

# Enterprise Data-as-a-Service

How Data Virtualization Enables Agility, Resiliency,  
and Hybrid Cloud Mobility

actifio

# Enterprise Data-as-a-Service

How Data Virtualization Enables Agility, Resiliency, and Hybrid Cloud Mobility

## EXECUTIVE SUMMARY

The world is changing fast. Businesses need new capabilities to keep up, new software to create competitive advantage.

To build it, they're moving to the cloud – private, public, or hybrid. These architectures accelerate business agility while reducing the cost of building, running, and protecting applications.

Actifio is helping enterprises get there faster. We **virtualize** data... free it from infrastructure, make it easier to access, and cheaper to protect. Data virtualization is becoming an accepted best practice among enterprise customers looking to establish a hybrid cloud architecture, making it:

- **Faster to begin...** freeing data from infrastructure, capturing it directly from applications, and letting you put it wherever you like over its full life cycle.
- **Faster to decide...** de-risking cloud vendor selection by making it easy to move data into, out of, and across private and third party clouds.
- **Faster to access...** letting you mount virtual copies instantly anywhere, rather than moving redundant copies from place to place each time you need one.
- **Faster to protect...** eliminating the need for siloed backup and DR point tools, as well as the complexity and cost that comes with them.
- **Faster to evolve...** providing a step-by-step migration path of use cases from vaulting, to analytics, to test data management for DevOps, backup, and DR.

Actifio is data virtualization for the enterprise. Actifio is **enterprise data-as-a-service**.

## Why Hybrid Cloud?

Today, the largest taxi company owns no cars, the biggest overnight lodging business has no hotel rooms, a top ten retailer has no stores, and the largest provider of on-demand entertainment needn't be delivered through your cable box. Each of these businesses—and many more—are being transformed by powerful new software enabling innovative new business models, and in the process rendering the infrastructure of established leaders obsolete. It is a time of opportunity, and a time of change.

What's not changing, though, is the IT budget in many of the same big businesses now threatened by disruptive new entrants. We hear it from top executives all the time:

<p>"I NEED TO MAKE MORE OF IT SOMEBODY ELSE'S PROBLEM."</p> <p><b>GLOBAL BANK VP, IT</b></p>	<p>"I NEED TO SAVE MORE ON WHAT I STILL NEED TO OWN."</p> <p><b>ENTERPRISE CIO</b></p>	<p>"I NEED TO FIND MORE LEVERAGE IN INNOVATION."</p> <p><b>TOP 5 MSP CEO</b></p>
--	--	--

The need for a digital enterprise to continuously develop new software, and to find new ways to leverage proprietary data is being inhibited by an IT model designed for an obsolete model of a physical enterprise. An unprecedented level of innovation across corporate IT departments worldwide, much of it centered on moving to the cloud, is underway to fix this.

Whether the planned cloud is public or private, by choosing a new Managed Service Provider (MSP) Partner or embracing a cloud service delivery model with internal resources, the vast majority of businesses (95% according to a recent RightScale survey of 1,000 IT professionals) are embracing a hybrid cloud model.

FIGURE 1: 95% of Respondents are Using Cloud

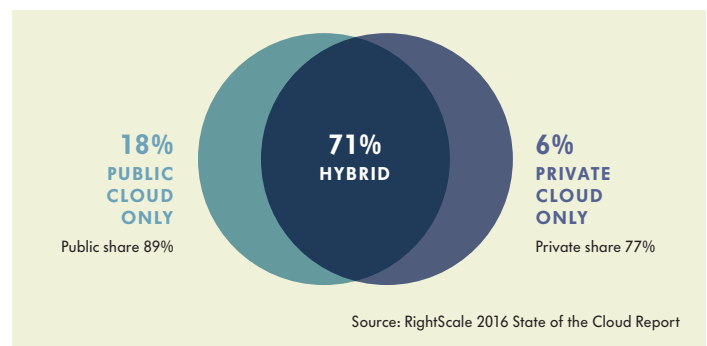
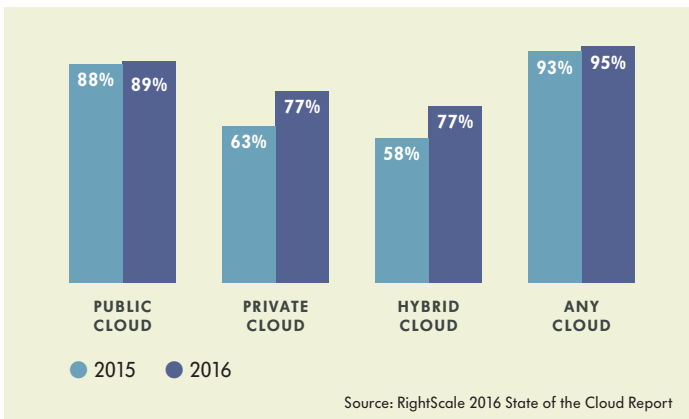


FIGURE 2: Universal Cloud Adoption: 2015 vs. 2016



## Hard Questions

While embracing the flexibility, scalability, and relative economy of cloud-based resources is pretty straightforward for a mobile app provider born in the cloud, it's much harder for an established enterprise IT organization. Hard questions abound:

- Which applications or use case should I start with?
- Who's the right cloud provider for me?
- Will I be locked into their platform?
- How do I make data available fast to those who should have it?
- What's the true cost of using a cloud provider over the long term?
- How do I protect my data from those who shouldn't access it?
- When will I get from here to there?
- Where should I even begin?

Answering these questions starts with still another: **Why?** More specifically... Why is answering these questions so darn difficult, for a technology that was supposed to make things easier? The answer is deceptively simple:

Changing our infrastructure is hard because we've architected everything up from our infrastructure.

Think about it. Start with your DevOps model. **How much application development cycle time is the result of someone in Dev waiting for someone in Ops to provision infrastructure or Data?** First, the Ops teams tackled the infrastructure provisioning problem by delivering Infrastructure-as-a-Service (IaaS) using private cloud deployments and are evaluating external public cloud services that are optimal for the business requirements.

Then there's the data. We manage, provision, and protect data much the same way we did before YouTube existed. With all that has changed since then, it takes a huge amount of Ops time, resources, and people to deliver the right data to the right dev/test teams at the right time.

Consider data protection... **How much of your current backup model is built around the boxes you're trying to protect?** How about DR and business continuity? If you're like most businesses, the folks who sold you your storage boxes went on to sell you a whole bunch of proprietary software tools that sit on top of them. Now a new generation of startups is trying to do the same thing, locking you into another proprietary box they built. It should come as no surprise that these vendors think outside the box: their proprietary software creates the need for more and more of their proprietary boxes, which create more and more copies of the data that require more and more of their proprietary software that uses more of their proprietary boxes. How do these boxes move with you when you have to go to the cloud?

As we said at the very start of this white paper, infrastructure doesn't matter anymore. Applications are what matter. Data itself is the new infrastructure, and the lifeblood of your business in 2017 and beyond. To truly deliver on the hybrid cloud IT model, hybrid cloud infrastructure needs to be combined with hybrid cloud data management.

## The Case for Data Virtualization & Data as a Service

What's required to fix all this—and to speed the path to the hybrid cloud—is a new data management model architected **down** from the applications you need, instead of **up** from the infrastructure you don't. It's time to virtualize your data; to decouple it from infrastructure just as you have decoupled compute from servers and network from switches, and then, to deliver that data on demand as a service to internal teams.

Do that, and the hard questions of the hybrid cloud get a whole lot easier. You can capture data directly from your applications, manage it wherever you like, and use it however and whenever you need to, no matter where it happens to be at the time: Your local data center, DR site, or MSP partner; in Amazon, Azure, Google, Oracle, IBM Softlayer (IBM Bluemix), GE Predix, etc. by enabling data to be delivered as a service.

Remember... If the first decision you're making in the move to hybrid cloud is the selection of one vendor over all the others, you're already setting yourself for the same problems down the road. **It's not about the infrastructure.**

Don't trade vendor lock-in on the box for vendor lock-in in the cloud.

Data virtualization eliminates this, and for that reason alone, it's fast become a generally accepted "best practice" for enterprise customers looking to move toward a public, private, or hybrid cloud model. Data virtualization does a whole lot more, though, transforming the responsiveness and efficiency of both business agility and business resiliency use cases by enabling data to be delivered as a service.

## Data Virtualization and Cloud Business Agility Challenges

How does data virtualization and DaaS contribute to application development, and business agility overall? Well, whether in the public or private cloud, fast access to high quality test data (usually database clones and subsets) is a critical requirement for Dev and Ops teams, and data virtualization can help.

## QUALITY

Dev and QA teams love the cloud for rapid code-testing and development environments, because they can provision the infrastructure on-demand, in the same self-serve manner they access other development tools and assets. This testing typically happens against low-fidelity test data, however, consisting either of synthetically generated sample data or a subset of the actual production data set, and thus lacking the richness and complexity of the production system to which it will eventually be deployed. As a result, Dev/Test teams often find defects late in their test cycles—in UAT, for example—impacting product quality or delaying the release.

**The solution to late-stage quality and timeline surprises is providing access to full, high-fidelity copies** of data throughout the development cycle, so more accurate testing can be done early and in parallel, i.e. “Shift Left!” Virtual databases representing high-fidelity, frequently refreshed full copies of production data deliver that solution, without the need to replicate the same physical infrastructure under each and every development team.

## SPEED

To overcome the quality and data fidelity problem, organizations often decide to clone data from production to test environments, or to restore it from backups. Creating physical copies for many test environments takes too much time, though, and moving large datasets from on-premises to the cloud or over any WAN can cause significant delays.

**The solution is to create a system that minimizes the physical movement of data**, to only the blocks that change incrementally. Virtual copies deliver this while providing instant access anywhere, enabling the simple provisioning data (via API or portal) in its native format for immediate use.

## COST

Creating all these extra physical copies requiring storage capacity, the bandwidth required to move data in and out of the cloud, and the amount of time DBA and operations admin teams spend creating clones or restoring copies all significantly increase the cost of making the right data available wherever and whenever it's needed.

**The solution is defining a single system of record for data** that eliminates excessively redundant physical data copies, establishing centralized control while automating access to the necessary data. Once again, the solution is data virtualization.

## CONTROL

Organizations often don't have an easy way to deliver the right data to the right users on the right machine in the right location. Masking sensitive data is essential to managing business risks and meeting government and compliance requirements, and can be a significant drain on time and resources. Make the process too long or time-consuming, though, and enterprising test/dev teams too often “go rogue” and hold on to physical copies of production data, unmasked, uncontrolled, and unknown to central Ops.

**The solution is automated data masking and provisioning with role-based access controls**, that improves the speed and ease with which the right teams can get the data they need (‘personal’ database copies to test against), while doing the right thing to keep Ops and InfoSec in control.

## Data Virtualization and Cloud Business Resiliency Challenges

IT leaders quickly discover the challenges of using cloud for data protection and DR—most often challenges related to their existing toolsets—and the difficulty of maintaining consistent enterprise SLAs across their new cloud footprint and existing data centers.

### LONG RTO / SLOW RESTORES

Existing enterprise backup products often have some cloud connection or support for cloud object storage, but they still have long recovery times, as data needs to be converted from proprietary, backup/archive formats into an application-native format. They claim instant recovery for a few VMs on-premises, but can't support instant recovery of systems at scale, particularly on AWS, Azure or Oracle Cloud. The result: Longer Recovery Time Objectives (RTO).

**The solution is data virtualization, which provides instant access to data that's been sent to the cloud**, in an application-native format, for scaled recovery operations. Since data virtualization preserves data in its native state, and doesn't need to move or create extra copies for restores, application recoveries get done faster.

### LONG RPO / INFREQUENT BACKUPS

Storage replication products that work fine on-premises often stop working across a hybrid cloud environment—usually because of storage incompatibility—and traditional backup products often don't support Recovery Point Objectives (RPOs) less than 12 to 24 hours.

**The solution is an incremental-forever data capture model** so more frequent backups can be done, shrinking backup windows while delivering lower RPOs. Data virtualization allows incremental changes to be merged into master copies, creating synthetic full copies for recoveries.

### RECOVERY FROM CORRUPTION / RANSOMWARE

Organizations that use native database replication tools unfortunately can't rewind to a previous point in time and recover. That means when corruptions or compromised files get replicated from the data center to the cloud, organizations can't recover clean data from either location.

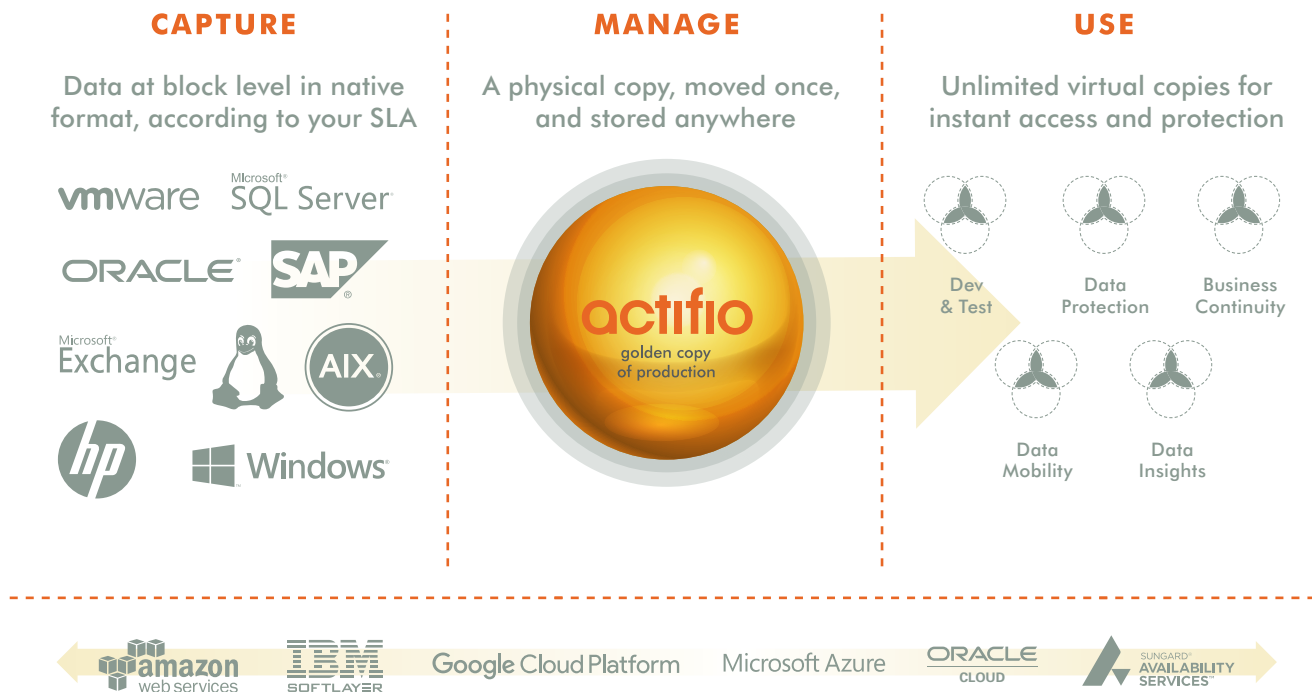
**The solution is easy rollback**, to a point in time when data was known to be good. This is a feature first available to consumers and coveted by the enterprise, and one baked into a data virtualization platform with synthetic full, instantly accessible copies.

### CLOUD TOOL LOCK-IN AND SPRAWL

Many new cloud-based tools are cloud- or function-specific, replacing lock-in to a box vendor with lock-in to a cloud vendor, and resulting in further tool sprawl.

**The solution is a data layer that's agnostic to the infrastructure underneath it**, whether on-premises or in the cloud, that creates and administers a single “golden master” copy of the data. Virtualized blocks of data and their associated metadata can then be used to satisfy multiple use cases and functional requirements: retention, recovery, resiliency, replication, etc.

FIGURE 3: Actifio Data as a Service



## Actifio: Data-as-a-Service for the Enterprise

Actifio data virtualization technology enables a Data-as-a-Service platform that lets enterprises to move faster toward cloud delivery models, satisfying what Gartner calls the critical “Bi-modal IT” requirements of building higher quality applications faster (Mode 2), while improving and reducing the cost of business resiliency and availability (Mode 1).

While traditional data management platforms are infrastructure centric and application agnostic, Actifio data as a service is the opposite: **application centric**, and **infrastructure agnostic**. It integrates with applications, hypervisors, and clouds via native APIs or lightweight connectors to:

- **Capture** application-consistent data at the block level, in native format directly from applications, according to SLAs set by the admin. These blocks—and the associated metadata indicating which blocks come from which applications, when—are used to assemble a “golden master” copy, updated incrementally forever, and serving as the system of record for application data throughout its entire lifecycle.
- **Manage** the lifecycle of that application data according to its SLA, which defines how often that data is to be refreshed with incremental captures (backups) from production, where and when it is to be moved, when it is to be masked/obfuscated/encrypted for security purposes, who can access it, and how long it is to be retained.
- Let you **Use** that data for multiple applications, for example to provision unlimited virtual copies for application development and testing, or analytics, to rapidly recover a 30 TB database, to restore the business

from discrete events like a site failure, application migration, or datacenter consolidation.

Actifio manages data for applications operating anywhere, preserves application data in its native format throughout its entire lifecycle, speeding up customers’ ability to leverage that data—on prem or in the cloud—on demand. Its SLA-driven model allows operational teams to more effectively manage policies across a hybrid infrastructure, even one including multiple clouds, data centers and remote offices. Wherever an Actifio solution is deployed—regardless the brand or form of its underlying storage—it can assure data is being managed according to a set of standardized (or fully customized) global policies. Actifio currently supports Amazon Web Services, Microsoft Azure, Oracle Cloud, Google Cloud Platform, IBM Softlayer (IBM Bluemix) bare metal, and dozens of MSP clouds based on VMware and Hyper-V.

Actifio’s ability to support almost any data use case from Test Data Management (TDM), Backup, Vaulting, Disaster Recovery (DR), Application/ Data Migration, Datacenter Consolidation or Analytics allows enterprise IT managers to avoid siloed toolsets that tie different IT groups to specific infrastructure, clouds, applications, storage systems or OS’s. The system delivers freedom, improving responsiveness while enhancing flexibility, breaking vendor lock-in, and reducing cost. That freedom empowers enterprise architects to move more quickly into the cloud, reducing the decision threshold of vendor and architectural choices through a data layer which is effectively decoupled from the constraints of underlying infrastructure, legacy, and proprietary toolsets.

## Actifio DaaS for Hybrid Cloud Business Resiliency

Enterprise customers have three important requirements to be satisfied for business resiliency across a hybrid cloud architecture:

1. Recovery Time Objective (RTO)
2. Recovery Point Objective (RPO)
3. Retention policy

Different application tiers have different RTO, RPO, and retention requirements, and enterprises often use multiple point tools to meet their needs, increasing operational complexity and total cost of ownership. Actifio, in contrast, provides a single data management architecture, delivering DaaS to support an almost infinite variety of RTO, RPO, and retention policy combinations, improving internal SLAs while reducing cost, complexity, and business risk—across whatever cloud-based or on-premises infrastructure that sits behind it.

Actifio delivers improved resiliency capabilities via software that runs either as a virtual appliance on AWS, Azure, Oracle Cloud, Softlayer (IBM Bluemix) bare metal, VMware, Hyper-V, popular enterprise clouds, or as a hardware appliance, on-premises. Actifio solutions provide:

1. Local/cloud backup with scalable instant recovery
2. Remote backup/vaulting for long term data retention (LTDR)
3. DR with scalable instant recovery

Let's take a look at how they work.

### LOCAL / CLOUD BACKUP WITH SCALABLE INSTANT RECOVERY

For applications either on-premises or in a cloud, Actifio lets you protect and retain data locally or in your primary production cloud—for a period of days, weeks, or months—to achieve rapid local data restoration. The system replaces traditional backup software, deduplication storage devices and storage array or cloud-based snapshots, as well as local backup tools and tape.

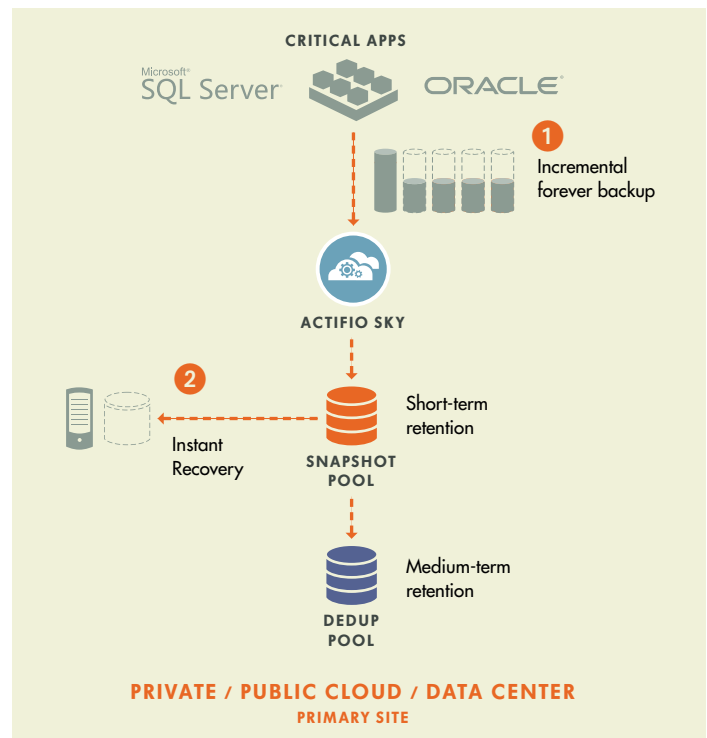
Deployed locally alongside the production application, Actifio performs incremental forever, application consistent backups, capturing data at the block level, using native application APIs such as VSS snapshots, Oracle RMAN APIs, or VMware VADP APIs. Each captured data block is stored as a rich, contextual object with metadata bindings pointing to the application this data belongs to, the time it was captured, and the SLA that defines its entire lifecycle.

For short-term data retention, these virtualized objects are maintained as a golden copy in the Actifio snapshot pool, an expandable pool of storage using any vendor or cloud independent disk storage.

By combining the baseline data and metadata of the initial full data capture with the changed blocks captured incrementally over time, Actifio delivers a single system of record for enterprise application data throughout its full lifecycle. Upon user/API request, Actifio uses the metadata to synthesize a full, point-in-time copy of the application data instantly, presenting it to a VM or physical server through a traditional storage interface.

This unique approach delivers not only instant access to application data at any time, but also a highly scalable model where provisioning/restore time is made independent of data size: a 50 TB database, for example, become accessible in the same time required to restore a 50 GB file.

FIGURE 4: Local Backup & Instant Recovery



For medium-term retention, application data is stored cost-efficiently in a deduped, compressed format in a “dedup pool,” another expandable storage pool using vendor independent storage. Depending on the SLA, application data is either moved automatically from the snapshot pool or directly from the application itself, into the dedup pool. Depending on the size of the data, change rate and the economics of storage costs, the dedup pool can be used to retain data from weeks to years.

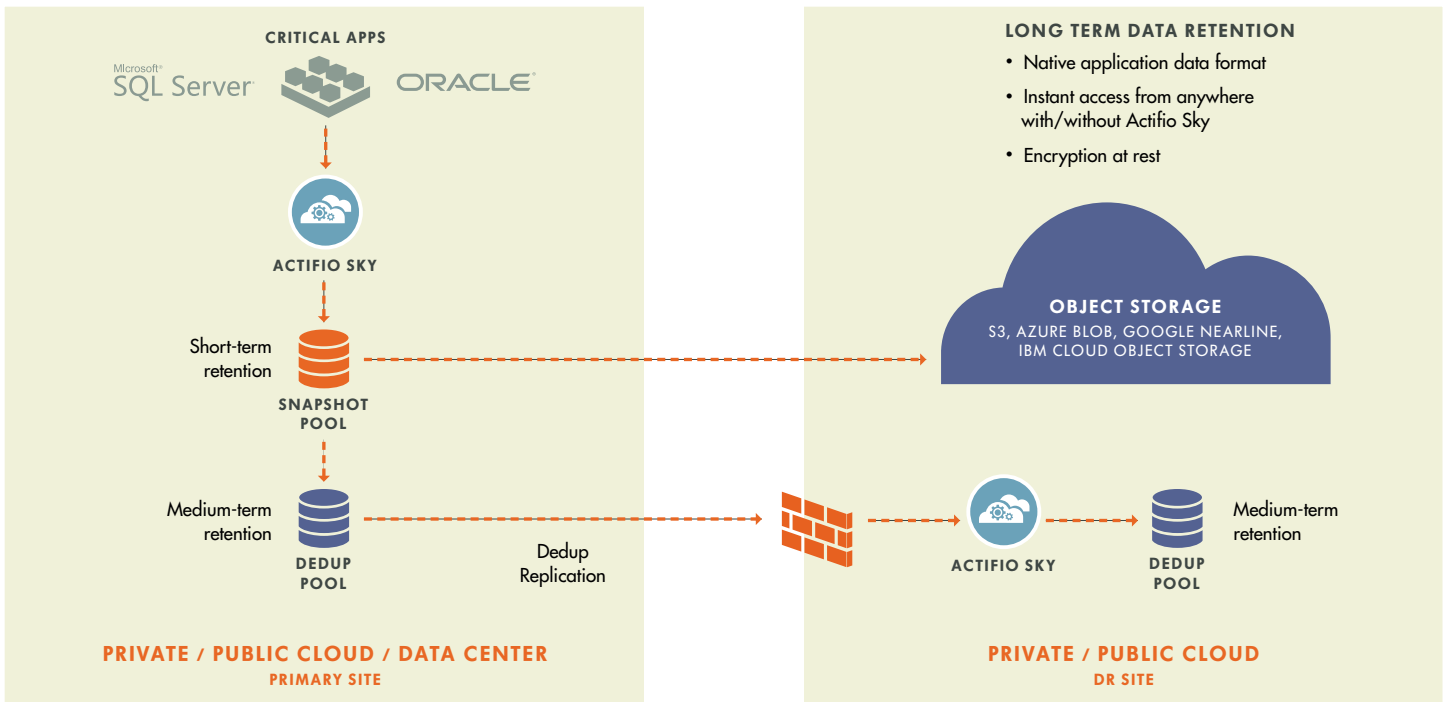
### REMOTE BACKUP / VAULTING FOR LTDR

If data needs to be taken off-site to protect against site failure, Actifio provides two different options, depending on retention time needed.

For storing data offsite for medium-term retention, Actifio's patented Dedup Backup Replication provides a highly efficient data replication technology sending only unique blocks of globally deduplicated and compressed data between the local and remote dedup pools. Data is then retained in the remote dedup pool for as long as defined by the SLA for the specific application. In the event access to data is required on the remote site, the deduped, compressed data in the remote dedup pool is rehydrated to a remote snapshot pool, and can be accessed directly by an application instantly.

For multi-year retention, Actifio uses cost-effective object storage through its patented OnVault replication, which sends data directly from the snapshot pool to an object storage bucket in the cloud. Another unique aspect of Actifio's object storage implementation is that applications can instantly and directly access read/write copies of vaulted data in cloud object storage via an Actifio Sky appliance—in the cloud, or on-prem—just as they would from a conventional storage device. And since the data is kept in its native format, it

FIGURE 5: How Actifio Delivers Long-Term Data Retention



is also accessible via a free, Linux-based utility, even if the customer chooses to stop using Actifio down the road. This makes Actifio a safer economic choice for customers with very long retention policies, as they're not locked in to any individual vendor for the potentially decades-long life of their vaulted data.

Actifio supports a broad variety of cloud object storage such as AWS S3, Microsoft Azure Blob, Google Cloud Nearline Storage, and IBM Cloud Object Storage and Bluemix.

**CUSTOMER SPOTLIGHT: A US COURTS SYSTEM**

This judicial branch of a state government began using Actifio for local backup and offsite disaster recovery of Oracle and SQL Server databases, and VMware VMs, as well as vaulting to a public cloud. It switched to the cloud to satisfy a compliance requirement to retain a copy of application data out-of-state. Its prior approach had substantial monthly operating costs from a traditional tape backup product and services, and redundant software licensing. A year later, the agency switched from its initial cloud vendor to Microsoft Azure, a change made fast and painless by Actifio.

**DR with Scalable Instant Recovery**

In the event of a site disaster, Actifio lets you leverage data in the remote site to quickly recover applications in the cloud, at DR sites, or in other clouds, with a wide range of RPO and RTO options to match the criticality of the application.

For faster RTOs, Actifio uses two other innovative replication technologies—**Dedup Async Replication (DAR)** and **StreamSnap**—each of which maintains a recent copy of data in its native format in the snapshot pool for instant mounting.

**DAR** begins with using the same sophisticated, bandwidth-optimized dedup replication technology used by backup replication to send data efficiently between the dedup pools. Once data lands on the remote dedup pool, DAR also automatically rehydrates the latest incremental image to the snapshot pool so that the application image is ready to be used instantly for low-RTO requirements.

For low-RTO and low-RPO requirements, StreamSnap is used. It sends data quickly and efficiently from the local snapshot pool to the target snapshot pool by replicating incremental changed blocks, compressed and encrypted before transport. StreamSnap delivers low RPO by taking fast snapshots of application data, and low RTO by efficiently and rapidly replicating to remote site for immediate use by an application.

With either replication technology, administrators are able to recover their critical applications for disaster recovery in minutes or hours, and can even run recovered applications on the Actifio appliance until they're ready to migrate services to production storage on the DR cloud, or fail back to the primary cloud or data center.

FIGURE 6: Bandwidth-Efficient Replication with Instant Recovery

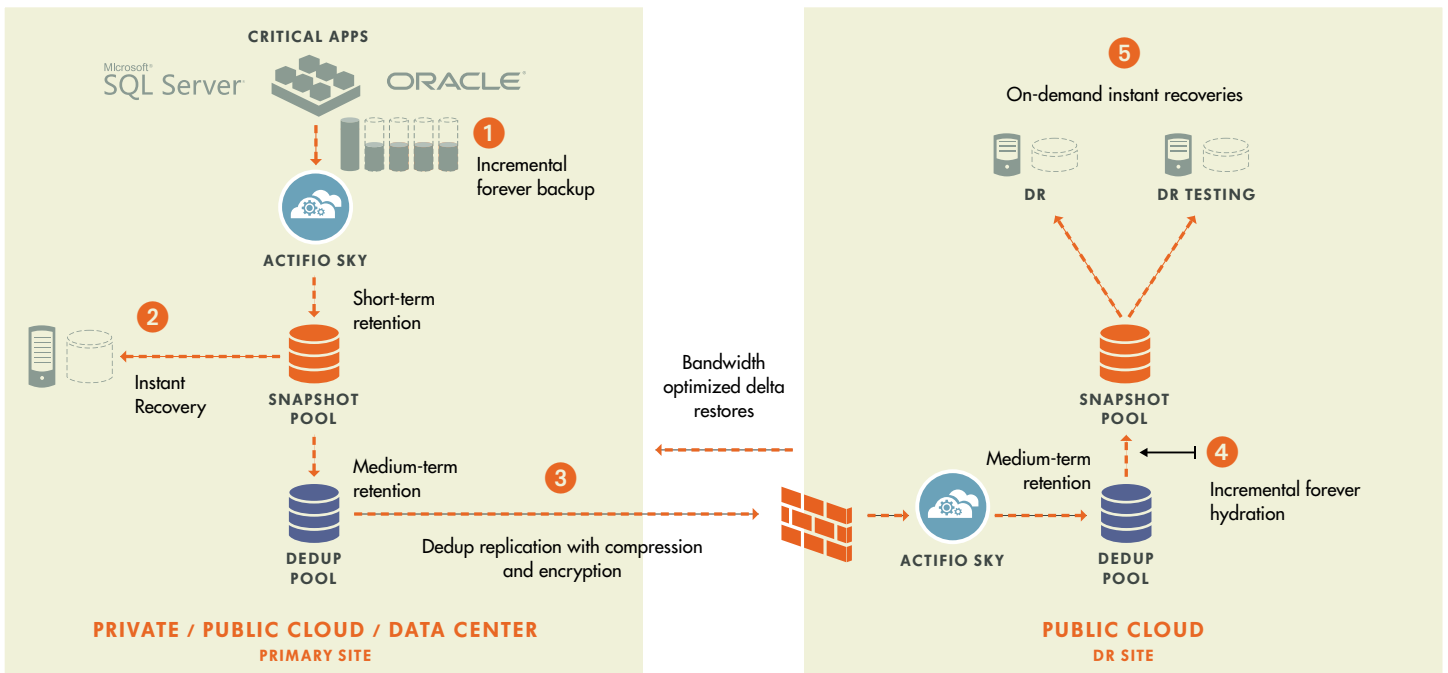
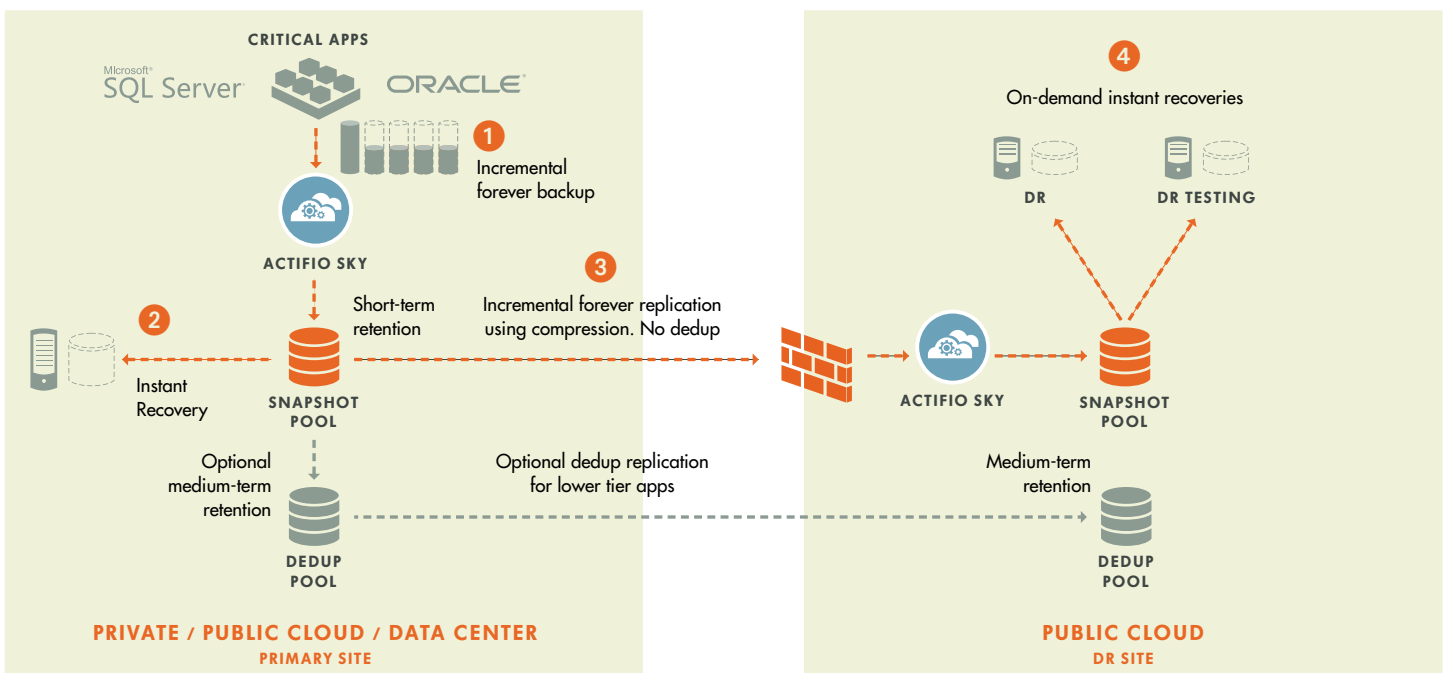


FIGURE 7: Scalable Instant Recovery with Low RPO





**CUSTOMER SPOTLIGHT: PC/NAMETAG**

pc/nametag, a global distributor of meeting and registration supplies, was making the transition from catalog to web-based sales. As part of this transition, the IT team needed to modernize their DR strategy, as well as decrease the cost and complexity of their DR processes. Their solution was to deploy Actifio Sky virtual appliances locally and to Amazon Web Services (AWS), where they could spin up DR compute resources on-demand, eliminating prior DR site colocation services, hardware and software expenses that were essentially a wasted insurance policy. Actifio enabled them to achieve recovery nearly 100 times faster than their legacy model, while simplifying operations and cutting costs.

In addition to speed, performance, and efficiency for business resilience, Actifio integrates powerful automation and orchestration tools to ease the challenges of enterprise data management. To tackle the complex task of orchestrating recoveries of large-scale VMware based environments and clouds, for example, Actifio created Resiliency Director™. Resiliency Director orchestrates compute, network, and data at the disaster recovery target site or cloud, and validates the results to provide a complete, global, enterprise-class solution for non-disruptive, automated recovery and test. Today, Actifio Resiliency Director is used extensively by many of the world's largest of Global Managed Services Providers and enterprises. Soon it will be extended to support additional cloud platforms, including the VMware Cloud on AWS to provide a one-click recovery and unmanned DR testing solution using cloud-based VMware resources.

**Actifio for Hybrid Cloud Business Agility**

For Dev, QA, and DevOps teams, Actifio Data as a Service can help simplify and accelerate:

1. The use of cloud resources for application development and testing,
2. Access to the right data by the right people in the cloud, and
3. Overall release cycles.

To do this in the cloud, customers deploy an Actifio Sky virtual appliance on the preferred cloud compute platform (AWS EC2, Azure, Oracle Cloud) from the cloud marketplace and provision the necessary storage for use by the appliance. An Actifio appliance is also deployed next to production systems (on-prem. or in the cloud) to capture application data using APIs such as VSS for Microsoft SQL Server and RMAN for Oracle Databases.

Using the data mobility features detailed above, Ops replicates an initial application-consistent copy of production databases, followed by incremental changes, to the cloud-based Actifio Sky instance, where the data will live in its snapshot pool. If data needs to be secured prior to provisioning in the cloud/remote site, a workflow can be set up to perform masking, ETL or other scripts operations to the data locally, or in the cloud

To serve up test data for DevOps teams, Actifio Sky creates virtual, application-consistent copies of databases, and instantly mounts them to the dev/test host resources, on-schedule or on-demand, so that appropriate users defined by Roles-based Access Controls have access to (only) the correct data, from the right (likely very recent) point in time. These virtual copies are read/writable and can be re-protected easily for bookmarking and roll backs, so destructive testing can be done with minimal lost time for recovering a prior environment.

FIGURE 8: Data as a Service: Test Data Management for Test/Dev in Cloud

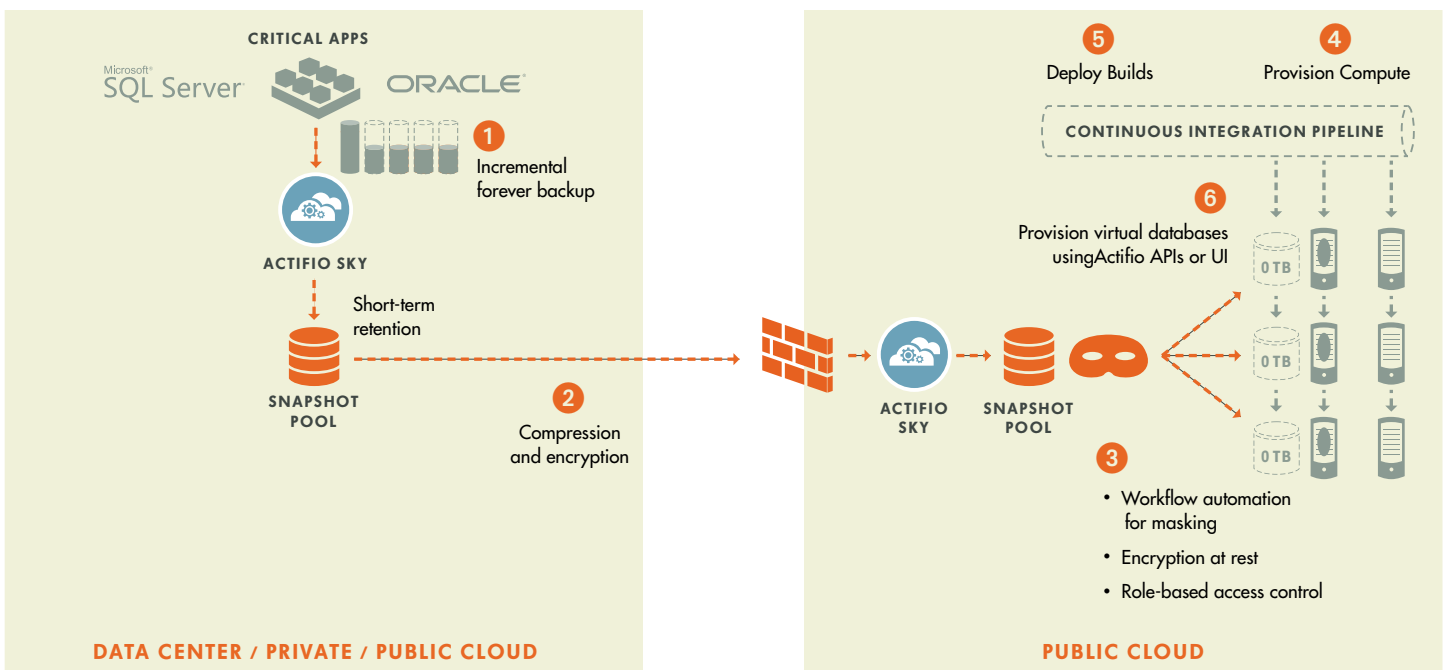
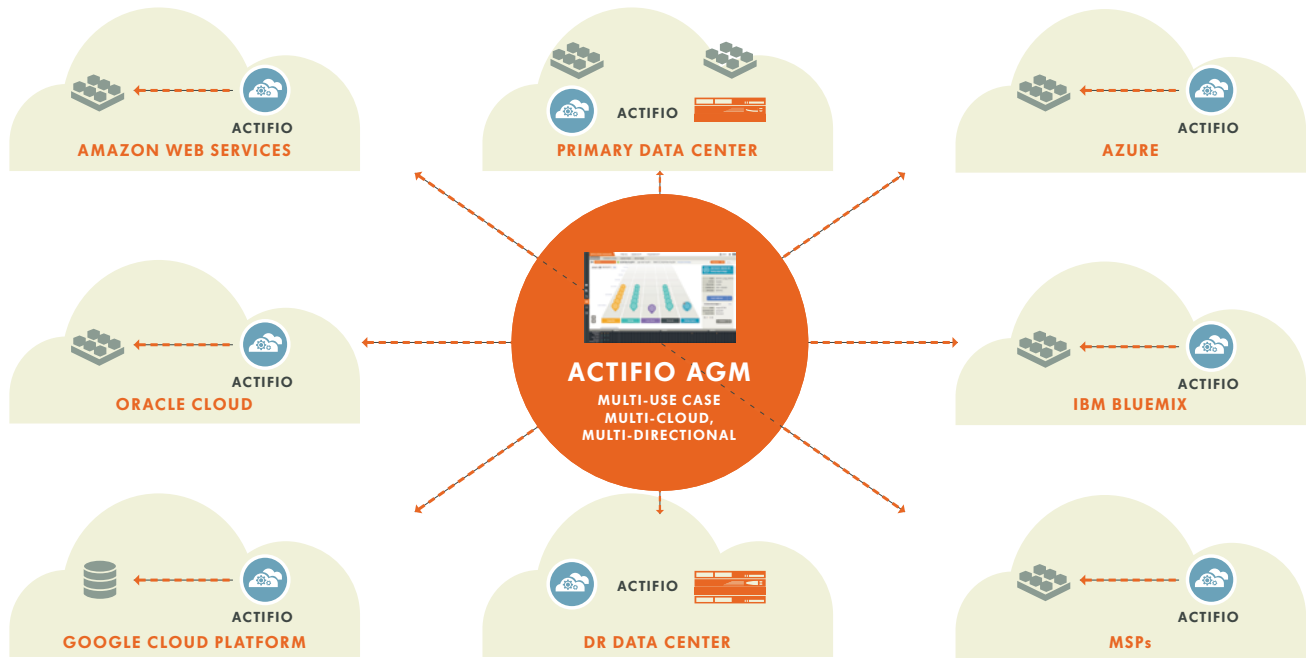


FIGURE 9: Web-Scale Data Cloud Platform and Solutions



**CUSTOMER SPOTLIGHT: TOP 20 GLOBAL BANK**

A global consumer & investment bank was preparing its second generation internal private cloud, focused on accelerating Agile development processes to deliver more features faster for business units. Two major challenges stood in their way: The need for faster provisioning of Oracle databases for DevOps teams, and the very basic problem of long restore times on large databases that violated internal SLAs. They knew their existing data management tools couldn't handle the task of protecting production databases, nor of provisioning those same databases for DevOps.

Their solution was deploying the Actifio Data as a Service platform to protect and provision/recover databases instantly, automatically from a Database-as-a-Service Cloud. The business results have already been enormous, achieving over \$1M in hard cost savings in the first year, mostly from reducing the storage needed to support all of their dev/test copies, as well as over 90% faster database provisioning time, and the ability to recover a multi-terabyte Oracle database in minutes.

The speed of the Actifio Test Data Management provisioning process is striking and often transformational for DevOps teams. Users commonly shrink a 5-day database provisioning process down to under 30 minutes. The productivity improvement across hundreds or even thousands of developers, testers, and DBAs resulting from this rapid "data-as-a-service" provisioning capability is massive, as is the enhancement of productivity and scalability

in Ops teams. Because they're able to test more, earlier and in parallel with real production data, several of Actifio's largest users have been able to add weeks of extra development time to their year... delivering more projects, faster, and at higher quality. A big "Shift Left" success.

The RESTful API in Actifio allows for easy integration with DevOps toolchains or Continuous Integration / Continuous Deployment pipelines. Automation tools such as Chef, Puppet, Ansible, Saltstack, Jenkins or others can instantly provision virtual databases on demand, wherever and whenever needed (visit [actifio.com](http://actifio.com) for a [demonstration video](#)<sup>1</sup>, and see how 10 virtual copies of a 1TB database can be provisioned in <5 minutes using Ansible via Actifio API integration.)

**Supported Public Clouds and Common Use Cases**

CLOUD	USE CASES	ONVAULT OBJECT STORAGE TARGET
AWS	TDM for Test/Dev, Backup, Vaulting, Recovery from Disaster, Ransomware or Corruption	S3, S3 IAS
Azure	TDM for Test/Dev, Backup, Vaulting, Recovery from Disaster, Ransomware or Corruption	Azure Blob
Oracle Cloud	TDM for Test/Dev, Backup, Vaulting, Recovery from Disaster, Ransomware or Corruption	N/A
Google Cloud Platform	Object Storage Vaults	Google Cloud Storage Nearline
Softlayer (IBM Bluemix)	TDM for Test/Dev, Backup, Vaulting, Recovery from Disaster, Ransomware or Corruption	IBM Cloud Object Storage (Cleversafe)

<sup>1</sup> <https://www.actifio.com/resources/demo-virtual-db-as-a-service-actifio-ansible-integration/>

## Summary

For enterprises, the desire to embrace a hybrid cloud architecture is just the first step in a journey that requires a fundamental re-thinking of IT strategy, a change that is re-shaping the IT business in the process. Applications are too often constrained by inflexible on-premises infrastructure, hampered by archaic data management models, and locked in to complex, confusing, and highly variable cloud vendor relationships.

The applications that power your business and the data that is its lifeblood must be at the core of this strategy... not the legacy infrastructure whose value has already begun its inevitable decline. A data virtualization layer—Actifio—is a critical tool in helping enterprise IT groups decouple their data and applications from their legacy infrastructure and to deliver Data-as-a-Service to improve the Agility and Resiliency of their business in a hybrid cloud model. With Actifio you can pursue a safe, pragmatic, use case-driven cloud adoption model, leveraging the right resources for your business right now.

For enterprises, the cloud is hybrid. Actifio enables an enterprise to truly leverage the economics, agility and specialized services offered by a multi-directional, multi-cloud, hybrid cloud architecture.

Want to see how it can work for you, risk-free? Drop us a line, and we'll discuss your needs and show you what we can do: [info@actifio.com](mailto:info@actifio.com) or **855.886.8997**. For additional information, visit Actifio at [www.actifio.com](http://www.actifio.com), or check out these resources to learn more:

- [451 Research Paper: Actifio Global Manager centralizes data management on-premises and in the cloud<sup>1</sup>](#)
- [Actifio OnVault Data Sheet<sup>2</sup>](#)
- [Actifio for AWS: Modernize Your Backup & DR in AWS<sup>3</sup>](#)
- [Database Virtualization for Test Data Management in AWS<sup>4</sup>](#)
- [Actifio—Azure Solution Brief<sup>5</sup>](#)
- [Demonstration: Oracle Application-Aware Mounts with Actifio Sky in AWS<sup>6</sup>](#)
- [Federal Agencies and the Actifio Cloud Advantage<sup>7</sup>](#)

- 1 [http://cdn2.hubspot.net/hubfs/214442/451\\_Reprint\\_Actifio\\_R1.pdf](http://cdn2.hubspot.net/hubfs/214442/451_Reprint_Actifio_R1.pdf)
- 2 [http://cdn2.hubspot.net/hubfs/214442/ACT\\_16-105\\_OnVault\\_Data\\_Sheet\\_E.pdf](http://cdn2.hubspot.net/hubfs/214442/ACT_16-105_OnVault_Data_Sheet_E.pdf)
- 3 [http://cdn2.hubspot.net/hubfs/214442/ACT-15-152\\_AWS\\_Backup\\_Modernization\\_WP\\_H.pdf](http://cdn2.hubspot.net/hubfs/214442/ACT-15-152_AWS_Backup_Modernization_WP_H.pdf)
- 4 [http://cdn2.hubspot.net/hubfs/214442/ACT-15-151\\_AWS\\_Test\\_Data\\_Whitepaper-G.pdf](http://cdn2.hubspot.net/hubfs/214442/ACT-15-151_AWS_Test_Data_Whitepaper-G.pdf)
- 5 <http://cdn2.hubspot.net/hubfs/214442/Actifio%20Sky%20in%20Azure.pdf>
- 6 <https://fast.wistia.net/embed/iframe/1fpj26itkx?videoFoam=true>
- 7 [http://cdn2.hubspot.net/hubfs/214442/ACT\\_15\\_185\\_Federal\\_Cloud\\_Brief\\_C.pdf](http://cdn2.hubspot.net/hubfs/214442/ACT_15_185_Federal_Cloud_Brief_C.pdf)

## About Actifio

Actifio virtualizes the data that's the lifeblood of businesses in more than 35 countries around the world. Its Virtual Data Pipeline™ technology enables businesses to manage, access, and protect their data faster, more efficiently, and more simply by decoupling data from physical storage, much the same way a hypervisor decouples compute from physical servers. For enterprise-class backup modernization, self-serve instant data access, or service provider business transformation, Actifio is the first and only enterprise class copy data virtualization platform.

For additional information, visit Actifio at [www.actifio.com](http://www.actifio.com) or contact Actifio at [info@actifio.com](mailto:info@actifio.com) or **855.886.8997**.

**actifio**<sup>®</sup>  
Radically Simple



©Actifio, Inc. All rights reserved. Actifio™ is a registered trademark of the Actifio Corporation. All other trademarks and service marks are property of their respective owners.