenecks at the "last mile" of the Internet connection.
- Increase the number of personnel using PDAs — they are a boon for the distributed workforce.
- Applications that are graphics-rich will have to be modified in order to reduce the potential for bottlenecks at the last mile of the Internet connection.
- Review the use of social media for workforce communication during a pandemic.

Business Impact: Preparing to operate for an extended period with reduced, distributed staff can only benefit organizations. Although businesses cannot economically plan for all contingencies, "soft disasters" (ones that affect staff access to facilities without necessarily damaging staff or

facilities) are likely to rise in frequency in the near term. Depending on the nature of the pandemic, governments may force quarantine provisions well ahead of any obvious local outbreak, disrupting operations as certainly as a major plague.

Benefit Rating: High

Market Penetration: More than 50% of target audience

Maturity: Mature mainstream

Recommended Reading: "'Swine Flu': Lessons in Pandemic Preparedness From the Financial Industry"

"The 'Swine Flu' Is a Reason to Plan, Not Panic"

"'Swine Flu' Means Test Transaction Document Recovery Plans Now"

"New U.S. Guidance on IT in Pandemics"

"Prepared for Avian Influenza: Our Interview With T. Rajah, CIO, CLSA"

"Avian Flu Demands a New Kind of Business Continuity Planning"

Work Area Recovery

Analysis By: Roberta Witty; John Morency

Definition: Work area recovery ensures that an adequate employee work environment — office space, phone, telecommunications, applications, special equipment, vital records and so forth — is available if the primary environment becomes unavailable. There are a number of options that can be used for work area recovery, including external facilities, such as hotels, third-party IT disaster recovery (DR) service provider sites, conference center facilities, mobile units, rentable satellite/temporary business office locations and drop-ship programs, to internal options such as work at home, cafeterias, training facilities, alternate internal office facilities and so forth. Even though many service alternatives exist, there are still many organizations for which no formalized work area recovery strategy exists or that have adopted work-at-home approaches. The latter point has been revalidated through Gartner client inquiry briefings as well as survey findings.

Position and Adoption Speed Justification: Some work area recovery needs are automatically addressed through the recovery plan. For example, the IT disaster recovery management (DRM) plan must, by definition, address IT recovery team member work area needs to recover the data center. However, unless the organization has made the transition from IT DRM to business continuity management (BCM), IT typically has not addressed all the business-unit aspects of its work area recovery needs. Multiple recovery strategies may be required — one for office-based staff, one for specialized workgroups such as customer service teams and call centers, one for work-at-home professionals, and one for mobile sales and service employees. Implementing each of these strategies may well require different combinations of in-house support effort, the use of fixed and/or mobile work area recovery locations and sufficient network capacity to support a large number of networked endpoints.

In addition, the emergence of cloud computing and desktop virtualization makes businesses and consumers used to the idea of being remote from their corporate IT resources while conducting business or carrying out their personal affairs. Increasing trends toward outsourcing, multisourcing and working within complex business ecosystems has encouraged organizations to become more decentralized, reducing their need for work area recovery.

For 2010, we did not change the position of work area recovery because no significant change has occurred in the available service provider options. Additionally, the increased use of desktop virtualization will improve the economies of teleworker access management, making the work-at-home option a far more affordable service management alternative.

User Advice: The rule of thumb for identifying work area recovery facilities for a regional disaster is a 150-mile radius from the production facility.

If the business process is mission-critical, then dedicated recovery space may be required. For instance, during a regional disaster, when a third-party recovery facility may be filled with a subscription service. Your mission-critical personnel will not have a place to work.

Special equipment needs must be identified before the selection of a facility is made. For example, call centers and trading desks are unique technical environments that not every service provider can supply.

Not all work area recovery locations are colocated with the data center disaster recovery site; therefore, network connectivity to the data center recovery site is required.

Consider the workforce operation models of your organization — mobile, work-at-home, on-site required and so forth — when considering all work area recovery options. Expect to have multiple options supported. Look to leverage internal locations — they are usually less costly than third-party services.

Increasing numbers of organizations have found that it makes sense to equip their workforce with laptop computers. Laptops cost slightly more than their desktop counterparts, but the difference in cost often is easily justified by reduced power requirements, the additional work that people do from home, and flexibility from the BCM and work area recovery perspective.

During the BCM planning stages, it can be useful to take the ongoing rate of technology and business change into account, recognizing that the work area recovery approach will be different today from what it will be in three years, and that it will likely change again by the five-year mark. Understanding how these changes are progressing (that is, the rising number of workers with laptops and work-at-home capabilities) provides a basis for thinking about changing work area recovery needs. For example, a contract with a third-party provider in place today can be reduced or eliminated within three years as you move to be in a position to take up other, and sometimes less costly, approaches.

Business Impact: Having a work area recovery plan is critical to ensuring that the organization can recover from a disaster — local or regional. Without your workforce, there is no business.

Benefit Rating: High

Market Penetration: More than 50% of target audience

Maturity: Mature mainstream

Sample Vendors: Agility Recovery Solutions; Caps Business Recovery Services; HP; IBM; Rentsys Recovery Services; SunGard Availability Services

Outage Management Systems

Analysis By: Zarko Sumic; Randy Rhodes

Definition: An outage management system (OMS) is a utility network management software application that models network topology for safe, efficient field operations related to outage restoration. OMSs tightly integrate with call centers to provide timely, accurate, customer-specific

outage information, as well as supervisory control and data acquisition (SCADA) systems for real-time-confirmed switching and breaker operations. These systems track, group, and display outages to safely and efficiently manage service-restoration activities. OMSs are commonly used for historical outage reporting and the calculation of reliability indexes, such as the System Average Interruption Duration Index (SAIDI) and the System Average Interruption Frequency Index (SAIFI). OMSs require an accurate network connectivity model that reflects distribution network topology from the substation to each customer connection. OMSs contain outage determination procedures, mostly on the network-tracing schemas that enable customer outage calls to automatically roll up to common protective devices (that is, fuses or breakers), enabling operators to group calls and identify all affected customers.

Position and Adoption Speed Justification: Extreme weather events — combined with aging infrastructure that results from protracted low-investment levels in the utility delivery infrastructure — are straining utility companies' efforts to keep mandated customer service levels, which maintains industry focus on OMS as a solution that can ensure customer service quality, while improving labor efficiency and infrastructure use. Consequently, regulators are likely to decide favorably on cost recovery for investments that can mitigate these issues.

Geographic information systems (GISs) are often used to manage a common-customer connectivity model that enables the OMSs to be synchronized with planning and estimating models for other utility applications. Some GIS vendors are providing OMS functionality as an extension of their GIS solutions (for example, GE Energy and Intergraph). Other OMS vendors are tightly integrating OMS with SCADA and distribution management systems (DMSs), and replacing customer-information-system-based trouble call systems. The trend is toward integration with enterprise asset management (EAM) systems to better manage work and associated financials. EAM uses outage and other operational data from EAM business intelligence to optimize asset performance through proactive maintenance and replacement programs. Advanced metering infrastructures (AMIs) are providing outage management functionality directly from customer meters, and having outages associated with a network model improves field operations and safety. Enterprise service bus architecture with industry-standard model-driven integration, such as the Common Information Model, is enabling effective integration with associated business processes.

OMSs are predicted to become obsolete before maturity, because the new breed of DMSs will eventually incorporate the OMS functionality as we know it. DMSs will include OMSs within real-time advanced distribution SCADA, which also will include automated restoration and self-healing "smart grid" functionality.

User Advice: The utility industry is getting increasingly focused on the development of the future energy delivery infrastructure — the intelligent grid (also known as the smart grid). Among other benefits, the smart grid should improve network resilience via the use of advanced control functions, such as self-healing and event avoidance. This focus has fueled buyer interest in solutions that can integrate smart grid technology into the OMS environment. As a result, OMSs are evolving beyond an emergency response decision support system into a component of an integrated smart grid platform. Integrating OMS products with AMI for outage notification and the callback function — as well as including the automated switching operation intended to minimize an affected area through the use of advanced protection systems and local automation — is pushing OMS toward an advanced distribution management system (ADMS). In addition, access to up-to-date consumption data, via the meter data management (MDM) component of AMI, provides a data repository for network analysis, which is required to study the planned switching actions' impact on feeder overload.

Business Impact: Customer service, field operations and asset management are the areas of biggest business impact.

Benefit Rating: Moderate

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Sample Vendors: ABB; CGI; GE Energy; Intergraph; Milsoft Utility Solutions; Oracle; Telvent Miner & Miner; Trimble

Recommended Reading: "Magic Quadrant for Outage Management Systems"

"Utilities Must Bolster Outage Management Systems for Catastrophes"

Print/Mail Business Continuity and Disaster Recovery

Analysis By: Pete Basiliere

Definition: Print and mail recovery involves the ability of any enterprise, large or small, to continue producing and shipping the physical documents needed for generating sales and collecting payments when the primary operation responsible for the work is incapacitated. Print and mail recovery can be provided internally or externally, with the choice independent of where the normal production takes place. Companies that have an in-house production operation may choose to have an outsourced service provider perform print/mail recovery or, if they have more than one production site, enable the sites to back up each other. Companies that have outsourced their print/mail production are likely to have the original outsourcer secure its own backup resource, taking advantage of the provider's multiple sites and staffing, or retain a second outsourcer for recovery needs.

Position and Adoption Speed Justification: A comprehensive continuity plan balances IT system availability with an understanding of the costs (financial, regulatory and reputation) associated with business outages to prioritize continuity procedures. Too often, however, the plans fail to consider or to completely cover the company's critical printed customer communications.

Nevertheless, many enterprises at least cover the few key transactional applications (invoices, statements and checks), but not every form of transaction (policies, notices) or customer communication (direct mail, marketing collateral). Even if they have a written recovery/continuity plan, most enterprises don't conduct regular, realistic tests involving live data, or they have ill-formed plans to use a local mailing organization with questionable ability to handle large volumes of sensitive correspondence on short notice or for long periods of time.

Well-qualified print and mail business continuity providers have the necessary capabilities and capacity to provide all-inclusive coverage. Several of these providers, such as the ones listed below, have been offering these services for many years. So, while most enterprises do not have a comprehensive print and mail recovery plan in place, the resources exist that will provide complete recovery if the enterprise chose to do so.

User Advice: Infrastructure and operations managers who plan for business interruptions must conduct regular, comprehensive reviews and tests of the plans to cover the company's critical printed customer communications. As part of the review, managers must determine the impact on the business if mailings are missed, including direct costs (lost revenue and fines for failing to meet regulatory requirements) and indirect costs (damage to relationships with customers, suppliers and business partners). Also, determine whether print and mail continuity services will be provided internally or externally, considering internal capacity, potential outsourcers' locations and services, and costs. If outsourcing the recovery, then evaluate potential print/mail recovery providers based on their physical locations, networking, computing and telecommunications

systems, printing and mailing equipment, and quality assurance procedures. Once the review is complete, test the data processing, printing and mailing of the selected applications to ensure they are accurately produced within your service-level requirements.

Business Impact: The possible consequences of a missing, incomplete or outdated continuity plan, and therefore the inability of the print and mail operation to function at 100% during the event, are potentially disastrous for every organization. Possible critical issues include:

- Inability to produce mission-critical financial transaction documents, such as invoices, statements and checks
- Significant interruptions in the supply chain (inbound paper, envelopes and supplies, as well as outbound mail and shipments) and inability to make postage payments
- Failure to produce required documents within prescribed time frames to meet regulatory requirements

Benefit Rating: High

Market Penetration: 20% to 50% of target audience

Maturity: Mature mainstream

Sample Vendors: Bowe Bell+Howell; Cosentry; FTC Continuity Services; Kubra; Mail-Gard; MBA Group; Pitney Bowes

Recommended Reading: "Supplier Communications Critical During a Pandemic"

"Predicts 2010: The Role of Business Continuity Management Continues to Expand and Extend"

"Incorporate Print and Mail Into Your Disaster Recovery and Business Continuity Planning"

Entering the Plateau

WAN Optimization Controllers

Analysis By: Joe Skorupa; Mark Fabbi

Definition: WAN optimization controllers (WOCs) optimize traffic across WANs through the use of bandwidth reduction algorithms (such as compression, data reduction and caching), network-level optimization (such as transmission control protocol manipulation, advanced traffic and bandwidth management, and quality of service), and other application-layer protocol spoofing and manipulation techniques. WOCs effectively reduce the bandwidth on the WAN, as well as mitigate latency issues associated with many current and legacy protocols. WOCs will continue to evolve and integrate more communication functionality required at the branch. Data reduction algorithms will include a richer set of caching and enterprise content delivery network (ECDN) functionality. WOC devices are also being virtualized with the ability to integrate more fine-grained monitoring and service-level agreement (SLA) management as well as other remote requirements, such as video streaming, Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP) and security functions. The WOC environment also includes soft WOCs, which can provide optimization services to individual mobile clients. This can be particularly useful during periods of disruption, including pandemics and natural disasters.

WOCs can be expensive for smaller locations, and complex networks that employ techniques, such as asymmetrical routing and mesh forwarding, can be complex to implement. When WOCs enable server centralization, loss of the WAN link results in loss of access to applications and other centralized services.

Position and Adoption Speed Justification: New algorithms and technologies have significantly improved the performance gains available with these technologies. Integration into complex networks and the challenges of full network deployment have improved significantly during the past 12 months. There will be increasing competition in this market as some of the basic functions become embedded in other WAN equipment and services.

User Advice: The technologies for WOCs continue to improve. Organizations wishing to consolidate branch office servers or dealing with real-time applications should consider WOCs as a way to deal with branch optimization requirements. Organizations may also consider Soft WOCs as part of their business continuity/disaster recovery (BC/DR) strategy.

Business Impact: WOCs can reduce the cost of WAN bandwidth, while delivering significant gains in application performance. They also have the ability to consolidate infrastructures and improve compliance through the centralization of data.

Benefit Rating: Moderate

Market Penetration: 20% to 50% of target audience

Maturity: Mature mainstream

Sample Vendors: Blue Coat Systems; Cisco; Citrix; Juniper Networks; Riverbed Technology

Recommended Reading: "Magic Quadrant for WAN Optimization Controllers, 2009"

E-Mail Continuity

Analysis By: Matthew Cain; Donna Scott; David Russell

Definition: Many organizations are coming to the realization that e-mail is their most important communication and collaboration tool, one that merits substantial redundancy in case of a catastrophic failure. "E-mail continuity" refers to off-site e-mail services, which, in the case of a failure at the primary location, enable e-mail services to continue operating at an alternative site with relatively fast recovery times — generally between one and four hours — and with minimal data loss. This implies that the organization has redundant infrastructure in the alternative location, to be used for failover purposes, coupled with continuous data protection and/or data or message replication.

Position and Adoption Speed Justification: E-mail continuity software products, infrastructure and services have been around for many years and continue to mature in terms of ease of configuration, failover initiation and failback. External service providers have also prospered and proved to be reliable. Native software replication functionality that protects against any data loss has been present in Lotus Domino for years and is also included in Microsoft Exchange 2007 and 2010. Overall, we estimate that 25% to 35% of organizations have adopted these technologies, with adoption higher among large organizations and lower among small and midsize ones. Due to the maturity of the technologies and the growth of the market, e-mail continuity is very close to the Plateau of Productivity.

User Advice: Organizations need to determine — through conversations with business users and other e-mail stakeholders — how much e-mail data they can afford to lose (a recovery point objective [RPO]), and how long they can afford to be offline (a recovery time objective [RTO]), if a disaster occurs. Many organizations confuse e-mail availability — the percentage of time an e-mail system is functioning properly at the primary site — with disaster recovery or continuity requirements. High-availability architectures (such as local clustering) will typically not provide e-mail services in the event of a major data center failure or disaster. The RPO and RTO will

determine what type of investment in e-mail continuity software, infrastructure and/or services is needed.

Business Impact: For many organizations, e-mail is a vital channel of communication with employees and customers — one that must be operational at all times. There are a range of options, including storage area network (SAN)-based replication, appliances, software-based and external service provider/cloud-based services — all with pros and cons — that need to be examined when making an e-mail continuity investment.

Benefit Rating: High

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Sample Vendors: CA; Cemaphore; Dell; Double-Take Software; EMC; IBM (Lotus Notes/Domino); Microsoft; Mimecast; NetApp; Neverfail Group; Teneros

Distributed Tape Backup

Analysis By: David Russell

Definition: Distributed tape backup solutions create copies of data on tape for the purposes of logical recovery (for example, from a virus attack), physical recovery (such as from a disk failure) and disaster recovery (e.g., from site loss). Typically, this function is delivered via an enterprise backup and recovery product that provides a robust, scalable catalog for tape media tracking and recovery purposes. Support matrices for operating systems, applications and tape devices can vary among the products, as can scalability, ease of use and licensing methodology.

Position and Adoption Speed Justification: The distributed tape backup market, while important, is mature, and the mainframe tape backup market is even more mature. During the past several years, this marketplace has begun shifting to disk-based backup, prompting speculation concerning the demise of tape-based backups.

Gartner believes that increasing tape speeds and capacity (such as the recently released LTO-5 technology), as well as the ability to use tape as a tertiary copy of data, remain business-relevant areas for employing tape in the backup process. Directionally, most companies are pursuing backup to disk (or at least to disk first, before tape) and are looking to restore recent and/or critical data from disk. There is still a solid value proposition for incorporating tape into many recovery environments and, sometimes, using only tape for a portion of the backup requirements. Gartner polling shows that, as of December 2009, 25% of large enterprises still back up directly to physical tape, and another 60% use a disk-to-disk-to-tape (D2D2T) methodology, meaning that tape is still used in 85% of the enterprise backup process.

Distributed tape backup products are sometimes augmented by more-sophisticated media management and vaulting solutions, which provide more-granular tape volume tracking capabilities.

User Advice: Tape backup should be used primarily for longer-term recoveries (for example, the recovery of data that's older than a few weeks or months), for off-site storage of disaster recovery copies and for additional backup copies of data that is also backed up to disk. Disk-based solutions are increasingly being used for short-term recoveries, such as the recovery of data created or modified during a small number of days, weeks or months, and, also for disaster recovery. Tape alone may still be appropriate for backup in some organizations due to cost and legacy investments; however, disk economics and data reduction improvements, such as data deduplication, are positively affecting the increased adoption of disk solutions, which are often

used in combination with tape. In the future, many vendors may offer the ability to store dramatically compressed or deduplicated data on physical tape media (ranging from 7:1 to 25:1 deduplication ratios), which could make tape-based backup more compelling as an additional or extended option to support disk-based solutions. Using the current capacity and cost of LTO-5 cartridges, this could lead to a price of $10 to under $3 per terabyte. To date, CommVault's backup application offers deduplication to tape, and other vendors may follow suit.

Business Impact: Business continuity and disaster recovery plans are critical components of IT and business strategies. Distributed tape backup represents a mature and economical means of providing backup, long-term archiving, disaster recovery and business continuity for vital corporate data.

Benefit Rating: Moderate

Market Penetration: More than 50% of target audience

Maturity: Mature mainstream

Sample Vendors: Arkeia; Atempo; BakBone Software; Barracuda Networks; CA; CommVault; EMC; HP; IBM; Innovation Data Processing; Microsoft; Symantec; Syncsort; Zmanda

Recommended Reading: "Enterprise Backup/Recovery Market Update: Change Driven by Virtualization and Data Reduction"

"Poll Shows Disk-Based Backup on the Rise, With a Few Surprises"

"Backup and Recovery Optimization and Cost Avoidance"

Appendixes

Figure 3. Hype Cycle for Business Continuity Management, 2009

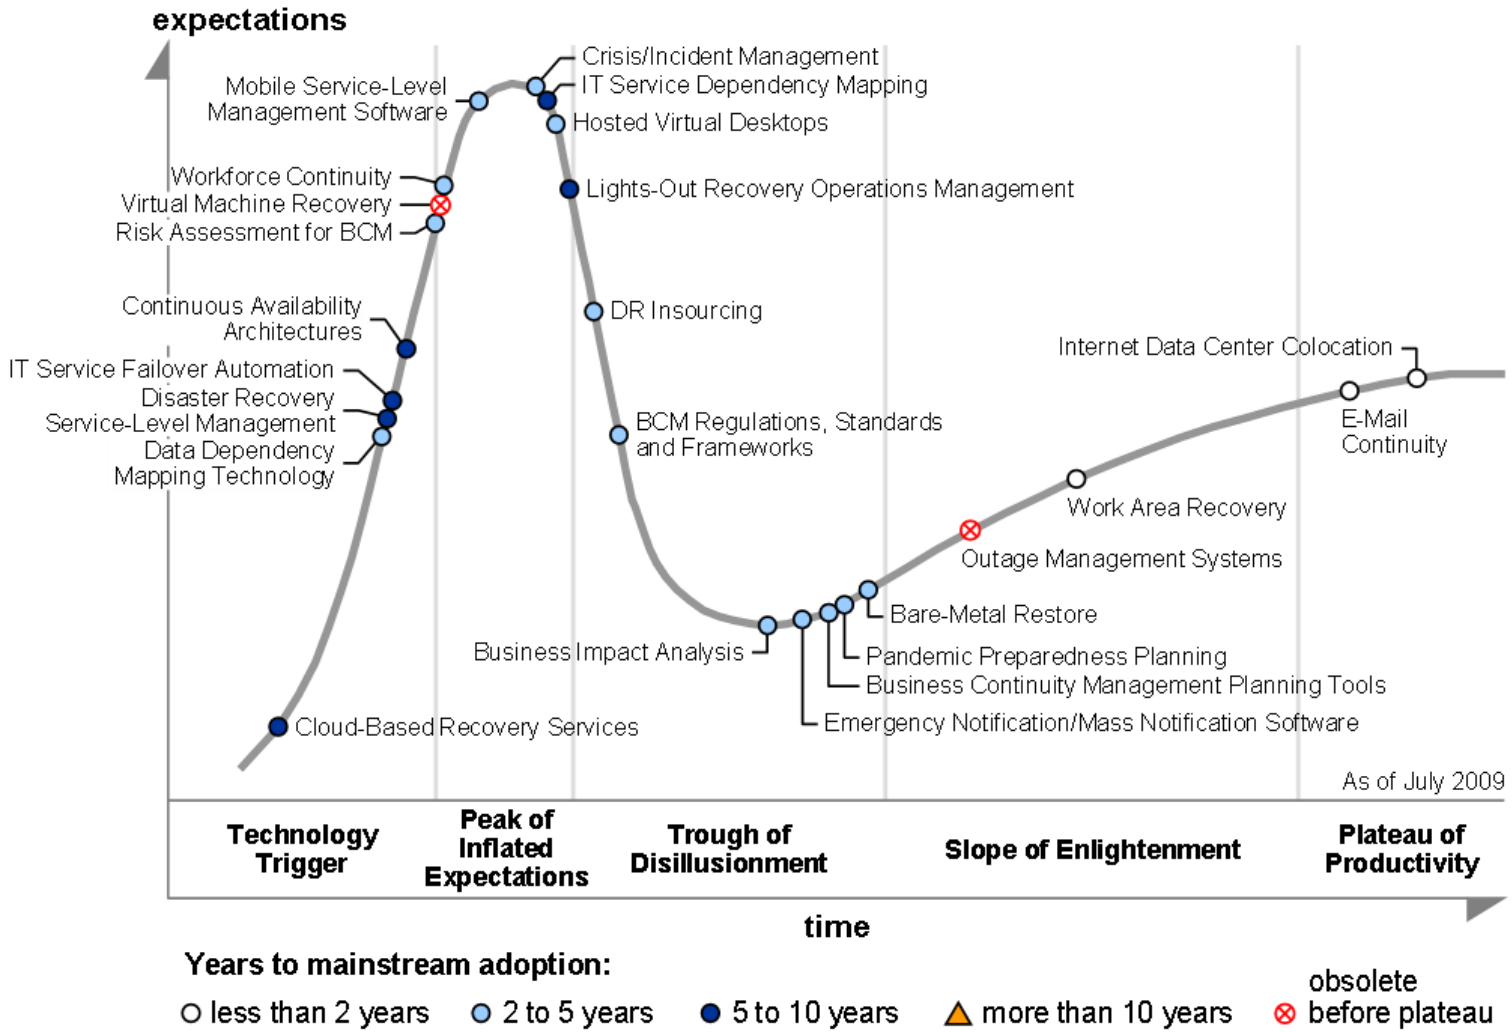

Source: Gartner (July 2009)

Hype Cycle Phases, Benefit Ratings and Maturity Levels

Table 1. Hype Cycle Phases

Phase	Definition
Technology Trigger	A breakthrough, public demonstration, product launch or other event generates significant press and industry interest.
Peak of Inflated Expectations	During this phase of overenthusiasm and unrealistic projections, a flurry of well-publicized activity by technology leaders results in some successes, but more failures, as the technology is pushed to its limits. The only enterprises making money are conference organizers and magazine publishers.
Trough of Disillusionment	Because the technology does not live up to its overinflated expectations, it rapidly becomes unfashionable. Media interest wanes, except for a few cautionary tales.
Slope of Enlightenment	Focused experimentation and solid hard work by an increasingly diverse range of organizations lead to a true understanding of the technology's applicability, risks and benefits. Commercial off-the-shelf methodologies and tools ease the development process.
Plateau of Productivity	The real-world benefits of the technology are demonstrated and accepted. Tools and methodologies are increasingly stable as they enter their second and third generations. Growing numbers of organizations feel comfortable with the reduced level of risk; the rapid growth phase of adoption begins. Approximately 20% of the technology's target audience has adopted or is adopting the technology as it enters this phase.
Years to Mainstream Adoption	The time required for the technology to reach the Plateau of Productivity.

Source: Gartner (August 2010)

Table 2. Benefit Ratings

Benefit Rating	Definition
Transformational	Enables new ways of doing business across industries that will result in major shifts in industry dynamics
High	Enables new ways of performing horizontal or vertical processes that will result in significantly increased revenue or cost savings for an enterprise
Moderate	Provides incremental improvements to established processes that will result in increased revenue or cost savings for an enterprise

Benefit Rating	Definition
Low	Slightly improves processes (for example, improved user experience) that will be difficult to translate into increased revenue or cost savings

Source: Gartner (August 2010)

Table 3. Maturity Levels

Maturity Level	Status	Products/Vendors
Embryonic	• In labs	• None
Emerging	• Commercialization by vendors Pilots and deployments by industry leaders	• First generation High price Much customization
Adolescent	• Maturing technology capabilities and process understanding Uptake beyond early adopters	• Second generation Less customization
Early mainstream	• Proven technology Vendors, technology and adoption rapidly evolving	• Third generation More out of box Methodologies
Mature mainstream	• Robust technology Not much evolution in vendors or technology	• Several dominant vendors
Legacy	• Not appropriate for new developments Cost of migration constrains replacement	• Maintenance revenue focus
Obsolete	• Rarely used	• Used/resale market only

Source: Gartner (August 2010)

RECOMMENDED READING

"Research Roundup: Business Continuity Management and IT Disaster Recovery Management, 2Q10"

"Understanding Gartner's Hype Cycles, 2010"

"Gartner for IT Leaders Overview: The Business Continuity Manager"

"How the Business Continuity Management Professional Can Survive the Worldwide Economic Crisis"

"Activity Cycle Overview: Business Continuity Manager Role, 2010 to 2011"

"Toolkit: BCM Governance and Implementation Responsibility Decision Matrix, 2010"

"Predicts 2010: The Role of Business Continuity Management Continues to Expand and Extend"

"IT-DRM Modernization Is Driving Increased Disaster Recovery Spending"

"Predicts 2010: New IT Disaster Recovery Technologies Are Emerging, but Most Are in the Early Stages"

"Data Center Conference Attendees Are Bullish on Virtualization and Cloud Computing for Improving Application Services Recovery and Availability"

"Cool Vendors in Storage Technologies, 2010"

"Data Deduplication Will Be Even Bigger in 2010"

"How to Understand and Select Business Continuity Management Software"

"Toolkit: Business Continuity Management Charter Best Practices and Template"

"Toolkit: Requirements for Crisis Command and Emergency Operations Centers"

"IT Disaster Recovery Sourcing Considerations"

"Define, Develop and Verify Plans for Application Availability and Recoverability"

"Toolkit Best Practice: Disaster Recovery Service Levels: What Makes Them Different and Why They Are Important"

"How to Calculate the Cost of Continuously Available IT Services"

"Backup and Recovery in a Server-Virtualized World"

"IT Service Dependency Mapping Tools Provide Configuration View"

"MarketScope for Emergency and Mass Notification Services"

"MarketScope for Business Continuity Management Planning Software"

REGIONAL HEADQUARTERS

Corporate Headquarters
56 Top Gallant Road
Stamford, CT 06902-7700
U.S.A.
+1 203 964 0096

European Headquarters
Tamesis
The Glanty
Egham
Surrey, TW20 9AW
UNITED KINGDOM
+44 1784 431611

Asia/Pacific Headquarters
Gartner Australasia Pty. Ltd.
Level 9, 141 Walker Street
North Sydney
New South Wales 2060
AUSTRALIA
+61 2 9459 4600

Japan Headquarters
Gartner Japan Ltd.
Aobadai Hills, 6F
7-7, Aobadai, 4-chome
Meguro-ku, Tokyo 153-0042
JAPAN
+81 3 3481 3670

Latin America Headquarters
Gartner do Brazil
Av. das Nações Unidas, 12551
9° andar—World Trade Center
04578-903—São Paulo SP
BRAZIL
+55 11 3443 1509

Gartner