



# Leveraging the Public Cloud for Enterprise Data Protection

## Why the public cloud has become the state-of-the-art target for backup, archiving, and disaster recovery

Workloads have been moving to the cloud for almost two decades, but only recently has the momentum in IT's commitment to the cloud begun to accelerate. For many IT organizations, the public cloud has emerged as the application platform of choice.

According to Enterprise Strategy Group's February 2015 IT Spending Intentions Survey<sup>1</sup>, three-quarters of all respondents currently utilizing or planning to adopt have at least one public cloud service, up almost 30 percent from 2013 usage levels. The report further notes that when all forms of public cloud computing are combined (IaaS, SaaS, PaaS), they become a "key area of focus" for 42 percent of IT organizations, with spending on such deployments likely to increase.

The reasons are obvious. Many companies' initial forays into the cloud have been successful. Early concerns about reliability, security, and scalability have faded. Enterprises clearly see the transformational potential of moving to cloud and are seeking greater efficiency and agility in their IT processes.

Given this success, the question then becomes: What other business processes can enterprises move to the cloud with equal or greater value? The answer is a process that is even older than the cloud: backup and disaster recovery (DR).

### Why Cloud-based Data Protection Makes Sense

Enterprises may have initially balked at the idea of entrusting backup and disaster recovery to the cloud. But the increasing success of cloud usage and IT organizations' growing familiarity and comfort with the cloud eliminate those concerns when it comes to secondary storage. Furthermore, the global ubiquity of the public cloud actually brings a higher level of efficiency to the challenge of backup and recovery.

- **Cost of Cloud Storage and TCO.** Cloud storage saves money and resources. Cloud computing falls under operational expenses (OPEX) rather than capital expenditures (CAPEX), and with storage, enterprises pay only for what they use. This

<sup>1</sup> ESG IT Spending Intentions Report, Feb 2015

helps enterprises align their capacity and cost models, and improves total cost of ownership (TCO). In the traditional model, hardware expenditures and storage compute investments are made episodically based on predicted demand. In reality, this can create unused storage capacity, or leave businesses short on capacity based on actual demand, placing service availability at risk. With cloud, businesses can adopt a capacity-cost model, where storage is available based on demand, and pricing is more closely aligned to actual demand. The OPEX model affects the overall project cost, and cloud storage ends up being cheaper, as well as ideally suited for secondary storage use cases such as for backup and disaster recovery.

• **Vendor and Infrastructure Consolidation.**

Enterprises can reduce vendor costs by eliminating multiple resource needs, including everything from on-site tape and secondary disk hardware, backup software, and their respective licenses, to ongoing maintenance costs. By using the public cloud rather than hosted or dedicated cloud storage solutions, enterprises eliminate the costs of cloud-specific gateway hardware.

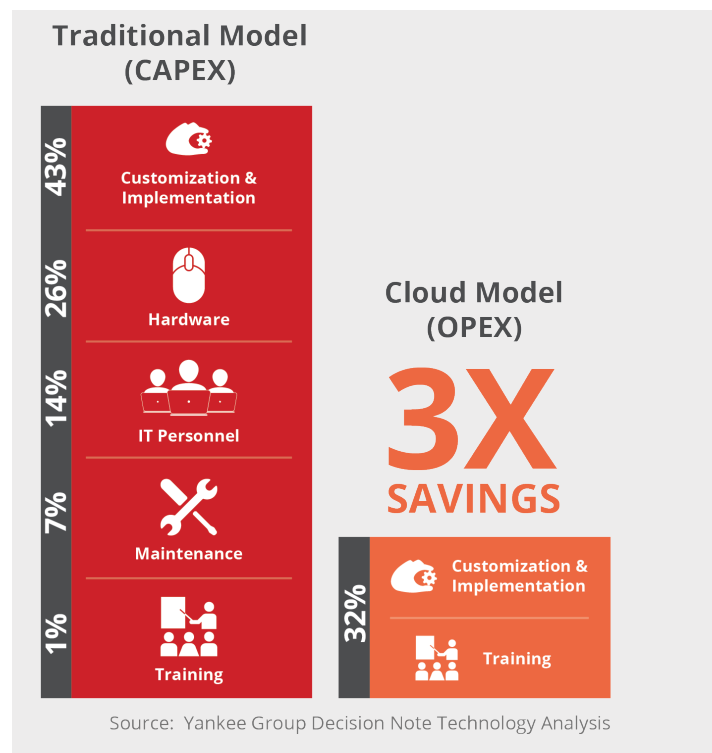
• **Reduced Administration.** On the resource side, enterprises can similarly and significantly reduce administrative overhead. They eliminate the need for the manual labor associated with managing tapes (loading, unloading, shipping off-site), not to mention the cumbersome task of confirming that backup tapes have recorded data properly. Cloud technology also offers a higher level of automation, which lessens the need for manual activities but also increases agility and reliability.

• **Global Availability and Consistency.** Moving to the cloud for secondary storage workloads



overcomes a lack of consistency in tools and processes, a residual effect in many organizations of mergers and acquisitions or decentralized IT management. Similarly, the

cloud enables enterprises to eliminate piecemeal deployment: They can launch the same quality of service simultaneously across the globe.



Once enterprises have deployed backup to the cloud, they derive other advantages. Centralized management eliminates the use of legacy silos and manual errors, and makes it easier for administrators to manage policies and restore data. At the same time, enterprises benefit from the higher levels of security and compliance derived from a cloud services provider that's focusing key resources on those issues so its clients don't have to.

• **Simplicity and Scalability.** The public cloud architecture can easily ingest and manage large volumes of data to scale as needed; it avoids the limits of on-site legacy hardware as well as the need to add new hardware. It's also a more efficient storage mechanism. Using object storage (as Amazon S3 employs in its Amazon Web Services cloud), the cloud stores only a single copy of data; other solutions, especially with inconsistent or non-standard approaches, can store up to 11 copies of data.

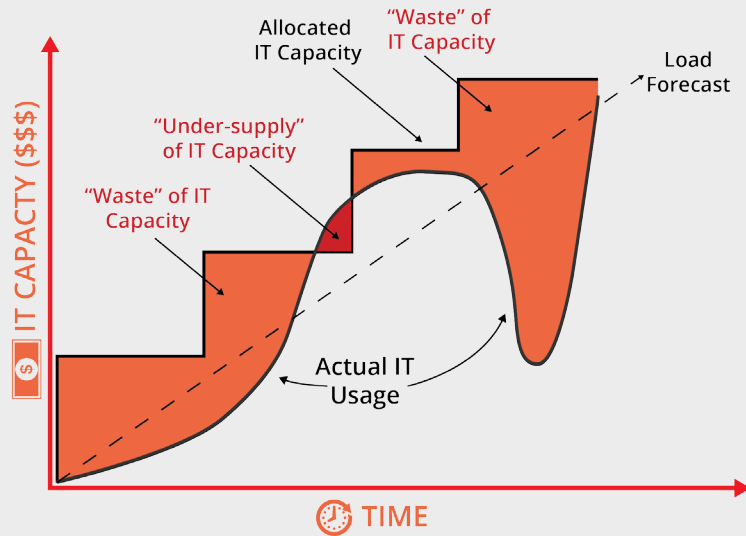
Much more than on-site systems, the public cloud offers global availability that is needed by companies processing massive volumes of data

and expanding growth. On-premises deployments constantly struggle with sizing storage, acquiring new storage, etc. For global enterprises, this is a boon. They can choose from multiple public cloud locations and get greater geographic reach. They can store their data using consistent internal requirements, but also while adhering to national and regional data regulations, such as the European Union’s stringent data-location requirements. Finally, they can store data close to where it’s used and needed to reduce latency and increase performance.

**The benefits of the capacity cost model.**

With a "pay as you play" pricing model, cloud storage offers ability to scale to peaks and ebbs in demand, ideally suited to backup and disaster recovery use cases.

Capital Intensive IT Model vs. Consumption of IT Model



**Thinking Benefits Beyond Just Backup**

The real benefit of using public cloud storage extends beyond backup, into utilizing cloud storage to solve other key areas for business efficiency and data protection.

• **Disaster Recovery.** Thankfully, the cloud has practically eliminated the old-time need of maintaining a duplicate hot site for business continuity — an expensive option that required constant monitoring between the two sites for consistency after OS upgrades and other changes. Public cloud storage brings the same kind of simplicity and efficiency to business continuity. Too many external threats are beyond IT’s control these days when it comes to potential data loss; even a location nowhere near a traditional floodplain can suddenly be rendered useless by a broken pipeline.

With data copied off-site, downtime (and its resulting impact on productivity) becomes even less likely. Cloud-based systems enable enterprises to restore data to different locations if workers are unable to access their traditional location. In fact, combining cloud server backup with DR, as **Gartner** analysts have noted, organizations can address two needs at once and get “disaster recovery as part of the deal.”

• **Test and Development.** With the right software, the public cloud can also be used as an extension of an enterprise data center, not only for backup and disaster recovery, but also as a platform-as-a-service (PaaS) option for test and development. By leveraging a copy of a virtual machine in the cloud, tests and validation can be run against a copy of the production data, making backups useful for more than just backup. Those applications still need to be confirmed and then downloaded again before they’re put (or put back) in production. The process isn’t markedly different, and as such, enterprises can derive an even bigger return on investment (ROI) by utilizing the cloud for multiple needs.



• **Data Analytics.** The cloud storage model promotes the thinking around connected data, for example, to leverage existing data for business good. Enterprises can not only store backup data but also analyze the information to understand core challenges around dormant data analysis, and storage growth and data classification.

## Leveraging Public Cloud with Cloud-First Approach

While the value of the public cloud for data protection and secondary storage is clear, the choices available to enterprises may not have as much clarity.



- **Cloud Gateways.** For instance, some legacy backup vendors require the use of a gateway appliance to link on-site systems with cloud systems. While this enables enterprises to govern the flow of data to the cloud, it also creates other problems. The gateway becomes a bottleneck, both from a networking and a reliability standpoint; if the gateway fails, all access to the cloud fails. In addition, such a hardware-based solution has no facility for ensuring deduplication of data in the cloud; it actually sends more data than is necessary for secondary storage.

- **Hosted Appliance.** Another common option is the hosted model, in which the cloud service provider essentially duplicates the on-premises architecture of the client. This essentially replicates that hot-site model, requiring a one-to-one match of equipment that can be both expensive to build and expensive to manage — because the provider is no longer managing aggregated storage on its own systems.

- **Cloud Native.** A third service option — built natively for existing public cloud service providers (e.g., Amazon) — provides a number of advantages for enterprises in search of greater efficiency and automation. It takes advantage of the native capabilities of public cloud such as storage, compute, etc. and their efficiencies to create a well-integrated offering from the ground up. And, because it does not require a translation layer between older deployments and a cloud-like a gateway appliance, it eliminates bottlenecks, boosting both performance and uptime.

This native cloud option also offers more versatility, allowing enterprises to use it as a convergence point for other important activities. It doesn't use cloud storage as a technological equivalent a warehouse; rather, the data can be re-used for many purposes — as noted — for disaster recovery and beyond.

## The Payoff

Relying on cloud-first software for secondary storage offers a variety of advantages: more efficiency for IT, more reliability for the business, and improved security for the data. But the most important aspect of converged cloud software is its total cost of ownership. By consolidating backup, disaster recovery, and PaaS opportunities in the cloud, enterprises can save anywhere from one-half to one-third of their costs.

**Relying on cloud-first software** for secondary storage offers a variety of advantages: more efficiency for IT, more reliability for the business, and improved security for the data.

Consider the cost elements: on-site servers and backup software; second servers for caching and restoration, used as the source drives for tape machines; pickup and delivery of tapes; and cost to store (and perhaps restore) tapes. Consider, too, the administration elements. Each of those steps requires maintenance, licensing, and oversight of multiple vendors; add in consulting costs in some cases.

Using a converged cloud-first option, the TCO is considerably lower. Why? Because it takes advantage both of the cloud's infrastructure efficiencies and of the cloud-first software's ability to eliminate the need for gateways, colocation, and other hardware obstacles. Backup may be secondary storage, but enterprise-grade storage is the new primary path to administrative efficiency.



## How Druva Helps

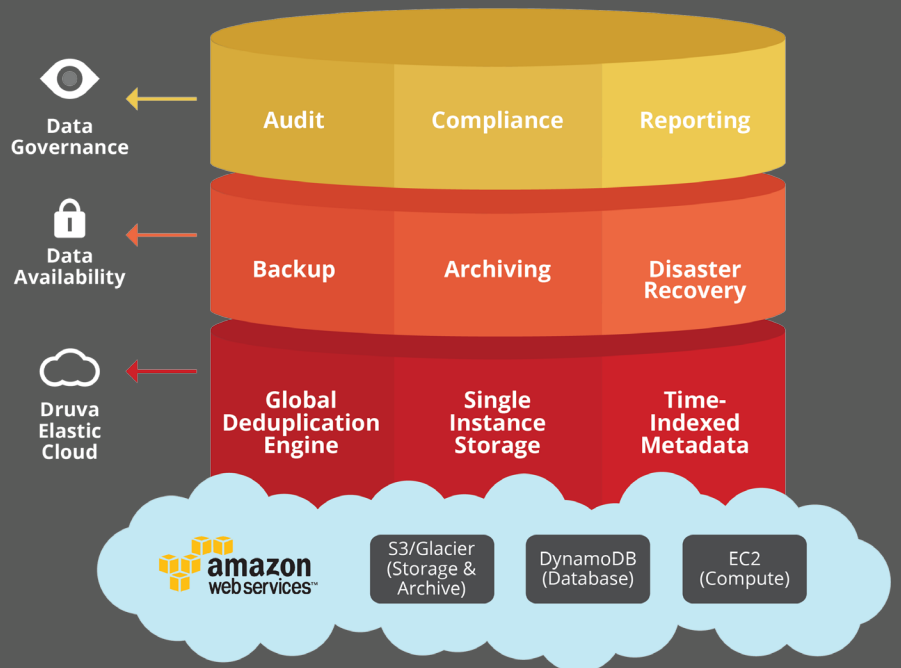
Druva is pioneering the increasingly popular area of cloud-first backup software. In addition to helping enterprises move beyond backup and be more efficient in disaster recovery, test and development, and even data analytics/compliance, it's also rewriting the rules around traditional backup infrastructure, which can have any number of limitations related to processing and storing data.

Druva's cloud architecture creates a single converged architecture that scales to support multiple workloads, all while natively leveraging Amazon Web Services storage technologies, such as S3 and Glacier. Its solution uses Amazon S3 storage for short-term hot data (30 days) and midterm warm data (90 days), applying intelligent metadata to track where data is stored. Archived data (more than 90 days) is stored on Amazon Glacier storage, for the least expensive and more reliable option.

The unified approach to backup and archiving provides hot-caching capabilities, delivering a LAN-speed local solution and removing the need for separately managed secondary storage. Unified hot, warm, and cold management means data moves between tiers automatically, as needed. Caching that utilizes global deduplication delivers high-speed multiple backup snapshots for fast recovery.

It also gives enterprises a higher level of comfort regarding security, thanks to split-key encryption and data scrambling (and enterprise data is inaccessible to Druva). Druva accommodates compliance issues as well, thanks to a model that addresses both regional data residency and privacy concerns. This enables even global companies to simplify their architecture by "storing local."

Druva has made significant strides in improving the architecture of data management in the cloud. It tackles the need for bandwidth management by globally deduplicating data on the client side. That way, its Phoenix software optimizes bandwidth for uploading data to the cloud. Both backups and restoration happen more quickly. In addition, it can accommodate retention needs based on your company and your industry's requirements.



Phoenix is a backup and disaster recovery storage solution even your CFO will love. Calculating the cost of hardware and backup software (along with associated maintenance and licenses), along with tape vaulting services, real-world examples show that Druva can provide enterprises backup — and more — for about half the cost of a backup-only solution.



For more information visit <http://www.druva.com/phoenix> or take the virtual product tour at <http://www.druva.com/phoenixtour>