White Paper

Azure Site Recovery: Automated VM Provisioning and Failover

Sponsored by: Microsoft

Peter Rutten Al Gillen

March 2015

EXECUTIVE SUMMARY

Data is the currency of modern business, and vigorously protecting this data is imperative, regardless of industry or company size — from small and medium-sized business (SMB) to enterprise to hosted service provider (HSP). The penalty for prolonged downtime or data loss can be severe. In the 3rd Platform era, in which mobile, big data, cloud, and social networking are transforming IT, data of systems of record is being joined by a tsunami of data from systems of engagement, and together they form the basis for data of insight from which competitive advantage is obtained.

In this evolving environment, data protection means ensuring that the infrastructure — whether onsite or in the cloud — that captures and processes data cannot go down (high availability, or HA); ensuring that if there is an infrastructure problem, no data is lost (disaster recovery, or DR); and guaranteeing that if data is lost, the most recent copies possible of that data exist somewhere, can be retrieved, and can be brought online instantly to maintain uninterrupted business continuity (backup and recovery).

The majority of mission-critical workloads are running on x86 systems today, and classic availability software that has been used in the past primarily on RISC-based Unix servers can be costly to acquire and complex to implement on x86 servers. Further, implementation of classic HA software often required built-in application awareness. These complexities caused enough financial and resource burden that medium-sized firms, and even some branch offices or remote sites of larger organizations, elected to avoid the expense and live dangerously — without an availability solution. But the consequences of an outage for these firms can be dire, including potential bankruptcy due to permanent data loss. Larger firms often choose to protect only their first-tier applications, not the lower tiers. But the ripple effects of a second-tier outage can still cause damage and significant pain for IT staff.

Today's datacenter is in some ways simplified, but in other ways complicated, by the virtualization of server hardware (compute and storage) and by the increasing role of the cloud — public, private, or hybrid. What is emerging in the industry today is a continuum of physical-virtual-cloud compute, and accordingly, solutions for HA, DR, and backup have started to morph and become more integrated in terms of their features and their underlying technologies.

Load balancing, failover clustering, and disaster recovery are no longer restricted to customer site to site or customer site to DR provider. They are also no longer restricted to physical servers, nor do they need to take place on-premises, partially or entirely. Another change is that DR is no longer restricted to blocks or files — it can now apply to applications and entire virtual machine (VMs). DR can occur within one hypervisor, between different hypervisors, between on-premises VMs and the cloud, and soon even between physical servers and the cloud. In other words, site to cloud is emerging as a viable alternative as virtualization and virtualization management are becoming closely intertwined with DR and DR as a service (DRaaS). IDC expects that future solutions will be aimed at further strengthening that continuum.

These technological advancements are hugely important for enabling firms to safely traverse the 3rd Platform shifts. DRaaS will allow SMBs to obtain levels of protection that they could not afford until now, and it will allow large firms to begin protecting other data than just tier 1 at acceptable cost. Large enterprises might even consider investigating a switch from using a DR provider to DRaaS when their DR contract is up for renewal.

This White Paper discusses how DRaaS has come to fruition and takes an in-depth look at one specific DRaaS solution: Microsoft Azure Site Recovery (ASR), which enables the failover of Hyper-V and vSphere VMs, as well as physical instances that are running on-premises, to Microsoft Azure.

IN THIS WHITE PAPER

This IDC White Paper discusses the changing nature of classic high-availability software, the growth in failover and restart solutions, and disaster recovery for virtual and physical machines as well as for data sets associated with virtualized and nonvirtualized datacenters. This White Paper also discusses Microsoft Azure Site Recovery, which is emerging as a solution that may be a good fit for existing and future Windows customers.

MARKET OVERVIEW

Welcome to the digital era, where businesses' fortunes hinge on the ability to take advantage of the surge in mobile transactions, making sense of the unstoppable growth of data sloshing through global networks and leveraging the emerging Internet of Things (IoT). This overload of information presents unprecedented opportunities that arise from capturing and analyzing data.

Capturing these transactions, cataloging and analyzing data, and responding to a changing business environment demand high availability of many business systems, making a widely consumable, easily implemented, and cost-effective HA environment a critical component of enterprise datacenters — as well as at IT facilities in SMBs. Unplanned downtime can have damaging effects on a business' customer relations, revenue, reputation, and regulatory compliance, with even a short amount of downtime costing an organization from thousands to hundreds of thousands of dollars, or more.

In-depth survey-based IDC research shows that the average annual revenue loss per hour of downtime in midsize companies varies significantly by industry sector: nearly $60,000 for manufacturing firms, $158,000 for healthcare businesses, as much as $400,000 for retail businesses, and nearly $10 million for financial firms. These results are from research within the United States, but samples determined that they are consistent with studies that IDC has conducted in other regions.

IT professionals at medium-sized and large businesses are trying to evolve their facilities and datacenters to keep up with stakeholder demands — from lines of business to customers to suppliers to shareholders — and respond to service requirements calling for "always on" operations.

The majority of mission-critical workloads are running on x86 systems today, and classic availability software that has been used in the past can be costly to acquire and complex to implement. Because of this financial and resource burden, medium-sized firms sometimes elect to avoid the expense and live dangerously — without an availability solution. But the consequences of an outage for these firms can be dire, including bankruptcy due to permanent data loss. Larger firms often choose to protect only their first-tier applications, not the lower tiers. The ripple effects of a second-tier outage can still cause damage and significant pain for IT staff.

Current Availability and Clustering Software Market Trends

IDC distinguishes between three types of availability software: load balancing, failover clustering, and disaster recovery. All three can be used both within a virtualized environment and exclusive of a virtualized environment.

Load Balancing

The primary task of load balancers, whether appliances, virtual appliances, software based, or a feature within an operating system (OS), is routing traffic to the most available server in a server farm or server cluster, as fast and efficiently as possible, and monitoring the availability of these servers to determine which ones are performing well, which ones are getting too much load and need to be relieved, and which ones may have dropped out entirely.

Load balancers can also distribute load based on the resources available on a server, start up standby servers if necessary, and buffer content being routed to a slow client so as to sustain performance, besides offering other features. Some load balancers can start up, turn off, and migrate virtual machines depending on network performance, and there are load balancers that can manage globally distributed traffic and make sure that the traffic is always routed to servers with the greatest resources and availability.

Load balancers exist as open source software and as proprietary solutions marketed by multiple vendors. Microsoft's Windows Server 2012 and 2012 R2 have a network load balancing (NLB) feature. The KVM hypervisor can be used with oVirt to create VM load balancing and is included in Red Hat Enterprise Linux and other Linux distributions. KEMP Technologies markets various load balancing solutions and is also the underlying load balancing technology within VMware's vSphere. Other vendors are selling load balancers with feature sets targeted at specific customer needs.

Failover Clustering

Failover clustering software delivers high availability by weaving multiple identically configured nodes together into a cluster. This is a very common approach in datacenters today. The cluster is designed to sense, via a "heartbeat" signal, that a node is failing (or has failed) and immediately acts to migrate the operational workloads on that node to a healthy node, maintaining full data integrity in the process. Often, a user will not notice this "failover" took place. Clusters, however, are costly because of the hardware redundancy involved, and in some cases, awareness needs to be built into the application software.

A cluster of two nodes, for example, requires two processors, special power switching technology to turn a node off if necessary, connections from both nodes to shared storage disks or a SAN, and connections from all nodes to a LAN. If the cluster stretches across remote datacenters, a WAN is also required. This typically means more than twice the capex than for an unprotected server.

Other complicating factors are managing which node has access to the shared resources and preventing "split brain," which occurs when two instances of a cluster are accessing data while unaware of each other's existence. Clusters can also contain single points of failure in their SAN fabric or network configuration. The three largest nonhardware vendors of clustering software are Microsoft (various products as part of Windows Server), Symantec (Symantec Cluster Server), and VMware (vSphere HA).

Disaster Recovery

Disaster recovery is based on replication, a technology with a rich history. Replication mirrors data across a network, either in real time (continuous replication) or at intervals (snapshot-based replication). The technology is typically used to move data from a local source location to one or more remote target locations.

Replication is usually block based, moving data from one physical storage array to another to provide DR with less data loss and shorter time to recovery than was possible with tape. It used to be synchronous — meaning that the source must receive confirmation that the data has been written to the target before sending the next block — and therefore slow. Today, replication can be file based, application based, and even VM based.

Furthermore, asynchronous replication allows for data to be written to the target as the source moves on to the next block while a log keeps track of synchronicity, making the process fast enough to bridge long distances with fewer latency-related slowdowns. As a result, replication over a WAN has become an ideal underlying technology for DR to prevent data loss in the event of a regional catastrophe.

There are various standalone availability and clustering products from software vendors, large and small, across these categories as well as solutions developed by hardware manufacturers (such as IBM, HP, Oracle, and Fujitsu) that are specific to their platforms. It should be noted that vendors often combine these functionalities and offer them as features of a larger package, such as system management or middleware, or combine them with yet more functionality such as replication or disaster recovery.

The Trend Toward Virtualization and Cloud

All three approaches described previously can increasingly apply to virtualized environments. For example, guest clustering delivers high availability based on more or less the same principles as a physically clustered environment, except that in this case, virtual machines are combined into a cluster (or pool). The shared storage that the clustered VMs are connected to can also be virtual.

Virtual HA will respond to a failure in the physical hardware or in a VM, respectively failing over to a healthy VM on the same hardware or to a VM on another physical node.

VMs can also fail over to a VM on a physical server in a remote datacenter to maintain business continuity in case of a regional outage. Increasingly, VMs on one type of hypervisor can now even fail over to a VM on another within a multihypervisor environment. However, in many cases, hardware redundancy — and the associated extra cost — still plays a role in a virtualized environment that must be highly available. The exception is where organizations are able to leverage underutilized servers for the purpose of providing a failover host for VMs running elsewhere in the corporate environment. All popular hypervisors, including Microsoft Hyper-V, VMware vSphere, and Citrix XenServer, offer HA capabilities.

Replication and DR have their origins in maintaining data integrity and availability of storage-based data. However, in the virtualized environments of today, the distinctions between protecting a VM (or a group of VMs on a given server) and replicating one or more VMs so as to recover the data a VM contains in case of a disaster are becoming blurry.

Hypervisors, it turns out, are just as capable of replicating a VM as they are of executing a VM failover. What's more, disaster recovery has been among the first HA technologies to blaze a trail to the public cloud and transform into DRaaS. DRaaS eliminates the hardware redundancy required with other solutions, replicating VMs to the cloud, as we see in the next section.

Differences in DR Adoption and Utilization Between Enterprises and SMBs

Enterprises and SMBs have different considerations when adopting and utilizing DR, including:

- **Use.** Enterprises may purchase dedicated DR solutions, while SMBs may be looking for solutions that combine various replication technologies, such as data protection, migration, and content distribution.
- **Business objectives.** Enterprises are typically operating with more demanding recovery point objective (the maximum period during which data may be lost) and recovery time objective (the maximum amount of time until service must be restored) requirements than SMBs.
- **Types of applications.** Enterprises have a broader range of applications that need to be replicated, including custom and legacy applications.
- **Availability of target locations.** While enterprises may have physical source and target locations that are miles or hundreds of miles apart, many SMBs cannot afford the extra expense of a physical target location; also, the distance between these locations plays a role because the longer the distance, the more bandwidth that is required.

- **Data volume.** Enterprises typically process greater data volumes than SMBs and may have greater bandwidth to move those data sets than SMBs. DR solutions need to support that volume.
- **Support for hypervisors and OSs.** Enterprises tend to have most or all of the major operating environments within their infrastructure, while SMBs may have selected a fairly homogeneous environment.

Based on these considerations, enterprises and SMBs will typically select different DR solutions, with cost, scalability, the overhead of a host, and ease of integration being primary drivers. SMBs achieve acceptable DR with host-based solutions, which use servers to copy data from one site to another in a homogeneous environment, or hypervisor-based solutions, in which one dedicated VM per host takes care of the replication function for all VMs on that host. Both approaches are low cost and easy to integrate. Host-based solutions provide SMBs with medium levels of scalability, while scalability with hypervisor-based replication tends to be low.

Fast-growing SMBs can choose a more costly approach with appliance-based or application-based solutions, each of which will provide high levels of scalability — the first approach without the overhead of a host and the second approach with the overhead. Appliances support heterogeneous environments and don't require a host, but — of course — they require the purchase of the appliance. Virtual appliances, which are more affordable and easier to deploy, albeit offering less scalability and performance, are becoming popular among SMBs.

Enterprises have been relying and will continue to rely heavily on using expensive array-based DR, in which replication is executed on the storage array, relieving the host from the replication overhead. Array-based DR is also popular for replicating multiple applications and/or large volumes of data.

The Evolution Toward DRaaS

Enterprises have gone through various technological shifts over time to keep up with the increasing demand on their infrastructures while enhancing availability and data protection. As a result, various approaches and various generations of technical solutions coexist in the datacenter, which has led to increasingly difficult-to-manage complexity. Disaster recovery via replication is just one of many protection technologies deployed.

Enterprise to Enterprise

DR has traditionally required site-to-site replication, and site-to-site replication is based on doubling the infrastructure: a production site with all the applications, server hardware, storage hardware, and networking necessary for business operations and a nearly identical recovery site that can instantly and automatically take over any and all functions that fail on the production site with minimal data loss. What this means for the enterprise is a doubling of the aforementioned complexity, especially if various operating systems and virtual platforms are deployed, as well as a doubling of the related capex and opex.

To avoid these investments and to simplify IT operations, enterprises have increasingly been attracted to disaster recovery providers that offer backup and disaster recovery services, whereby the enterprise pays an annual fee for a dedicated section of the provider's infrastructure to deliver DR and backup services. Most of these providers have been around for many years, are mature, and provide robust services, and their infrastructures facilitate easier scaling than with a site-to-site environment. Also, most of them will contractually guarantee levels of data protection, something that is a sought-after sense of security for large firms. They are not cheap, however, and tend to be mostly suitable for large enterprises – SMBs typically cannot afford them.

Enterprise to Cloud

Increasingly, traditional DR providers have been experiencing competition from cloud providers with a DRaaS offering, a service for failover of on-premises as well as cloud infrastructure to their cloud using replication, and from pure-play DRaaS providers. These DRaaS providers, including DR provider incumbents IBM and SunGard as well as a number of recent entries and start-ups, have been marketing faster recovery times at lower cost and with more flexible terms; interest in these modern offerings is growing rapidly.

In a DRaaS environment, customers typically pay for only the storage component in the cloud, as cloud-based VMs remain dormant until they need to be booted up for a recovery. A random list of participants in this market includes Amazon, Google, Rackspace, Zerto, Microsoft, Egenera, Windstream, and Symantec.

DRaaS vendors list as drivers for their services greater resiliency of the infrastructure, improved service levels, lower opex and capex, easier testing, faster time to market, and a simple pay-per-use model. Some observed limitations are that DRaaS environments are still maturing and that not all providers are fully equipped to provide failover, failback, and testing automation; replicate between various storage technologies; replicate between different hypervisors; provide high levels of security; and ensure compliance. Specialized firms, however, are working on improving DRaaS provider infrastructures. Another concern is that from an economics perspective, DRaaS may be a viable alternative for enterprises, but just like traditional DR providers, DRaaS can be costly for medium-sized companies.

Automated VM Provisioning and Failover

As mentioned previously, replication has evolved substantially and is now used not only for HA/DR, migration, content distribution, and data protection but also for live migration of a VM from one hypervisor to another while automatically converting its data. As part of this evolution, replication technology has been integrated into other solutions, including hypervisors and HA/DR solutions. Virtualization vendors such as Microsoft and VMware have begun providing hypervisor-based replication, a highly available host-level function that can be used to replicate VMs at the VM level. For simple management, the replication function has been integrated into orchestration tools such as Microsoft System Center Virtual Machine Manager (VMM) and VMware vCenter.

The Role of Backup as a Service

It is important to mention backup as a service (BaaS) within the context of DRaaS because the former is closely related to the latter, or should be, to ensure a holistic data protection strategy. Among the multitude of data protection systems that enterprises typically deploy are a variety of legacy and modern backup strategies, including traditional forms of backup, use of tape backup systems with often complicated rotation schemes (still gaining in absolute terms but declining as a share of raw storage capacity), increasing use of disk backup, use of snapshots to obtain frozen-in-time copies of a data set so that writing to the data can continue, and use of service providers for backups and vaulting. It is not unusual for IT to dedicate substantial resources to manage these mixed and complicated environments and still experience data loss or prolonged unavailability of data, not in the least because the volume of data that needs to be protected is burgeoning.

IDC research shows that the use of disk-based data protection instead of tape has increased to improve recovery and to reduce or even eliminate the backup window. Often a disk-based solution goes hand in hand with a deduplication solution, either as software or as a purpose-built backup appliance, allowing firms to greatly reduce their reliance on tape-based backup by reducing the overall storage footprint through deduplication. More recently, the adoption of the cloud for data protection has grown rapidly for reasons of cost (lower capex), ease of use, scalability, better integration with a virtualized environment, and better service-level agreements (SLAs).

BaaS providers deliver cloud-based backup services, including the infrastructure, the applications, the necessary processes, the overall management, and the SLAs, as required by the customer. BaaS providers may use public clouds (Google, Amazon, Microsoft), private clouds (the provider manages the enterprise's private cloud), or hybrid clouds (for example, an onsite deduplication appliance combined with backup in the cloud).

FOCUS ON MICROSOFT AZURE SITE RECOVERY

In response to many customer requests for DR on Microsoft Azure, Microsoft developed Azure Site Recovery (ASR), which enables the failover of Hyper-V and VMware VMs, as well as physical instances, that are running on-premises to Microsoft Azure. Azure Site Recovery protects mission-critical applications with automated replication-based DR of physical and virtual machines. Servers can be protected to targets that are on-premises, at a hosting service provider, or on the Azure cloud. Microsoft is essentially entering the DRaaS market with ASR by enabling failover of a VM or physical machine to Azure, providing customers — SMBs, enterprises, and HSPs — with the ability to achieve DR without needing to invest in a costly duplicate infrastructure.

Data replication with ASR can be fully automated through the policies set by IT and then executed in coordination with the following well-known Windows Server technologies:

- **Hyper-V Replica** is a built-in replication technology at the host level that asynchronously replicates a VM from a source site to a target site across a LAN and/or a WAN.
- **System Center** is a management tool for infrastructure provisioning and monitoring, application performance monitoring, automation, and so on across on-premises, HSP, and Azure environments. System Center is not a prerequisite for Hyper-V replication of on-premises VMs to Azure, which is good news for SMB customers that do not always run System Center.
- **SQL Server AlwaysOn** is a tool that provides flexible options for selecting the appropriate type of high availability and DR for an application.

- **Array-based replication** takes advantage of the replication capabilities (as a feature or an add-on) of various leading SAN array products that customers — especially enterprises — may have installed already from such vendors as EMC, NetApp, and HP, which Microsoft leverages by integrating that replication channel with System Center.
- **ASR-integrated InMage** technology includes site-to-site replication and failover of physical servers as well as site-to-site replication and failover of VMware VMs. Microsoft has announced that by mid-2015, ASR with InMage technology will also enable replication and failover from on-premises physical servers to Azure as well as from VMware VMs to Azure.

How ASR Works

The ASR tool itself resides on Microsoft Azure and remotely monitors VMs in a customer's datacenter on an ongoing basis. Recovery Plans, which contain IT's recovery instructions in case of an outage, such as which server and service to bring back first and how fast, are kept in the Azure Management Portal. IT has the ability to design very simple recovery plans or highly customized scenarios using PowerShell scripts.

PowerShell, an integral component of Windows Server 2012, is a task automation and configuration management framework from Microsoft with a command-line shell and a scripting language. Unlike in traditional DR environments, IT can use ASR to test recovery plans as often as desirable without causing disruptions in the operational infrastructure. The testing is noninvasive and can be done without the cost, complexity, and downtime of a traditional DR test.

ASR comes with encryption capabilities. If the VM source and recovery target are both on-premises, the data replication between source and target as well as the communications with ASR (which remains offsite) can be encrypted. If the replication target is on Microsoft Azure, then the replicated data is encrypted as well. Encryption of data at rest is also provided for DR to Azure to support customers with regulatory requirements. Replicating to Azure requires a Site Recovery Vault on Azure; however, no live VMs are needed, as a failover automatically spins up the required VMs. This is a cost benefit to customers, not only because they do not need to pay for running the VMs in Azure but also because they save on licensing fees for Microsoft workloads through DR benefits covered under Microsoft's Software Assurance (SA). For each licensed instance customers run, the SA allows them to run one instance of the software on a backup server for disaster recovery.

ASR supports customers' compliance requirements by facilitating frequent DR drills via test failovers, which, according to Microsoft, do not pose a risk to data and will not impact production. The test results are automatically captured in a presentable Excel report that meets auditing requirements.

Microsoft has also taken on one of the tougher aspects of DRaaS, the tedious task of ensuring that the networking between source and target is correct and functional. ASR provides a feature for network customization, enabling IT to map virtual networks between source and target sites.

ASR Advantages for Midsize Firms

Microsoft believes that it delivers several distinct advantages compared with other DRaaS providers. For midsize, companies, these benefits are:

- Targeted pricing, because midsize firms cannot always afford traditional DR services:
  - ASR using Hyper-V but without VMM is available without requiring System Center or System Center VMM. There are no up-front costs or termination fees, and users "pay only for what they use." This offering should also attract small businesses that typically do not have VMM or System Center.
  - ASR with VMM requires System Center (and System Center VMM), but this is arguably still a more cost-effective approach for medium-sized companies than engaging traditional DR services to protect their mission-critical workloads.
- The ability to protect as few or as many VMs as the business requires, whether 2 or 2,000, a level of scalability that Microsoft believes is a differentiator and beneficial for smaller firms.

ASR Advantages for Large Enterprises

For large enterprises, which may have tier 1 workloads well protected already, benefits include:

- ASR provides an affordable opportunity to protect lower-tier workloads and remote office and branch office applications, which are neglected at many enterprises.
- Enterprises that are about to renew their agreements with DR providers might want to do a comparative assessment between ASR and their current provider to decide whether they could benefit from Microsoft's DRaaS offering.
- Enterprise customers can take advantage of Azure Site Recovery leveraging InMage technology, described in more detail in the section that follows, which provides DR for heterogeneous IT environments.
- Enterprises can take advantage of adding storage array-based replication between SAN devices that host virtual machine data for their tier 1 workloads. With the same solution, enterprises can manage their tier 1 workloads requiring synchronous replication from array-based replication as well as other workloads that can be protected with software-based near synchronous replication.

Azure Site Recovery Leveraging InMage Technology

Last year, Microsoft added a new service to ASR that leverages InMage technology and that is designed to facilitate site-to-site replication and failover of physical servers as well as site-to-site replication and failover of VMware VMs to Azure. Microsoft is expected to announce that ASR leveraging InMage will also facilitate replication and failover from on-premises physical servers to Azure as well as from VMware VMs to Azure.

Microsoft acquired InMage in mid-2014 and has begun integrating the technology into ASR, an indication that it intends for ASR to become a disaster recovery service for heterogeneous IT environments with multiple hypervisors and various operating systems (Windows, Linux) on physical servers, enabling the workloads they run to fail over to either a secondary site or Azure.

Before the acquisition, InMage had been the underlying technology for several commercial DRaaS services, including HP and SunGard (and it continues to fulfill that role). Today, the technology can be downloaded from the Microsoft Azure Site Recovery portal with standard ASR licensing.

How ASR with InMage Works

Using InMage technology, ASR performs in-OS replication. It places an agent inside the OS of a server or a VM that is being replicated, and the agent duplicates disk writes as they happen and sends duplicated data to a dedicated-purpose server (physical or virtual). This way, the replicated server is barely used for the replication activity, which should allow for the process to be near synchronous (meaning with a latency of mere seconds). The dedicated-purpose server takes care of the caching, compressing, and encrypting of the data and then forwards the data to a target server, which has virtual disks attached to it and which writes the data to these disks. (In case of a failover, these virtual disks disconnect from the target server and establish a link with the VMs that are replicating the source VMs, providing a near instant failover response. The philosophy behind this technology is that because it is executed inside the OS, it doesn't matter what or where the OS is. It could be in a VM on any hypervisor, it could be Windows or Linux on bare metal, and it could be in a cloud, including — but not limited to — Azure.

Current ASR Scenarios

To summarize, ASR currently operates with the following replication and failover scenarios:

- On-premises VMM site to on-premises VMM site with Hyper-V replication, which requires System Center and System Center VMM, Microsoft's management tool for configuring and managing a host, networking, and storage to create and deploy VMs and services to private clouds
- On-premises to Azure with Hyper-V replication with System Center VMM
- On-premises to Azure with Hyper-V replication without System Center VMM
- On-premises VMware site to on-premises VMware site with InMage
- On-premises VMM site to on-premises VMM site with SAN replication using storage array-based replication between SAN devices that host VM data

Microsoft is expected to announce physical server to Azure as well as VMware to Azure replication and failover leveraging InMage technology. Public preview is set to begin in late March or April 2015.

Azure Backup

Microsoft Azure not only provides DRaaS to customers but also delivers BaaS; together, they represent a continuum of data protection. Azure Backup is targeted at small businesses, departments in large firms, and firms with remote offices that require reconciled backups. Both Windows Server Backup for small-scale application and data protection and System Center Data Protection Manager for a more advanced and powerful backup of physical and virtual environments have been integrated with Azure Backup. This allows customers to deploy Azure as an offsite backup depository using those same tools. Azure Backup encrypts all data, both in transit and in storage.

Microsoft believes that Azure Backup offers customers the advantage of requiring fewer onsite media that need to be kept up to date and secured. Compared with disk or tape, Azure Backup provides the advantage of a smaller footprint in the datacenter, and compared with tape alone, the Azure cloud reduces the cost and complexity of tape media as well as the cost of transporting and storing them offsite. IDC does not believe that tape will disappear as a backup medium; rather, IDC believes that with a service like Azure Backup, customers gain a broader set of options to achieve the most efficient, cost-effective, secure, and manageable mix of backup strategies for data protection, compliance, and accessibility.

FUTURE OUTLOOK

As data becomes the currency of modern business, vigorously protecting that data is imperative, regardless of industry or company size. The penalty for prolonged downtime or data loss can be severe. In the 3rd Platform era, traditional data from systems of record is being joined by a tsunami of data from systems of engagement, and together these data sets form the basis for data of insight from which competitive advantage is obtained.

In the future, site-to-site DR solutions, which require costly infrastructure redundancy, will be implemented primarily if security concerns are paramount, typically for very sensitive or mission-critical data. Site-to-site DR requires a doubling of server, storage, and networking capex and opex that most firms will consider an unappealing, or even an insurmountable, barrier. Instead, DR providers have helped firms obtain DR with guaranteed SLAs, greater scalability, little or no capex, and sufficient security for their data protection needs, albeit still at a price that only large enterprises can afford and that has led them to primarily protect their tier 1 data. Many SMBs have implemented their own mix of solutions without a second site and without a DR provider, potentially risking downtime and data loss in the event of a catastrophe.

Now, with the cloud as the target site for DR, and with DRaaS offerings emerging as a viable, more affordable, and even more scalable alternative, SMBs have an opportunity to start protecting their data in a much more robust, yet affordable, fashion, while large enterprises can protect the lower tiers of their data and applications.

What's more, DRaaS providers such as Microsoft (with ASR) will start providing a continuum from physical to virtual to cloud, allowing for the replication and failover of physical servers and VMs to VMs in the cloud, even as the latter remain dormant until automatically spun up by a failover. Microsoft ASR will also be enabling DR from within multiple hypervisor environments, Windows Hyper-V as well as VMware vSphere, and failover to multiple targets — for example, to Azure and to a secondary site.

CHALLENGES AND OPPORTUNITIES

Any new technology includes challenges and opportunities that influence the attractiveness and adoption of that technology. With regard to HA, DR, and backup, customers face the following challenges and opportunities:

- **Challenge:**
  - Continued resistance to using public cloud infrastructure because of lingering concerns over security and regulatory compliance
- **Opportunity:**
  - Microsoft has addressed many customer concerns, including having achieved certification for many types of industry regulatory compliance rules, and has provided encrypted data transfer resources to help provide further protection for data in transit. In fact, IDC sees backup, HA/failover, and DRaaS as being attractive new scenarios to justify the adoption of cloud computing as an addendum to corporate infrastructure, especially given the low cost of entry for customers to adopt cloud services to extend their existing infrastructure.
- **Challenge:**
  - An increasingly competitive DRaaS environment
- **Opportunity:**
  - Every major cloud service provider, ISP, or hoster has some form of DRaaS, HA/failover, and BaaS offerings. All vendors realize that offering this suite of services represents an on-ramp to bring customers to their respective cloud offerings. It will likely be a buyer's market, with vendors willing to wheel and deal to sign up customers so they have an early beachhead with the adoption of cloud by these customers.
- **Challenge:**
  - Complexity of implementing a cloud-based service as an extension to existing (well-understood) datacenter operations
- **Opportunity:**
  - There is considerable benefit to customers that adopt cloud for availability, DR, and backup in terms of capex savings related to hardware, storage, and network infrastructure and the ability to avoid major investments in new datacenter infrastructure when current facilities become obsolete or reach their capacity.

RECOMMENDATIONS

IDC makes the following recommendations for customers:

- **Small business customers:** Historically, small customers have suffered from limited budget, little IT bandwidth, and a lack of technical skills to implement sophisticated software to give their applications a viable availability and recovery plan. As a result, the adoption of availability and recovery solutions was typically low or nonexistent in this market segment.

  The world is changing very quickly, and the cloud-based solutions that are coming onto the market today are not only cost effective but also relatively easy to implement. Organizations that previously had few or no viable options for building a viable availability and recovery strategy now have a wealth of opportunities from which to choose.

- **Midmarket customers:** Midmarket companies have classically fallen in between small customers, where availability and recovery were effectively ignored, and large organizations, where there was a clear understanding of the use of (and value of) availability and recovery. In some cases, midmarket customers have implemented limited recovery programs, but unless they were contracted with a cloud service provider, they were unlikely to have good coverage on key application workloads.

  Like small business customers, midmarket customers are one of the great beneficiaries of the changing nature of the availability and recovery market. Where there was previously no viable strategy for a cost-effective, actionable solution, today there is.

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-insights-community.com
www.idc.com