Data protection designed for the cloud era

Ashar Baig

September 30, 2013

TABLE OF CONTENTS

Executive summary

While natural disasters account for most news headlines, they comprise only a small percentage of IT downtime. The majority of IT outages come from accidents, sabotage, and technical failures. Additionally, most data-loss events and outages are due to the failure of single hard-disk drives, machines, or servers.

These outages reinforce the critical importance of thorough disaster recovery (DR) and business continuity (BC) planning so an organization can rapidly recover from data-loss events and resume business operations as quickly as possible. Since the survival of a business depends on rapid BC, failure to develop a BC/DR plan can result in lost productivity, customers, and revenue as well as decreases in customer satisfaction, sales, reputation, and stock price.

This report examines the limitations of various cloud-based data-protection solutions in the market today and describes a best practice for enabling local area network (LAN)-powered and cloud-powered near-instantaneous BC by virtualizing the IT environment. Doing so creates both a local and a cloud copy of all applications and data, which allows employees to follow a data-loss event. It will illustrate for CIOs, IT decision-makers and managers, system administrators, and cloud service providers that:

- Constantly evolving IT infrastructures require data-protection solutions with easy-to-use interfaces. At the same time, economic considerations require moving away from IT specialists to IT generalists, as well as a deliberate shift to cohesive solutions.
- Following a disaster, the time to receive “thawed data” from a public-cloud provider, acquire new hardware, install operating systems and applications, restore files from portable storage media, and test to ensure everything works could spell the difference between an organization remaining competitive or going out of business following a disaster.
- Many cloud-enabled solutions have been architected for enterprise data-protection environments as well as cloud-washed without any significant change to the underlying technology architecture or functionality.
- Recently, some data-protection vendors have addressed the difference between data restore and data recovery to resume BC by introducing offerings that enable near-instant recovery.

Situational analysis

Data loss is an inevitable reality for IT. According to well-regarded published research about real-world data-loss events and their consequences on businesses, here are some hard facts:

- 25 percent of personal computers are forecasted to fail each year.
- 24 percent of all companies will experience a complete data loss each calendar year.
- Among those companies that experience a complete data loss, 70 percent will go out of business within the first year of that data-loss event.
- Companies that experience data loss will need an average of 30 hours to recover.

The estimated cost of downtime for a small- to mid-sized company is $74,000 per hour[1]. Hence, the average cost of data loss to a small- and medium-sized business can be $2.2 million for each data-loss event. The intangible costs, which can be even greater, come from lost productivity, lost customers, impact on reputation, falling stock price, etc. These numbers further support how critical BC planning is for rapid recovery from a data-loss event.

[1] *Source: Aberdeen Group, "Four steps to setting the right budget for downtime prevention"*
The cost of downtime: Cost per hour of downtime
All respondents: $138,000
Small companies (fewer than 100 employees): $6,900
Mid-sized (between 100 and 1,000 employees): $74,000
Large organizations (more than 1,000 employees): $1,130,000

The nuts and bolts of BC planning

Following a disaster, IT's primary objective is to resume normal business operations as quickly as possible. This business requirement has created two cloud use cases:

- High availability, which is the ability to access corporate data almost instantaneously following a data-loss event in a private-cloud or public-cloud environment
- Using the cloud as a DR storage medium instead of building a DR site

Therefore, today's DR is more than just recovering lost files. Businesses measure data-protection solutions with the DR and BC yardstick.

Backing up data is not enough anymore

Traditionally, enterprises protected against data loss by backing up data and sending a copy off-site nightly, during the infamous ever-shrinking backup window.[2] But nightly backups have some serious shortcomings:

1. **Failure-prone restores:** Because the task of nightly data backup is not highly visible inside the enterprise, IT often delegates junior or entry-level members of the IT organization to do it. However, those staff members often lack data-protection knowledge or experience, which increases the probability of restore failures. Few IT pros confess any knowledge or experience with data protection, and they rarely volunteer for the job. The result is a self-fulfilling prophecy—high failure rates for data restores due to high probability of data corruption.
2. **Lack of self-service:** End users must rely on IT specialists to conduct restores. For cloud backup, the problem is further exacerbated when corporate IT has to open a service ticket with the service provider (in order to abide by its standard workflows) and encounters limited staff that cater to dozens or hundreds of customers.
3. **Failure to complete backups:** The reason the backup window typically occurs for a few hours in the middle of the night is that nightly backups are not competing for wide area network (WAN)

[2] The backup window is not actually shrinking. Amid the tsunami of data growth whereby data is growing on an average of more than 60 percent per year, IT is asked to back up more and more data during the same backup window.

bandwidth with other high-priority applications. IT seldom increases the time allotted for the backup window. With data growing an average of more than 60 percent per year, the backup administrator is challenged to complete the backups during the allotted time.

4. **Slow backups and restores:** Backup sets and restore packages are typically very large and, when conducted over a wide area link, can be agonizingly slow.

Technology refresh: every 3 to 5 years

CIOs constantly worry about the capital-intensive nature of the ubiquitous technology refresh. These capital expenditures (capex) include new hardware along with the human capital required to perform ongoing maintenance and support, data-center cooling expenditures, real estate costs, etc. CIOs are eager to reduce their capex by graduating to exponentially more attractive operational expenditure (opex)-based Software-as-a-Service (SaaS) models. This is an entirely new business model with its own set of value propositions for enterprise data-protection environments. These value propositions include reduced costs and pay-per-use pricing models that enable a truly dynamic data-protection environment.

Need for instant BC

One critical BC planning criterion is stipulating the order for recovering each application within the DR plan and the time required for each. These recovery time objectives and recovery point objectives must be clearly stated in the DR and BC plan. The plan must factor in hardware procurement times, an alternate data-center site for BC, resources needed for BC, and the time it would take to render the recovered data usable. Therefore, all pre- and post-disaster planning and execution boils down to one thing—rapid BC following a disaster.

Cloud backup was initially designed for safe and efficient storage of operational and compliance data. Data backed up off-site to the public cloud or the DR site is de-duplicated, compressed, and often encrypted, so the data bits are stored in a deep freeze in the cloud provider's data vault to maximize storage efficiencies. This data-storage mechanism is safe but requires hours or sometimes days of "thaw time" before the data can be usable.

In the event of a disaster, this data needs to be recovered, retrieved, and groomed before BC resumes. Besides the thaw time, in most cases, the data is too big to transport across the wide area link. The public-cloud provider must recover the data on portable storage media and ship it to the customer, who then loads the de-duplicated, compressed, and encrypted data onto its servers and grooms it to resume BC. Meanwhile, the customer must acquire new hardware, install operating systems and applications, restore files from portable storage media, and test to ensure everything works. The time required for this exercise can spell the difference between the organization remaining competitive or going out of business following a disaster.

Hence, data restore and data recovery to resume BC are related but disparate concepts. Cloud data backup and restore does not mean the immediate application availability businesses require for resuming BC. Recently, some data-protection vendors have addressed these concerns by introducing offerings that enable near-instant recovery. This innovative new approach to data protection ensures BC in mere minutes after a data loss because the data and applications are always ready to run for testing or recovery.

Need for local and cloud copies of the data

The majority of data-loss events and outages are due to the failure of a single hard-disk drive, machine, or server. Therefore, the fastest way to recover from this kind of data loss is having a local copy of the data for near-instantaneous LAN-speed recovery. Cloud-powered, near-instantaneous BC can protect against natural disasters or other events that can cause complete site failures.

Typical hardware procurement times range from four to 12 weeks for most organizations, but in a disaster, new hardware may not be immediately available to restore the data recovered from the cloud. Again, in this scenario, cloud-powered, near-instantaneous BC immediately following a disaster could be the difference between a company staying in business or folding. Cloud-powered virtual BC using an internet connection is the game-changer technology for resuming business operations almost instantaneously. It should be a key component of every organization's data-protection strategy.

Proliferation of data copies

Companies create duplicate files by creating copies of data:

- When employees email their colleagues
- When employees save a file copy on their local drives (or even worse, at a consumer-grade online cloud-storage service "just in case")
- For various organizational stakeholders
- For testing, development, and support
- For short- and long-term archiving, migrations, disaster recovery, point-in-time copies (e.g., snapshots or replications)
- When they back up data to disk or tape (on-site or off-site)

This proliferation of data copies creates data sprawl that results in bloated opex and a management nightmare. Furthermore, capex balloons dramatically because these copies of primary data require more storage media and the management of these copies increases opex.

Lack of DR testing for BC assurance

DR testing or dry runs are hardware-intensive and time-consuming and take up valuable IT resources. Most backup service providers charge up to $25,000 for each DR dry-run they conduct. Due to the high cost and complexity, IT personnel rarely perform these drills, even though businesses need assurance that their applications will resume when the need arises. But without dry-runs, IT has no verification that backed-up data is actually recoverable within their standards.

Higher IT expectations equal higher data-protection expectations

Today's workers demand 100-percent uptime and 100-percent data availability. Their expectations have skyrocketed because they expect the same level of always-on service from their enterprise IT that they get

from cloud-based services like Microsoft Office 365, Facebook, LinkedIn, and Gmail. Additionally, the modern corporate workforce is geographically dispersed across worksites and time zones, and flexible schedules have become the norm. But IT has budgetary realties and cannot always meet end users' inflated expectations. As a result, user dissatisfaction increases while user expectations decrease.

IT security challenges

Employee dissatisfaction with IT has given birth to shadow IT, a term used to denote users deploying unsanctioned software or using unsanctioned devices. The primary motivator behind bring-your-own-device (BYOD) and bring-your-own-software (BYOS) is the belief that IT lacks the resources it needs to move quickly enough on an important business project or that IT does not fully grasp the evolving technological needs of the business.

Although geographically dispersed employees and partners are often not connected to the corporate LAN, IT still must secure the corporate IP on their laptops and handheld devices. This is further complicated by employee use of consumer-class cloud-storage services, like Dropbox, Box, SkyDrive, and Google Drive, instead of corporate servers. IT has less and less control over how and where information is being shared and protected.

The weaknesses of current solutions

Most data-protection solutions available in the market today were designed for enterprise environments in which data never passed the corporate firewall. These solutions possess closed proprietary hardware and software interfaces that prohibit data-protection solutions from integrating with third-party systems. They are high-touch and mostly agent-based and were designed for a single-tenant environment. They lack features that are required by cloud-based environments, such as service level agreement (SLA)-based monitoring and management, tiered storage repositories, role-based access to the storage vault, in-depth reporting/notifications/audit trail, elasticity that grows as clients and the base grow, secure data at rest and in-flight, high availability with N+1 and replication, and a web portal for centralized management.

Disparate solutions: application- or technology-specific

Today's data-protection landscape is littered with disparate siloed systems that protect corporate intellectual property. IT environments are plagued with incongruent systems that:

1. Protect physical servers, desktops, and laptops on the corporate LAN
2. Protect virtual servers, desktops, and laptops on the corporate LAN
3. Protect laptops in the field
4. Protect handheld devices like tablets and smartphones

At the same time, these IT environments require solutions for:

1. Batch data backups during the allotted nightly window for granular backups
2. Image-based data backups, e.g., snapshots and replications for faster backups
3. De-duplication for data reduction
4. Short- and long-term archiving
5. Continuous data protection (CDP) for applications that are too critical to wait for nightly backups
6. WAN-optimization for efficient transport of data off-site
7. On-premise backup for a local copy of data for LAN-speed restores to recover from accidents and equipment failures
8. Cloud backup or DR site for an off-site copy of the data

These dissimilar systems and solutions are often application-specific and often have their own isolated storage pools and antiquated management interfaces. Disparate data-protection solutions and archaic management interfaces balloon opex and have the tendency to create IT specialists who are experts in only one particular application or tool.

Amid constantly evolving IT infrastructures, robust data-protection solutions require easy-to-use interfaces that minimize the IT burden. Furthermore, the current economic environment dictates moving away from IT specialists and toward IT generalists, lowering the IT learning curve, and a deliberate shift from incongruent solutions to cohesive solutions to reduce opex.

Bolt-on solutions: byproducts of acquisitions

In another category, some larger vendors have made acquisitions to fill technology gaps, such as de-duplication, protection of virtual environments, protection of laptops in the field, image-based replication and snapshot solutions, WAN-optimization solutions, data-archiving, local backup appliances for a local copy of data, etc. However, these bolt-on solutions present disparate management interfaces and inconsistent feature sets.

Agent-based solutions: resource-intensive, difficult to deploy, and inherently unsecure

Agent software is a piece of code that is installed on each server and needs to be protected. The agent captures the information to be backed up from the source machine it resides on and sends it to the media server, which collects backup data from all agents and stores it.

Agents are resource-intensive, a pain to deploy, and inherently unsecure. A “light” agent typically uses 2 percent of the hardware and software resources (CPU, memory, backplane, etc.) during a 24-hour period. When this agent is running, it can use as much as 20 percent to 25 percent of a system’s resources. Additionally, a single agent adds more than 15-percent server overhead to each application.

Since agents are application-specific, an agent must be deployed and installed on each server it protects, which is time- and labor-intensive and carries high management overhead. Almost every agent has

administrative privileges, requiring a port in the firewall to be opened for it. This effectively creates a backdoor in the server—an invitation to tap into the agent and have your way with the server. With no "in-flight" encryption mechanisms, agents also put data at risk during transmission from the remote office to the data center.

Agents require lots of tape because they must conduct full backups once a week (typically on a weekend), with incremental backups during the week. They also carry security risks and are manpower-intensive, factors that dramatically impact the total cost of ownership and performance of the agent-based data-protection environment.

Cloud as a backup-storage target solution

The majority of the cloud-backup solutions available in the market today use the cloud as a storage target for data backed up off-site. In fact, the majority of these solutions were architected for enterprise data-protection environments, but when "cloud" became a popular term, most of those vendors changed their marketing collateral and labelled their solutions as "cloud-enabled" without any significant change to the underlying technology architecture or functionality. Under the covers, they possess the inherent limitations of single-tenant solutions; they lack varying levels of SLAs for different customers, require proprietary hardware or software that cannot integrate with third-party systems, and lack scalability. Furthermore, due to batch-processing, these legacy data-protection solutions lack the agility that newer image-based data-protection solutions have, so they often struggle to meet backup windows. What the market needs is data protection built for the cloud and that leverages the cloud to protect organizational IT infrastructure.

The solution: end-to-end data protection designed for the cloud era

Businesses now demand a complete data-protection solution, a single integrated platform that performs all the data-protection tasks listed below, coupled with an intuitive management interface. This solution must securely protect files and folders on physical and virtual servers, desktops, and laptops. The result will be a drastically simpler data-protection platform that lowers costs, the learning curve, and maintenance costs; provides better protection; and improves regulatory and corporate compliance. It would include:

- Backup and recovery
- Business continuity
- Disaster recovery
- Cloud virtualization
- Data retention
- Data tiering
- Archiving
- Cloud storage
- De-duplication
- WAN-optimization

Automation and self-service to empower end users

Today's economic reality dictates cost-cutting and doing more with less. This trend leaves IT under-staffed, regardless of the size of the business. The ideal data-protection solution is fully automated and empowers end users to conduct their own backups and restores without having to contact IT or a service provider for support.

Cloud-powered virtual BC

Following a disaster, the typical hardware procurement times exceed three months per downtime, which would put most companies out of business. The ideal data-protection solution would leverage the cloud to enable near-instant BC. This can be achieved if the data is virtualized and stored in its native format, which is directly readable by its application.

Local copy of the data

As stated previously, the majority of data loss and outages are due to the failure of a single hard-disk drive, desktop, laptop, or server. The fastest way to recover the data is to have a local copy for near-instantaneous LAN-speed recovery. As with cloud-powered virtual BC, this can be easily accomplished if the data is virtualized, stored in its native format, and directly readable by its application on a local appliance for near-instantaneous LAN-speed BC.

Agentless architecture

Eliminating agents greatly speeds up implementation time and application performance, removing the need to purchase new client licenses. The first API specifically designed for data protection was VMware's vSphere APIs for Data Protection, which the company introduced three years ago. Because it was starting from a clean slate, VMware wanted to build the optimal solution, so it examined the inherent limitations of agents and decided that the virtual data-protection environment would be agentless. Today, even the agent-based data-protection solutions are agentless when protecting virtual environments. Logically, the ideal data-protection environment thus should be agentless in both virtual and physical environments.

Image-based data protection leveraging replication and snapshots

Another technological reality is the concept of snapshots. Legacy applications typically store each backup in an encapsulated and proprietary file so that different versions of backups can be maintained. A fresh approach would be a backup solution that stores data in its native format so that access is easier and open.

Security

Data security is paramount in an ideal data-protection environment. All data, either in-flight or at-rest, should be secure. Since encryption and decryption are highly resource-intensive and often double the time it takes to back up and recover the data, the data is logically best kept in its native format to ensure rapid BC. This data must be securely transported (tunnelled) and kept in a safe environment to defend against its being compromised.

Future-proof data protection

Cloud-powered SaaS future-proofs the end customer by eliminating the need for a technology refresh every three to five years; the end customer has no need to plan for a server or storage upgrade, so capex and opex are drastically reduced. Furthermore, it eases data-migration pains by shifting the burden of data migration from in-house IT to the service provider. A SaaS solution also future-proofs companies against rapidly changing hardware and software infrastructures, APIs, and interfaces.

Eliminate downtime

Often companies focus on backup so much they forget the reason backups were conducted in the first place, which is simply to resume business operations as quickly as possible following a data-loss event. Hence, the focus should not be on backing up files. It should be on protecting applications that are vital to business resumption—accounts receivable, accounts payable, enterprise resource planning (ERP), customer relationship management (CRM), etc. By protecting these applications, IT eliminates downtime and allows employees to resume productivity following a data-loss event.

Eliminate storage inefficiencies and data sprawl: single-data copy, snapshotted

Like its various organizational stakeholders, IT makes multiple copies of production data for backup, replication, snapshots, mirroring, CDP, cloning, cloud storage, etc. It also generates dozens of mish-mashed copies of production data, which are cumbersome to manage. As a result, storage silos get littered across backup disks and tapes, virtual tape libraries, storage vaults, archive systems, snapshot repositories, virtual volumes, data-analytics platforms, and data clones.

The data sprawl caused by these copies forces companies to purchase more storage hardware and related products to protect and manage the exponentially growing volume of data copies. This untenable situation increases the storage spend by an order of magnitude.

A more efficient, less complex, more seamless, and flexible method of copy-creation is needed. By leveraging data-protection hooks, a single copy of data can be made for the entire enterprise. The data-protection solution could then leverage snapshots and replication to provide multiple point-in-time versions of the data to maintain data consistency, increase data availability, and provide longer data retention. This process also lowers capex and opex dramatically.

Those vendors who have architected their products from the ground up to leverage the cloud for DR and BC and who offer a single solution for all data-protection needs would be better positioned to serve the needs of IT, not just now but also in the future.

Data-protection best practices

Data protection should not be about backups and restores but rather about keeping the business running nonstop. An intelligent way to achieve this is by virtualizing the applications in their native format to provide local and cloud access to these applications. This methodology does not require backup windows but still provides instantaneous BC, allowing organizations always to meet recovery time objectives and recovery point objectives.

Cloud-powered BC benefits

- All-inclusive pricing delivers predictable monthly data-protection costs amid seismic data growth.
- Controlled data-center capex
    - Hardware upgrades and refreshes
    - In-house IT expertise
    - Real-estate cost for DR site
    - Data center power and cooling
    - Software upgrades
    - Ongoing maintenance, testing, and upkeep
    - Ongoing support costs
- Hardware upgrades and refreshes are user- and application-disruptive.
- Proactive 24/7/365-monitoring by cloud-service providers to detect failures
- Capacity on-demand
- Geographic diversity
- Easy and instant access to data

- Reduced infrastructure complexity
- Enables near-zero backup windows
- Instantaneous BC with no IT downtime
- De-duplication, WAN-optimization, and considerably reducing data copies drastically slashes network bandwidth and storage costs.
- Backups and restores can be initiated by end users or IT generalists.

How it works

An on-site appliance installed at the customer premises virtualizes the failed server(s) and resumes operations immediately until new hardware can be prepared. This local appliance leverages built-in image-based replication to replicate data securely and within minutes to the service provider's data center off-site. That data center itself should be secure with some or all of the certifications listed below:

1. SAS 70[3]/SSAE-16 Type II audit-certified
2. TIA 942 Tier 4-certified

The data-protection technology must leverage reverse incremental backups to store easy-to-find latest versions of backups as well as brick-level restores and to offer granularity to recover just a single specific file. In the event of a natural disaster affecting the entire site or the entire region, virtualized applications in their native format would be available instantaneously to work off of the service provider's cloud using an internet connection.

Simple unified solution

An effective data-protection platform should be engineered from the ground up as a unified platform architected to provide feature parity across all components and technologies and must provide a cohesive management interface. In the event of data loss, this intelligent data-protection platform will provide

[3] *In 2011, SAS-70 was superseded by SSAE-16, which retains the original purpose for SAS-70 compliance.*

applications on-demand instantly to all organizational stakeholders for resumption of business operations. This next-generation data-protection platform should possess the following attributes:

- **Next-generation interface**: The interface should include comprehensive management and visibility through a single pane of glass for all data-protection needs. Businesses prefer a web-based management console that is simple and easy for quick troubleshooting or instant resumption of business tasks. The graphical user interface must be designed to focus on managing by exception to showcase quickly what needs to be fixed versus looking through endless logs to pinpoint potential problems. This comprehensive visibility and intuitive management facilitates decision support and reduces the overall IT burden and cost.

- **Compliance**: Built-in security standards of this next-generation data-protection platform must meet or exceed the guidelines set forth by regulations and regulatory bodies, including the Health Insurance Portability and Accountability Act (HIPAA), Sarbanes-Oxley Act (SOX), Gramm-Leach-Billey Act (GLBA), Payment Card Industry Data Security Standard (PCI DSS), Financial Industries Regulatory Authority (FINRA), Municipal Securities Rulemaking Board (MSRB), etc.

- **Rapid release cycles**: Companies demand evolving tools that exploit the latest technologies to keep up with the fluid demands of today's always-on businesses. The data-protection vendor should have platform-nimble and agile release cycles that incorporate customer feedback and market requirements.

Summary

In today's turbulent economy, IT budgets are under constant scrutiny. The idea of using a plethora of disparate data-protection products with inconsistent management interfaces is tremendously imprudent and fiscally irresponsible. Capex and opex savings start with radical simplification of the data-protection platform to drastically reduce data sprawl created by storage silos and data copies proliferated across the corporate LAN and various cloud services.

To guard against natural disasters and other data loss events, organizations must employ a data-protection platform that facilitates instant resumption of business operations. Today, being able to recover the data and spend days of thaw time to condition that data so that it is usable by its application(s) is not enough. Businesses have no appetite or tolerance for downtime because each minute of downtime results in productivity loss, lost customers, lost revenue, and negative customer satisfaction, sales, reputation, and stock price.

A SaaS offering future-proofs enterprises against the high capex associated with the server/storage upgrades that are part of the dreadful technology refresh every three to five years and the rapidly changing hardware and software infrastructures, APIs, and interfaces.

Furthermore, an outsourced SaaS offering shifts the burden of data migration from IT to the service provider, drastically reducing opex.

In our view, the solution that best symbolizes the instant-BC vision is a unified platform that has feature parity across all components and technologies with a cohesive management interface. That management interface can provide applications on-demand instantly so that all stakeholders in the organization can resume business operations following a data-loss event. This kind of single unified solution has disrupted the data-protection industry because a single product enables a new level of agility, nimbleness, flexibility, and efficiency that is unfathomable with traditional siloed data-protection approaches. This kind of bold, fresh move can only happen when there is no legacy baggage, no legacy revenue streams, and no existing customers to support with legacy products.

Key takeaways

Traditional data-protection solutions that have been available for the last three decades are broken and require a complete overhaul.

1. Disparate data-protection solutions are cumbersome to manage; each has its own storage silos, drastically increasing capex and opex.
2. Larger data-protection vendors have filled their technology gaps with acquisitions over time. These bolt-on solutions present disparate management interfaces and inconsistent feature sets that drastically increase capex and opex.
3. The RTOs enabled by data-protection solutions in the market today are too slow and require multiple steps before the recovered data is usable, resulting in unacceptable business downtime.
4. Agent-based data-protection solutions are human- and capital-intensive, thus ballooning IT expenditures.
5. Making multiple copies of the production data generates dozens of mish-mashed copies that are cumbersome to manage, so storage silos are littered across various storage media. This data sprawl causes organizations to purchase more storage hardware and related products to protect and manage the exponentially growing volumes of data copies.
6. Restore failure rates are high because IT shies away from expensive and resource-intensive DR testing.
7. Nightly backups are not frequent enough to meet the needs of today's enterprises. In addition, they are often corrupted and result in high restore-failure rates.
8. Product limitations are further exacerbated by data-protection and data-security challenges such as rogue IT, BYOD, BYOS, and the ubiquity of consumer-class cloud-storage services.

These factors plus ballooning end-user expectations amid frequently slashed IT budgets create a data-protection and data-security perfect storm. Data-protection vendors who deliver a solution that is engineered from the ground up as a unified platform with an intuitive interface with SLA management and that enables applications on-demand instantly to all organizational stakeholders for resumption of business operations instantaneously will be tremendously successful in the battle for cloud backup.

About Ashar Baig

Ashar Baig is an industry analyst focused on cloud and storage technologies. He is the president of Analyst Connection, a full-service analyst firm that provides analyst services as well as management, marketing, and sales consulting services, including product management/product marketing advice and sales-optimization consulting. Prior to Analyst Connection, Ashar was the senior analyst and consultant at Taneja Group, where he focused on data protection, cloud storage, and public/private/hybrid cloud space. He also led Taneja Group's managed service providers (MSPs) consulting for vendors. He often advises vendors looking to target MSPs or MSPs looking to grow their revenues.

Ashar possesses more than 18 years of high-tech industry experience, having held senior roles at North America's leading companies, such as Taneja Group (consultant and industry analyst roles), Asigra (BCDR), Artisan Infrastructure (IaaS), Camouflage Software (data security), IBM (HPC and grid computing), TELUS (RBOC), HP (mobile messaging), Siemens (CRM), and Intel (networking and communication).

Ashar is passionate about driving industry standards and evangelizing them. Until recently, he was the chairman of the Special Interest Group (SIG) focused on cloud backup, recovery, and restore (BURR) within the Cloud Storage Initiative (CSI) of the Storage Networking Industry Association (SNIA). He is often referred to as "Mr. MSPs" due to his passionate evangelism and educational speaking engagements at various industry events. He also volunteers for SNIA's Analytics and Big Data Committee (ABDC) and Data Protection and Capacity Optimization (DPCO) Committee. Previously, Ashar held senior positions within EC2 and ECTF.

Ashar is the founder and manager of the LinkedIn Cloud Backup group. He actively blogs on the LinkedIn Cloud Backup group page as well as on the Analyst Connection website.

About GigaOM Research

GigaOM Research gives you insider access to expert industry insights on emerging markets. Focused on delivering highly relevant and timely research to the people who need it most, our analysis, reports, and original research come from the most respected voices in the industry. Whether you're beginning to learn about a new market or are an industry insider, GigaOM Research addresses the need for relevant, illuminating insights into the industry's most dynamic markets.

Visit us at: pro.gigaom.com.