White Paper

Amazon Web Services: Enabling Cost-Efficient Disaster Recovery Leveraging Cloud Infrastructure

By Lauren Whitehouse and Jason Buffington

January, 2012

This ESG White Paper was commissioned by Amazon Web Services LLC and is distributed under license from ESG.

Contents

The Cloud is Ideal for Disaster Recovery

Cloud infrastructure is gaining in popularity for organizations that have very real business needs that rely on information technology (IT) infrastructure but struggle with the costs—both capital and operational—of expanding their data centers. But before even discussing the cloud, it's worthwhile to settle on a definition. ESG defines a cloud infrastructure as follows:

> *A computing model in which the equipment—including servers, storage, and networking components—used to support an organization's IT operations is hosted by a service provider and made available to customers over a network, typically the Internet. The service provider owns the equipment and is responsible for housing, running, and maintaining it, with the client typically paying on a per-use basis.*

Today, more than three-quarters (82%) of organizations have plans to leverage cloud-based services to some extent over the next five years.[1] The cost, agility, and flexibility benefits are too obvious to deny, particularly for tasks such as disaster recovery (DR)[2].

DR is an ideal use case for taking advantage of the cloud. While many organizations remain cautious about placing production services in the cloud, they are often more comfortable testing those waters for DR—especially since the cloud alters the economics of DR so radically. Some organizations implement DR only for their most critical applications to minimize risk and keep expenses in check. Using the cloud enables companies to extend DR services to additional workloads, further reducing their exposure to business interruption. Others that currently operate failover sites are finding the costs skyrocketing because of continual data growth. But for many (particularly smaller) organizations, the cloud actually makes DR possible for the first time.

This paper outlines the problems organizations face when implementing DR, describes how the cloud changes the game, and provides some insight into a suite of components from Amazon Web Services (AWS) that make "DR in the cloud" simple and cost-efficient.

Risks are Abundant

Why is DR so important? Every organization is vulnerable to a range of outages and disasters. Computer viruses are everywhere; applications and disk drives are vulnerable to faults; data can become corrupted; and of course, human error is an ever-present threat. While none of these seem like "disasters," the interruptions they cause can wreak havoc on daily business activities. Given that, imagine the impact of "real" disasters such as fires, floods, power failures, or weather-related outages. Whether caused by technical problems or natural phenomena, this unplanned downtime must be immediately addressed by IT organizations in order to restore business to a fully operational state.

The absence of a DR strategy and plan comes with significant risk. Business downtime can result in huge losses in productivity. Breaches of customer service agreements, whether explicit or implied, can cause irreparable damage to reputations and/or financial assets. In ESG research on data protection, 74% of survey respondents state they could withstand three hours or less of downtime for tier-1 (business-critical) data before experiencing adverse business affects, and more than half (53%) could only tolerate one hour or less of tier-1 downtime[3].

Aggressive recovery objectives are common in today's business environment where global operations and 24/7 productivity are typical. IT must deliver on contracted recovery time objectives (RTOs, defined as the amount of time between an outage and operational resumption) and recovery point objectives (RPOs, defined as the amount of data loss that the organizations can tolerate in the transition). Once the implications of downtime and data loss

[1] Source: ESG Research Report, *Cloud Computing Adoption Trends*, May 2011.
[2] DR is defined as "the act or process of invoking pre-planned procedures after a pre-determined period of time has elapsed in order to recover critical business system functions, application servers, and applications from loss, which may include retrieving data from backup copies or replication targets as a result of a catastrophe."
[3] Source: ESG Research Report, *Data Protection Trends*, April 2010.

are understood, IT organizations can determine application availability requirements and the proper protection mechanisms to apply to ensure compliance. To meet service level agreements (SLAs), organizations often utilize various layers of data protection using a mix of replication/mirroring technologies for primary system and storage IT continuity. Snapshots and daily backups are commonly used for operational recovery and DR. However, the result is "over insurance" for top tier applications and data (resulting in extra up-front expenses that may not be needed), and "under insurance" for lower tiers (resulting in extra expenses after a disaster).

Challenges to Implementing DR

The value of DR is not in question; every organization is concerned about its ability to get back up and running after an outage or disaster. But implementing DR can be expensive and complex, as well as tedious and time-consuming. However, consider the alternative—numerous organizations have simply vanished because they were unable to get back to full operations rapidly after a disaster. Others have been impacted by lost revenue and damage to their reputation.

One example of a DR debacle is the August 2010 incident in the Commonwealth of Virginia. A major IT system failure disrupted services statewide in Virginia, affecting 27 of the 89 state's agencies. As a result of the SAN going down, 13% of the state's file servers failed. The Commonwealth lacked a continuity procedure, which delayed recovery of failed systems for 18 hours and resumption of services for over a week. State agencies, such as the Department of Motor Vehicles and the Department of Taxation, were heavily impacted by the outage and unable to process requests—or book revenue for the state—for days to weeks. Virginia's IT contractor was cited for its responsibility in the incident and recently paid the Commonwealth of Virginia $4.7M in damages to cover the cost of improvements to Virginia's network infrastructure and data protection systems, and repair the losses incurred by the outage. Having a DR strategy in place could have avoided this outcome.

Hot, Warm, or Cold Standby - Which to Choose?

Traditionally, staging for disaster recovery requires access to physical resources and data copies housed at a geographically remote location. Three common remote site strategies are: cold, warm, and hot standby sites. Their thermal descriptions match up with how long it takes—and (inversely) how much it costs—to recover applications and data, in order to resume IT operations. Cold takes longer but is less expensive, hot is faster but more costly.

- Cold standby typically involves maintaining offsite copies of data on tape, as well as access to similar hardware systems on an as-needed or rental basis for bare metal recovery. The network at a cold standby site must be configured to match the primary network configuration, including VLAN, VPN, DNS, and firewall rules. After an interruption at the primary site, cold standby restores application availability within hours to days, often with delays during the recovery processes due to tape functions being time-consuming and prone to errors. Many organizations neglect to take on the expense and effort of testing a cold standby site, and, consequently, recovery is done via trial and error—and may not succeed. Some disadvantages of a cold standby approach are that failing back is more complex, and, since systems and data are not available until needed for recovery, it cannot be used for other tasks such as testing one's DR plan or mitigating smaller system outages.
- Warm standby maintains data copies on disk (often using virtual machine images to speed restore and recovery), enabling access to infrastructure for application delivery within minutes to hours of a primary site outage. Periodic testing is easier because resources are on disk. A warm site is often used for replication or mirroring of data and systems for recovery purposes. Virtual machine images that encapsulate the operating system, application, data, and configuration settings make it simpler to synchronize between the primary and warm standby site.
- Hot standby offers the fastest restore as data and applications are maintained offsite on running systems. This enables continuous application availability within seconds of an interruption at the primary site.

Multiple application instances can be running and receiving regular updates, while only a single instance provides access to services and content. The hot standby site is available to immediately take over operations in the event of a failure at the primary site. Server clustering and synchronous replication are good DR options, but are also the most expensive to maintain. In addition, hot standby solutions require routine upgrades and maintenance to the offsite systems along with a high-speed network running between sites for replication, adding additional complexity and expense.

Consequently, certain DR strategies require a complete set of additional hardware at a geographically remote site that is capable of serving nearly the same IO demands as one's original production server farm – effectively doubling infrastructure and operational costs for many, and rendering DR impossible for some. In addition, having sufficient staff and time available to manage the remote site, in addition to the primary site, is often deemed not feasible. Some organizations are blocked from DR simply by their inability to attract or retain sufficient skilled IT personnel to manage growing storage and the associated backup, archiving, and DR scenarios.

Many companies— mostly small and mid-sized firms—don't have a remote-site option because they lack access to corporate-owned or co-located properties to house duplicate resources. ESG research indicates that less than 20% of midmarket companies (defined as those with 100 to 999 employees) have corporate-owned secondary sites available to them[4] (see Figure 1). This research also indicates that 48% of midmarket companies and 23% of enterprise organizations lease space to house DR infrastructure. Maintaining a second corporate-owned or leased site often results in high costs and significant waste. Costs include the purchase and maintenance of additional server/storage/networking infrastructure, along with the rental fees for data center floor space, energy to power and cool the systems, IT staff time, etc. In many cases, these expensive resources then sit underutilized or even idle most of the time. That is a significant expense for something that may never get used.

Figure 1. Type of DR Site – By Company Size

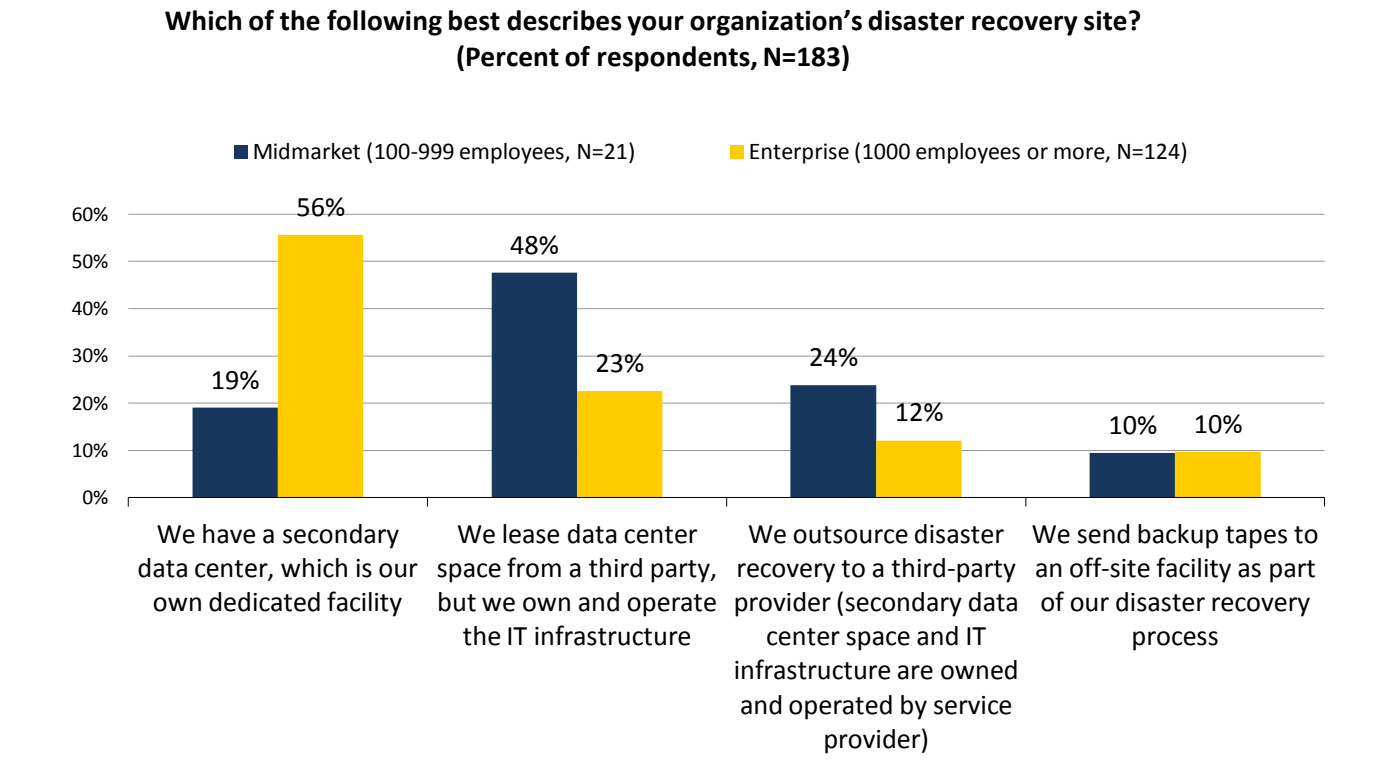

Source: Enterprise Strategy Group, 2011.

[4] Source: ESG Research Brief, *The Impact of Virtualization on Disaster Recovery*, publication pending, 2011.

DR Testing Can Be Problematic

Some consider DR testing optional because it can be inconvenient, disruptive, and expensive, but it should be a requirement. By not testing a DR plan and equipment, more risk is introduced. Since downtime may be required in testing procedures, many organizations have good intentions when it comes to testing their DR plans, but are prohibited by the need to make the DR site available, spin up new hardware, and find off-peak time (often weekends, which may require additional staff compensation) to do the testing. The Director of IT at AWS customer Sage Manufacturing described their previously time-consuming process just to set up for DR testing:

> *"In the past, spinning up a test environment would consist of procuring new hardware and turning it on. This would typically take about a month from the word 'go' to bring the hardware in and set it up. Now it takes about 12 hours."*

Testing DR systems may disrupt normal service delivery, and it can take significant time as IT organizations practice recovering a lost site, documenting the steps, and checking configurations. Oftentimes, those that do test use non-realistic workloads and subsets of relevant data in order to simplify and speed the process. However, this can render the tests unreliable.

Virtualization and Cloud Infrastructure Services Enable DR

In the face of these challenges, server virtualization and cloud-based services are like knights in shining armor. The separation of physical devices from computing and data services enables DR to be feasible where it wasn't before, primarily because of the reduced hardware costs and the associated expenses of maintaining that hardware. Virtualization reduces the amount of remote site hardware necessary, as well as the energy and management resources to run them. Cloud computing offers the ability to scale up and down easily, delivering a true "elastic" capacity that optimizes resources while minimizing expense. There are no capital expenses in a cloud-based solution. The point is to pay as you go—and pay only for what you use—and not be tied to a long-term commitment. In addition, cloud computing can improve time-to-market conditions because service creation is streamlined and delivery automatic. Organizations can focus their IT resources on business differentiators instead of spending time managing general infrastructure resources. Flexibility is also increased because virtual machines are portable. Organizations can fail back from a disaster to the original location or to a new one, in the cloud or on-site, without regard to what the original hardware was.

With cloud-based computing and storage, organizations have access to a DR platform without building one. They don't need additional corporate-owned infrastructure assets like servers and storage arrays, and they can leverage massive deployments that provide economies of scale to benefit all subscribers, whether the cloud is hosted in house or with a service provider. Redundancy can be automatic, with stored objects replicated across multiple geographies without significant additional cost.

A subscription-style model makes an "outsourced" DR services approach highly attractive, particularly in comparison to a do-it-yourself (DIY) model. A subscription-style approach makes resources available on demand, so there is no need to over provision assets. Instead, resources are more fully utilized and can seamlessly scale. In addition, using the cloud as a DR platform makes it more affordable to create a truly durable implementation by replicating systems and data across multiple geographies. The cloud can also render the DR platform geographically agnostic, eliminating time constraints and making the location of equipment irrelevant.

Of course, no conversation about cloud-based services would be of interest without discussing its fundamental benefit: reduced costs. With the cloud, the capital costs of equipment, as well as the operational costs of floor space, energy, staff, updates, and maintenance can be eliminated. As Figure 2 shows, cost reduction is the "great expectation" of IT organizations. Of those ESG survey respondents who cite the use of cloud computing as an

explicit 2011 IT cost reduction strategy, nearly one-third (31%) expect cloud services to significantly impact their IT strategies[5].

Figure 2. Expected Impact of Cloud Computing Services, By Usage of Cloud as a Cost Reduction Strategy

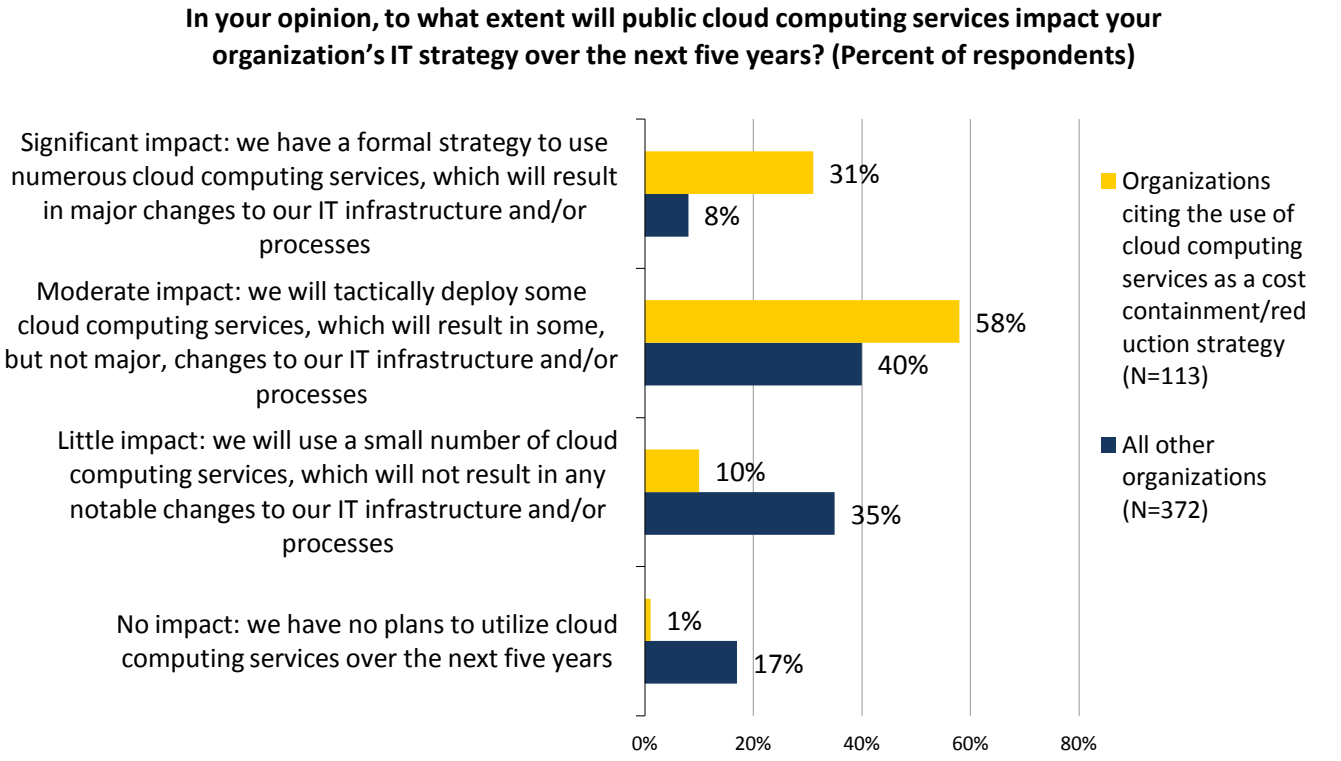

Source: Enterprise Strategy Group, 2011.

Another significant cost impact of cloud computing is the pay-as-you-go subscription model, which creates an opportunity for more predictable budgeting and eliminates the need to make a long-term commitment. Upfront capital investments are eliminated since organizations only pay for DR infrastructure when they use it. Staffing costs are minimized since the responsibility for purchasing and managing the DR infrastructure is outsourced to a cloud service provider. Monthly payments are based on services used, making it easier to budget.

Replication to a cloud infrastructure makes DR processes much simpler, and lowers RTOs. Cloud-based DR enables organizations to asynchronously replicate data and rapidly fail over to an instance of their application(s) at a cloud infrastructure service provider. Since virtual machine images can be stored there, recovery time is dramatically reduced. Applications and data can come online in minutes. ESG survey respondents confirm this, as more than half (51%) agree or strongly agree that replicating virtual machine images for DR enabled them to lower RTOs[6] (see Figure 3). Testing for DR is simpler as well. DR simulations can be created more quickly and less expensively in the cloud versus an on-premises or co-location facility. In ESG research, 43% of respondents found that deploying server virtualization simplified DR testing, with larger sites (45%) more likely than smaller sites (32%) to reap the benefits.[7]

[5] Source: ESG Research Report, *Cloud Computing Adoption Trends*, May 2011.
[6] Source: ESG Research Brief, *The Impact of Virtualization on Disaster Recovery*, publication pending 2011.
[7] Ibid.

Figure 3. Sentiment Regarding Virtual Machine Replication for DR Lowering RTOs

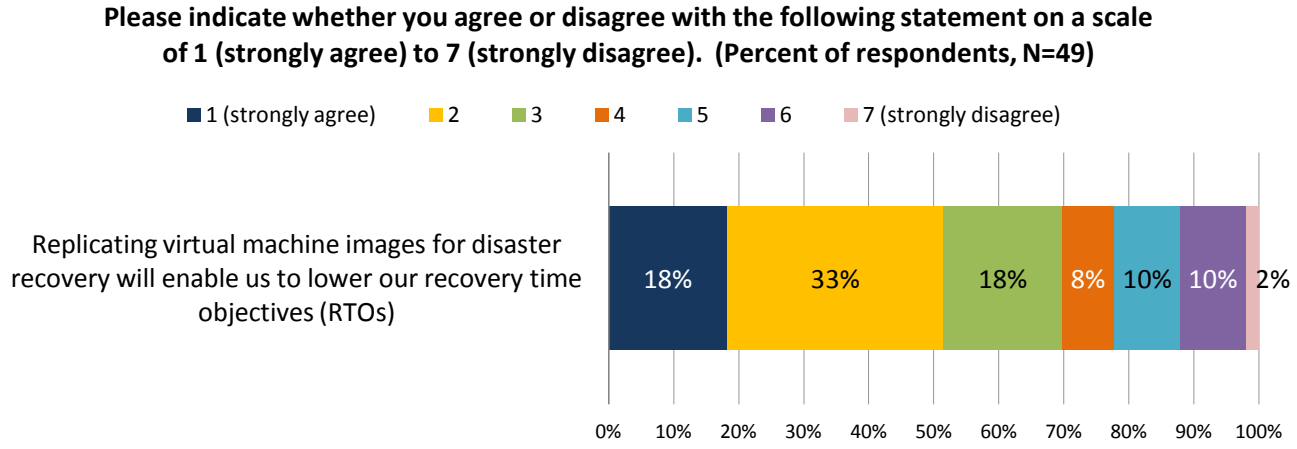

Source: Enterprise Strategy Group, 2011.

Security and compliance initiatives can also be improved in the cloud. Service providers enjoy economies of scale that make it more cost-effective to build data centers that conform to stringent security qualifications, implement extensive data encryption algorithms, and more. Lastly, many service providers specialize in delivering features and certifications that are needed to address regulatory requirements in financial and other industries. When looking at security and compliance with cloud-based solutions, ESG highly recommends that every customer conduct their own due diligence in order to fully understand any geo-political or industry regulations that might apply to how or where their data is stored.

Is the Cloud for Everyone?

Despite the obvious benefits, many organizations have not adopted cloud-based models yet. As Figure 4 shows, the key factors preventing wide-scale adoption of cloud computing tend to be issues regarding security, loss of control, and investments in existing infrastructure. Many organizations are reluctant to give up control of their assets, based on fears regarding data security and privacy. In addition, some feel that their investments in infrastructure and staff should be sufficient to do the job without adding to the cost by implementing a cloud deployment[8].

However, with the right service features, proper education, and customer validations, cloud service providers can overcome the real and the perceived barriers. Attitudes are changing as more and more organizations adopt cloud-based strategies and demonstrate their value. A common method of initiating cloud computing for organizations that have concerns is by using it to augment their existing computing environments. They may start by offloading certain applications to a software-as-a-service (SaaS) model, or by experimenting with adding compute or storage capacity in the cloud instead of buying additional on-site equipment. These are critical assets and important business processes, and organizations want to attain a certain comfort level before committing wholeheartedly to a cloud-based deployment.

Others dip their toe in the water by implementing a virtual private cloud: a private cloud that exists within a shared or public cloud. Security-sensitive companies get a flexible, secure IT resource pool transparently connected to their enterprise via a virtual private network (VPN).

[8] Source: ESG Research Report, *Cloud Computing Adoption Trends*, May 2011.

Figure 4. Factors Preventing Wide-Scale Adoption of Public Cloud Computing

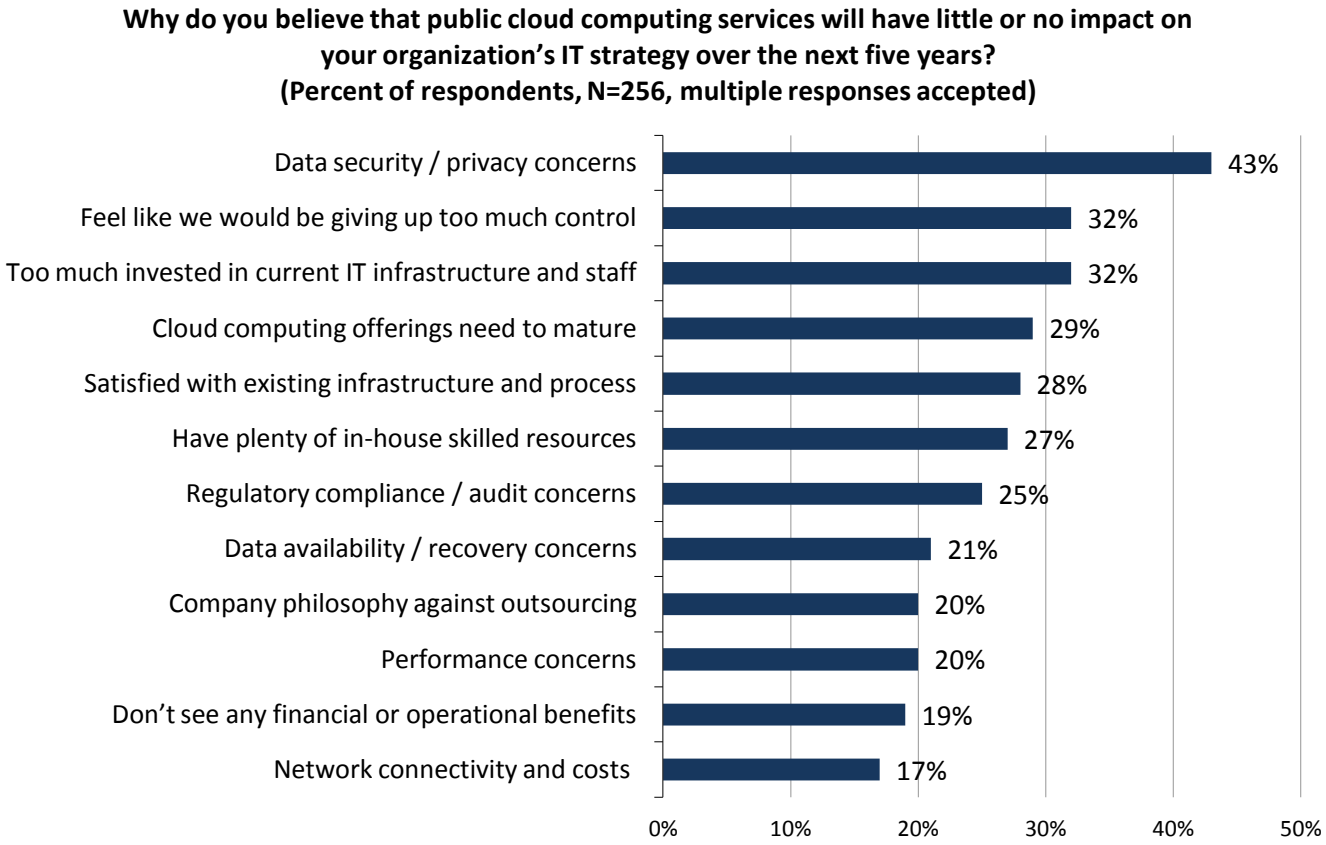

Source: Enterprise Strategy Group, 2011.

Amazon Web Services

AWS has been offering Infrastructure-as-a-Service (IaaS) to organizations of all sizes using a consumption-based business model since 2006. Consisting of a collection of modular cloud service components, AWS provides a global, elastic IT infrastructure that is well known for scalability, reliability, and security. These modular services can be used independently or combined to meet specific computing and storage requirements.

AWS makes it easy to get started. One option to seed the cloud repository is to import large amounts of data into AWS using portable storage devices transported via third-party logistics. Data can be exported in a similar fashion too. Ongoing data movement thereafter occurs using network links.

Subscribers to AWS find that cost reduction is achieved in a number of ways. First, companies will not incur costs associated with building out a second data center or committing to a collocation lease. For example, What's Up Interactive, an AWS customer, claims that it was able to avoid an estimated $1 million in outfitting a DR site. Second, there are no fees for inbound data transfer, which eliminates costs associated with moving data to a secondary and/or tertiary site. Next, AWS' pay-as-you-go pricing reduces monthly costs since only resources actually used are paid for. Finally, there are a variety of purchase options that allow you to reserve capacity for a DR failover while further lowering the cost of the utilized resources.

Durability – Amazon Web Services Regions and Availability Zones

Safeguarding data in the cloud is of utmost importance, not only for data protection but to gain customer trust. AWS Availability Zones are key to its durability for availability. When data is copied to the AWS cloud, it is stored in one of the eight global AWS Regions[9]. Within each Region are multiple Availability Zones—each in a distinct geographic area with data centers built on different floodplains, weather patterns, power grids, etc. Availability Zones are interconnected using fiber, so data is copied to each Availability Zone on fiber rather than the Internet—offering better performance and security/privacy. Each snapshot or file is copied multiple times across the Availability Zones. Combining the Availability Zones with AWS's regular integrity checks and ability to self-heal provides their "11 nines" durability design. Customers should consider this as well when calculating costs per gigabyte–most organizations develop their cost estimates assuming they will make duplicate copies of everything, but with AWS this is already in place. AWS never moves customer data outside of the user-defined Region to meet customers' regulatory and/or legal compliance requirements.

The breadth of AWS services is extensive, but three services in particular are key for DR: Amazon Simple Storage Service (Amazon S3), Amazon Elastic Compute Cloud (Amazon EC2), and Amazon Elastic Block Store (Amazon EBS).

Amazon S3

Amazon S3 is a cloud-based object store available that is through Web services interfaces such as REST and SOAP. It is used as a cloud storage container for backup data and images. Organizations can write, read, and delete virtually an unlimited number of objects containing from one byte to 5 TB of data each. Amazon S3 is similar to a traditional on-premise SAN or NAS device, but, as an AWS cloud implementation, is far more agile, flexible, and geographically redundant. Amazon S3 is designed to deliver "11 nines" of durability per year—this is accomplished by automatically making redundant copies in multiple Availability Zones, reducing the chance of data loss to one in 150 billion. Amazon S3 is also designed to offer 99.99% availability of objects, equal to just under one hour of *yearly* downtime.

No-Fee Data Import

Some of the biggest expenses of disaster recovery are those associated with physical secondary sites. When transitioning to cloud-based DR, moving backup data can represent a large up-front cost. However, AWS has eliminated fees for data import into Amazon S3. This can represent a significant cost savings at the onset, as well as long-term.

Amazon EC2

The Amazon EC2 is on-demand computing power for which subscribers pay by the hour with no long-term commitment. It enables organizations to boot up new AWS virtual servers in minutes and rapidly scale up or down. Supporting Windows, Linux, FreeBSD, and Open Solaris, Amazon EC2 can be used with all major Web and application platforms. An Amazon EC2 environment includes the operating system, services, database, and application platform stack required for a cloud-hosted application service. The virtual application stack can be started, stopped, restarted, or rebooted from a Web-based console using Web service APIs, with 99.95% availability per region using Availability Zones.

Amazon EBS

Amazon EBS provides persistent, block-level storage volumes for Amazon EC2 applications within the same availability zone. Amazon EBS is well suited to Amazon EC2 applications that require a database, file system, or access to raw block-level storage. Amazon S3 backups can be used to reconstitute applications in the cloud using Amazon EC2 computing power with attached Amazon EBS storage. In addition, data copies backed up in Amazon S3 can leverage Amazon EBS to run additional computation or test/development workloads.

[9] At the time of publication, AWS maintains eight regions: US East (N. Virginia), US West (N. California), US West (Oregon), GovCloud (US), EU (Ireland), Asia Pacific (Tokyo), Asia Pacific (Singapore), and South America (Sao Paulo).

Amazon EBS is used just like any other block storage. One or more volumes (1 GB up to 1 TB) can be mounted as devices by Amazon EC2 instances. Data can be striped across multiple Amazon EBS volumes for IO and performance improvements just like onsite storage, and data can be optionally imported or exported using physical bulk transfer instead of over network connections.

For DR, point-in-time snapshots of Amazon EBS volumes can be copied and maintained in Amazon S3 storage, limiting any data loss to what was created since the last recovery point and recovery time interruption. Its consumption-based pricing model is based on allocated volume size and IO requests, so costs are incurred only for what is used (Amazon S3 usage costs for snapshots are a separate fee).

An example demonstrates how simple it is to fail over to the AWS cloud. An application that is backed up to Amazon S3 can be copied over to Amazon EBS storage and attached to an Amazon EC2 instance. Just like that, the application with its current data set is restored and can operate in the cloud. The application is up and running while the necessary work on reconstituting the downed system or storage is performed in the data center—virtually eliminating downtime. Monthly fees to store the snapshot on Amazon S3 are incurred, as well as Amazon EC2 time and Amazon EBS storage while in use.

AWS Direct Connect

Enterprise data can also be backed up to the cloud via AWS Direct Connect, a dedicated network connection that links a primary customer data center or co-location environment directly to AWS. AWS Direct Connect allows the public Internet to be bypassed when connecting to AWS, which may reduce network costs, improve bandwidth throughput, and provide a more consistent network experience.

Amazon Virtual Private Cloud

Amazon Virtual Private Cloud (Amazon VPC) enables subscribers to create an encrypted virtual private network for connecting to the AWS cloud infrastructure, making the AWS cloud an extension of the corporate data center. Subscribers can provision an area of the AWS cloud that is isolated from others, and define a customized virtual network with complete control, including choosing private IP addresses in the range of choice, creating public or private subnets, configuring gateways and route tables, managing inbound and outbound access via control lists, and more. Resources can be seamlessly scaled, and usage fees are based on VPN connection hours and standard transfer charges.

AWS Import/Export

To seed the cloud, large amounts of data can be imported into AWS (and exported from it) using subscriber's portable storage devices transported via third-party logistics via the AWS Import/Export option. AWS transfers data directly into Amazon S3 using AWS high-speed internal network and bypassing the Internet. For large data sets, this is often more rapid than Internet transfer and could help avoid the need to upgrade connectivity.

AWS Storage Gateway

The new AWS Storage Gateway is a service that connects an on-premise software appliance with cloud-based storage. The appliance is downloaded from the Amazon website and launched on an on-premise virtualization host.

The AWS Storage Gateway enables customers to use existing applications to store data on Amazon S3 by exposing a standard iSCSI interface to the customer's on-premise application server. The customer points the application to the Gateway, which will store a primary copy locally on their on-premise storage (SAN or DAS) and also store a snapshot in the cloud on Amazon S3, as an Amazon EBS snapshot. Subsequent snapshots will store only differential changes.

Should a disaster occur, customers can restore their applications using Amazon EC2, pull their snapshots into Amazon EBS, and have the application up and running right away. Now, instead of paying for servers and infrastructure that sits idle in a disaster recovery site, customers only pay for snapshots stored in S3. If they want to recover the application in the cloud using Amazon EC2 and Amazon EBS, those subscription-based fees are only incurred when needed in the case of a disaster.

Security and Compliance

Security concerns are often a top barrier for those considering cloud implementations. Organizations depend on their data and applications to run businesses, and their level of trust that their data is secure from intrusion, as well as fully protected from loss, are critical. The economies of scale that service providers offer can make it easier for them to provide the highest levels of physical and digital security—levels that most organizations cannot implement on their own. AWS offers some of the most advanced security features in the industry.

It starts with shared responsibility between the subscriber and AWS, enabling flexibility and the levels of customer control required for certain industry-specific compliance requirements. AWS operates and manages the physical infrastructure and its security, as well as host operating systems and the virtualization layer. Customers assume the responsibility for guest operating systems and application software, including updates and security patches. Customers are also responsible for configuring the AWS-provided firewall. Subscribers can enhance security as they choose to meet more stringent compliance requirements with additional host-based firewalls and intrusion detection features, as well as encryption and encryption key management.

AWS services are built on an environment with extensive and validated security and controls, including:

- Service Organization Controls 1 (SOC 1) Type 2 report[10] (formerly SAS 70[11] Type II report), with periodic independent audits to confirm security features and controls that safeguard customer data.
- ISO 270001 Certification, an internationally-recognized security management standard that specifies leading practices and comprehensive security controls following the ISO 27002 best practice guidelines.
- PCI DSS[12] Level 1 compliance, an independent validation of the platform for the secure use of processing, transmitting, and storing credit card data.
- Relevant government agency and public sector compliance qualifications, such as an ITAR-[13]compliant environment.
- To support customers with FIPS 140-2[14] requirements, the Amazon VPC VPN endpoints and SSL-terminating load balancers in AWS GovCloud (US) operate using FIPS 140-2-validated hardware.

[10] The audit for this report is conducted in accordance with the Statement on Standards for Attestation Engagements No. 16 (SSAE 16) and the International Standards for Assurance Engagements No. 3402 (ISAE 3402) professional standards. This dual-standard report can meet a broad range of auditing requirements for U.S. and international auditing bodies. The SOC 1 report audit attests that AWS' control objectives are appropriately designed and that the individual controls defined to safeguard customer data are operating effectively.
[11] Statement on Auditing Standards No. 70
[12] Payment Card Industry Data Security Standards
[13] International Traffic in Arms Regulations, required to manage data related to defense articles or services.
[14] Federal Information Processing Standard publication; FIPS 140-2 defines four levels of security (1-4)

Leveraging AWS Components for DR

These components and features combine to make DR viable in AWS for both on-premises workloads and AWS-based workloads. Below are some standard use cases to demonstrate the range of AWS options.

- **Offsite Backup**. Amazon S3 and its "11 nines" durability design provide a good destination for backup data. On-premises data is typically transferred over the Internet and is accessible from any location. Other data movement options include AWS Direct Connect and the AWS Import/Export service. Snapshots of Amazon S3 volumes can be stored in Amazon EBS, providing a point-in-time backup that can be restored on premises or in the AWS cloud. Of course, systems running in AWS can also be backed up to Amazon S3 using snapshots of Amazon EBS volumes. In addition, numerous commercial backup services use Amazon S3 as their target.
- **"Pilot Light" Service**. In this service, the most critical elements of a computing environment are already configured in AWS for fast recovery, should the need arise. For example, core database servers would be replicating to Amazon EC2, and other infrastructure components in AWS could be quickly provisioned to restore the complete system. Other business services could be pre-configured in servers to be started up quickly. Additional Elastic IP addresses can be pre-allocated for networking, and Elastic Load Balancing can be used to distribute traffic to multiple instances to enhance performance. Less critical systems can have installation package and configuration information stored in Amazon EBS snapshots to speed application server setup. AWS CloudFormation allows users to create and manage AWS resources, automating provisioning and updates to further automate failover procedures. It converts the entire DR process into a version-controlled script—simplifying testing, restore and failover, and introducing cost savings. With AWS, automated provisioning and configuration can speed recovery time.
- **Warm Standby**. This solution extends the "Pilot Light" elements and preparation, and further decreases the recovery time by having some services always running. Business-critical systems can be duplicated and always on, running on minimal Amazon EC2 instances of the smallest possible size, making systems fully functional but not scaled to production levels. They can be scaled quickly to handle production loads in a disaster, but in the meantime can be used for testing, QA, and other non-production tasks.

These are just a few examples of the many ways that AWS components and services can provide DR. Many other services are available, and AWS will work with customers to identify the most appropriate implementations.

The Bigger Truth

DR is a critical piece of any data protection plan, but its cost and complexity often hinder deployments. There are tremendous advantages to using cloud infrastructures to reduce costs and simplify DR operations.

Services such as AWS help to eliminate upfront capital investments, since there is no need to invest in a complete set of failover infrastructure. In addition, using the cloud eliminates the need to purchase recovery and failover hardware for a secondary site. The on-demand, subscription-based model allows organizations to expand and collapse levels of infrastructure as needs change, matching costs to actual need instead of making a large investment that may never be used. In many cases the cost of each replicated server can be funded from a monthly operational budget.

Significant operational savings are generated by eliminating the need for skilled staff to procure, install, configure, and maintain a secondary site. By outsourcing to an expert such as AWS, organizations can free up their in-house staff to focus on more business-oriented projects, instead of deploying and maintaining a DR infrastructure. In addition, a cloud-based infrastructure can have far greater durability than a home-grown one because of automated geographical distribution of assets. All that an administrator needs to do is establish policies, monitor resources, and run reports from the AWS portal. It's simple to manage while offering robust and durable DR services.

For many organizations, the cloud opens the door to DR, enabling a significant reduction in risk. After years of worrying about being unprepared to deal with data loss, the emergence of cloud-based DR brings a sigh of relief. Many organizations simply do not have access to the resources or facilities needed for a failover site. Using a public cloud, DR can finally be introduced or extended to workloads that have been left unprotected. These subscription-based services make it affordable to fully protect all of their applications and data, not just the top tier.

So what are the next steps an organization should take to start or expand their DR efforts using the cloud?

- First, take stock of what would help the organization reduce risk and make outages easier to handle. Is the current protection plan sufficient? Can DR testing that is realistic and non-disruptive to production operations be executed? What are the risks? Data corruption? Component failure? Power outage? Are there specific natural challenges in the geographical area, such as seasonal tornadoes or flooding?
- Next, review compliance and security needs; levels of control required; and computing, storage, and network needs. Be sure to check on regional and industry-specific compliance requirements.
- Then, look at budget requirements. Can the organization afford self-maintained DR?
- Finally, investigate services such as AWS, and compare them for levels of security, scalability, reliability, and flexibility. Will they be easily incorporated into the current paradigm? Compare the cost of establishing a DIY standby site to one where cloud services from AWS are leveraged. How much added protection can be attained using AWS on a pay-as-you-need subscription basis versus setting up and managing a remote data center and keeping it available at all times.

While there are certainly a number of cloud-based services available, AWS offers economies of scale as well as extraordinarily high-level services. AWS has extensive investments in security and protection, and because of these features, has been selected by some of the most demanding and security conscious customers. Organizations such as NASA, the National Renewable Energy Laboratory at the U.S. Department of Energy, and NASDAQ all leverage AWS services for DR. AWS offers fine-grained control to the customer, as well as numerous building blocks designed to build the type of DR solution most appropriate given recovery time and recovery point objectives–and budget. Services are available on demand, and customers only pay for what they use, so costs remain low. In many ways, services such as AWS are the perfect match for DR, where significant infrastructure components are seldom needed, but when they are needed, they must be available quickly. The ability to spin up virtual machines and

storage and then collapse them back down when they're no longer required provides an insurance policy without breaking the bank.

AWS offers highly scalable, reliable, secure, high-performance infrastructure resources for DR, and ESG expects that many organizations, once they have experienced these services, will begin to trust the cloud enough to expand their use even beyond DR. This utility-type design can eliminate underutilization and over provisioning, while leading organizations to become more flexible in response to business challenges.

Enterprise Strategy Group | **Getting to the bigger truth.**

20 Asylum Street | Milford, MA 01757 | Tel:508.482.0188 Fax: 508.482.0218 | www.enterprisestrategygroup.com