

Technical Review

Microsoft SQL Server Recovery and Performance with Actifio on Google Cloud

Date: August 2020 **Author:** Tony Palmer, Senior Validation Analyst

Abstract

ESG tested backup, recovery, and performance of a disaster recovery (DR) copy of a 1TB Microsoft SQL Server 2017 database leveraging both high-performance disks and cloud object storage using Actifio on Google Cloud Platform running production transactions. We measured recovery times and calculated cost savings.

Whether you are a CIO, VP of IT, or a database administrator (DBA), you are looking at ways to reduce costs while ensuring the performance, resilience, and integrity of your high-performance databases like SQL Server, SAP HANA, and Oracle.

As organizations leverage the cloud for backup, disaster recovery (DR), and analytics. Actifio's technology allows IT to use object storage in the cloud and recover directly from it, reducing costs while maintaining SSD-class performance on par with production systems.

The Challenges

Databases are central to driving the digital economy. IT organizations, including DBAs, struggle to balance performance, costs, application agility, test/dev, availability, and fast recovery. These use cases require database copies. IT must provide those copies quickly, easily, and cost-effectively. At the same time, IT is working to implement self-service database provisioning, cloud integration, and instant database recovery solutions.

Cloud today is an integral part of most organizations' strategy to increase agility while reducing costs. Whether protecting on-premises data by replicating to a public cloud, or replicating data on one cloud region to another, managing SLAs must be balanced against controlling costs.

Public clouds offer costly, high-performance, SSD-class storage and economical cloud object storage. A compelling solution would offer the cost effectiveness of object storage while delivering the performance of SSD-class storage.

ESG research reveals that cloud-based data protection has become mainstream, with 61% of organizations reporting the use of cloud as a backup target in 2019, up from 25% just three years prior. When asked about the primary benefit they realized as a result of this strategy, 31% reported reduced costs and 21% improved SLAs/recoverability, making them the most cited responses.¹ In a separate research study, ESG found that only 37% of organizations send secondary copies to public cloud services for non-protection purposes (development/test, analytics, etc.). Among the top barriers or challenges to re-using secondary data for business (non-protection) purposes reported by organizations are cost (38%) and deployment complexity (25%).²

SSD-CLASS PERFORMANCE AT THE COST OF CLOUD OBJECT STORAGE

ESG validated that organizations leveraging Actifio for MS SQL Server cloud backups and recovery can reduce the cost of backup and DR storage infrastructure by

 **88%**

with RTO and RPO measured in minutes, while enjoying the same high levels of performance as they get with their production instances.

¹ Source: ESG Master Survey Results, [Data Protection Cloud Strategies](#), June 2019.

² Source: ESG Master Survey Results, [The Evolution from Data Backup to Data Intelligence](#), January 2020.

Organizations need a solution that provides cost-effective, high-performance backup and DR and simple self-service data delivery for developers, data scientists, and application owners to achieve faster development with reduced risk.

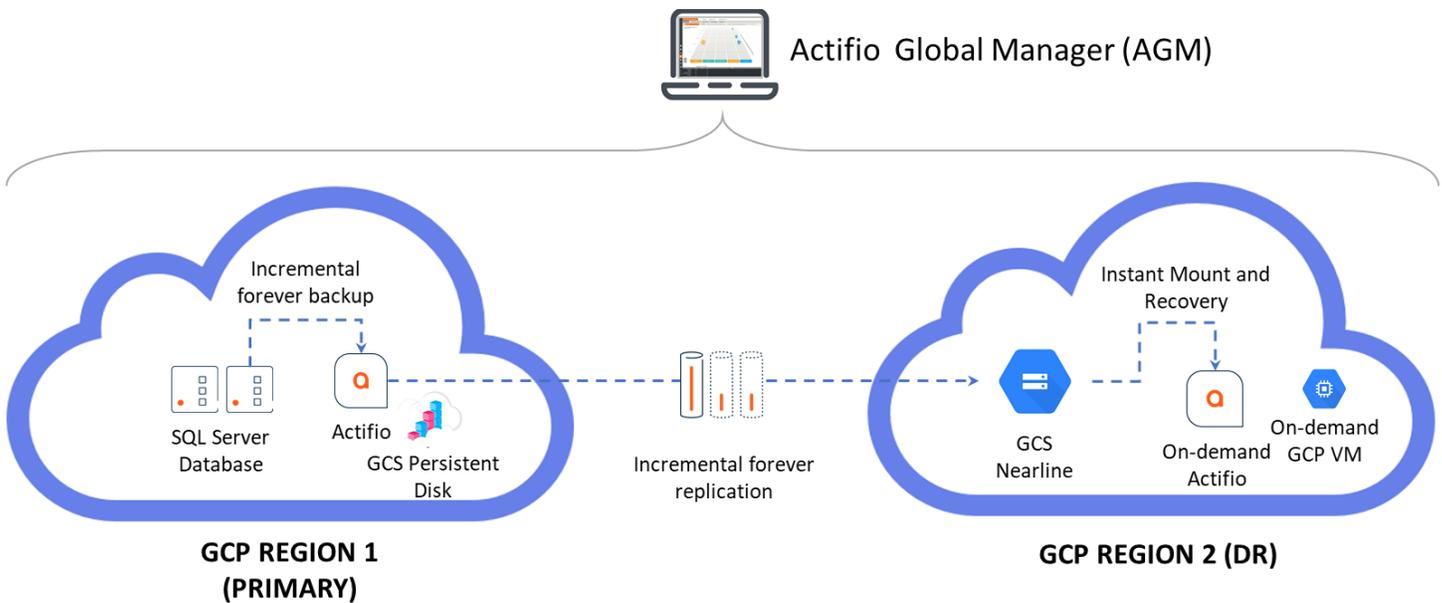
About Actifio

Actifio pioneered multi-cloud copy data management software, designed to enable organizations to virtualize and deliver their data instantly, anywhere. A software platform engineered for the enterprise, powered by their Virtual Data Pipeline technology, Actifio helps accelerate adoption of hybrid, public, and multi-cloud strategies; build higher quality applications faster; and improve business resiliency and availability.

ESG Testing and Validation

ESG performed evaluation and testing of Actifio software with SQL Server on Google Cloud Platform (GCP). This involved backing up a 965GB SQL Server 2017 database to SSD-class Persistent Disk in one GCP region (US West) and then replicating across a WAN to another region (US East) where it was stored on GCS Nearline storage. During the testing, we looked at backup and restore performance, time savings, and compared transaction performance at the primary site on Persistent Disk storage with transaction performance at the DR site on Nearline storage. The test bed is shown in Figure 1.

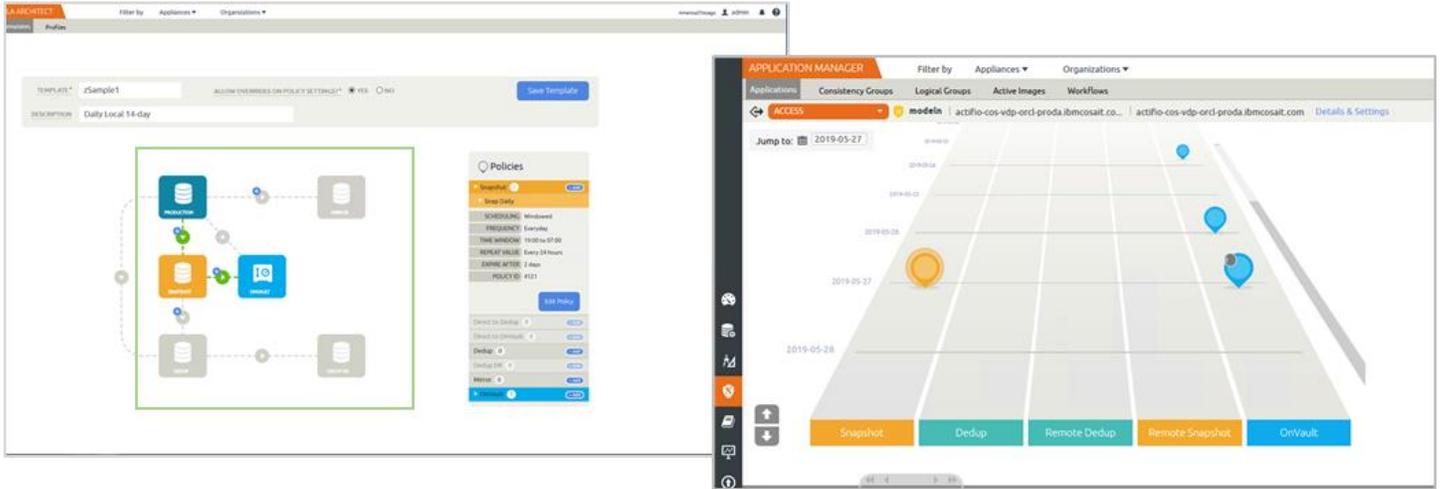
Figure 1. The ESG Test Bed



Source: Enterprise Strategy Group

First, we measured the time to complete the first backup to Persistent Disk in the local region. We created a policy for backup and DR replication using Actifio Global Manager (AGM).

Figure 2. Policy-driven Copy Data Management with AGM



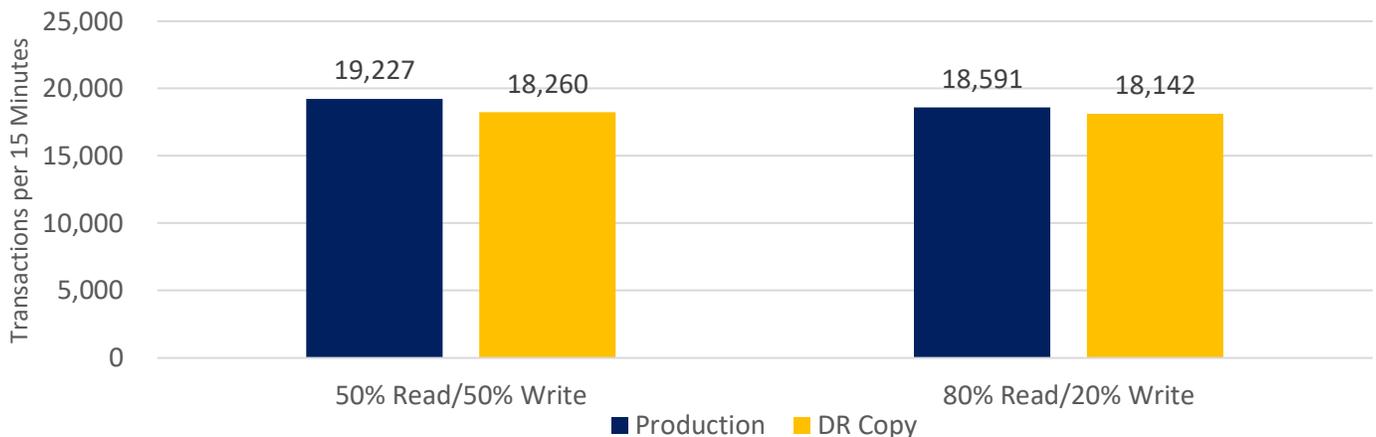
Source: Enterprise Strategy Group

In just 90 minutes, Actifio completed the first full backup and replicated to Nearline storage in the remote GCP region. From this point on, Actifio backed up and replicated in an incremental forever manner with application consistency. This shrinks the backup window and enables a tighter recovery point objective (RPO), especially important for large SQL Server databases.

To test DR, we spun up an on-demand Actifio appliance in the DR region. Next, we used AGM to bring up an on-demand SQL Server host and triggered Actifio’s app-aware mount technology to mount the backup image from GCS Nearline directly to the SQL Server host through the Sky appliance without having to rehydrate the backup to GCP Persistent Disk storage. At this point, Actifio brought the SQL Server database online in region two. It’s important to note that this is the same process for performing a migration or bringing up a test/dev instance. The recovery for the 965GB database took just five and a half minutes from start to finish.

Finally, the performance of the DR database was measured with the production version as the benchmark. This was a hard look at how well the DR database performed while running on lower cost Nearline Disk versus the production database running on SSD Persistent Disk block storage with the same server hardware. We ran various combinations of read-write workloads against the production database in region one and the DR copy running in region two using Load Runner to assess Actifio’s accelerated caching and optimizations for object storage.

Figure 3. Performance Comparison—Production versus DR



Source: Enterprise Strategy Group

It's important to note that SQL Server cannot use object storage directly; it requires block storage, which was being emulated by Actifio as a virtual block device. All SQL Server reads are read by Actifio from GCS Nearline storage and cached in Persistent Disk storage, ensuring future reads are served by the cache. All SQL Server writes are stored in GCS Persistent Disk. Actifio uses caching in Persistent Disk to deliver performance on par with production systems, as can be seen in the graph above. This is significant, considering the latency and performance sensitivity of databases and the apps that depend on them.

In addition to instant mount and recovery directly from cloud object storage for DR, Actifio has a "VMware Storage VMotion"-like migration feature in the cloud called "Mount & Migrate." Once turned on post-mount and -recovery, this feature—while users are accessing the database—starts copying blocks from the mounted virtual volume to the desired storage tier. Once the copy is done, the volume can be switched to the new storage tier for the SQL Server database with minimal downtime, less than a minute in the testing we observed.

Three-year Infrastructure Cost Analysis

ESG modeled and compared the costs that could be expected when deploying a disaster recovery plan with Actifio with a traditional method where backups are replicated to a DR location, reconstituted, then presented to an online server instance.

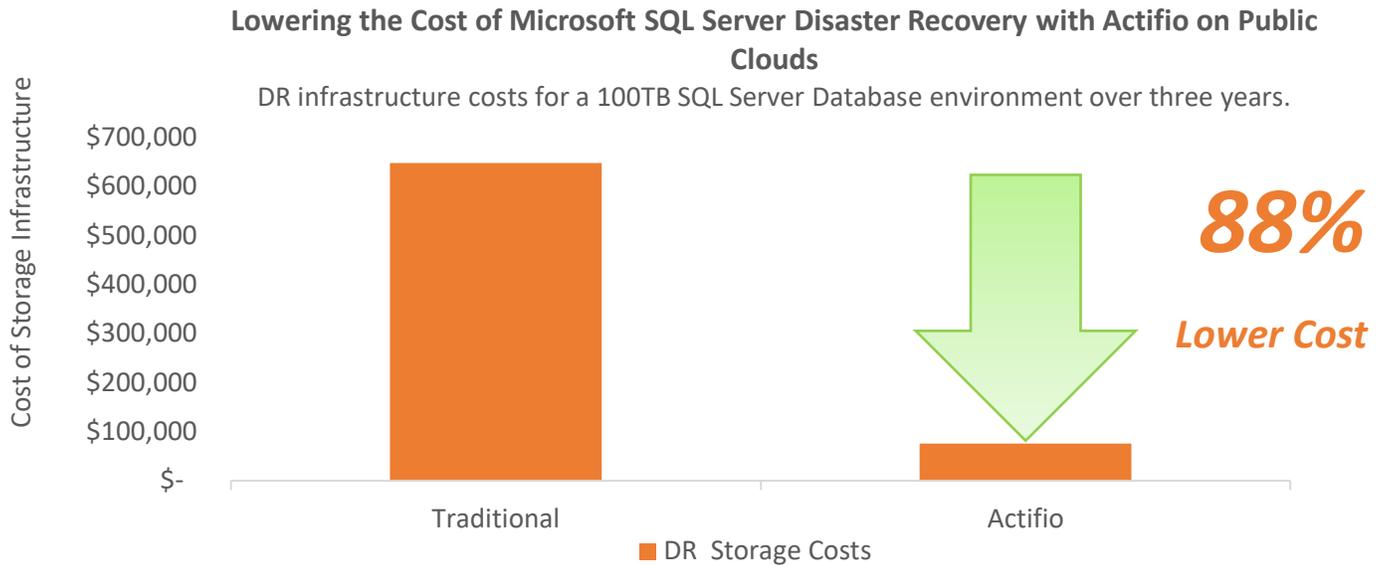
Direct SQL Server replication is a hard sell for many organizations due to the extremely high ongoing costs, so we did not include it in our three-year model. Direct replication requires an always-on second SQL Server instance with appropriate licensing, the same specs as the source instance, and Persistent Disk. Each of these comes with a significant monthly cost. By ESG's estimation, the DR VM instance and the Regional Persistent Disk storage alone would cost more than \$3,600 per month.³ This does not include the SQL Server license cost.

ESG modeled and compared the VM and storage-related costs that could be expected when deploying DR with a traditional backup/replication methodology compared with using Actifio. ESG modeled the expected costs for a company that needed to support a 1TB SQL Server production environment. Calculations used a simplified model based on costs that would be incurred over a three-year period without taking into consideration capacity and performance growth requirements or soft costs like administration.

The model looked at the hard costs associated with producing a DR copy of the SQL Server database using Actifio. This was compared to the more traditional model of running backups and replicating full copies of the database for each DR copy that was needed. We did not include the costs for the production database. In the analysis shown in Figure 4, we assumed a fully burdened storage cost of \$170 per TB per month for the traditional model using Persistent Disk and DR copies on Nearline at a cost of \$10 per TB per month. With Actifio, the DR copies are served directly from Nearline with a small Persistent Disk cache (20% of total data set size) that is only accessed when testing or executing DR.

³ All pricing was estimated using the [Google Cloud Pricing Calculator](#).

Figure 4. DR Cost Analysis—Actifio versus Replication to High Performance Storage



Source: Enterprise Strategy Group

The cost of storage for Actifio over three years is \$76,080 compared to \$648,000 in a traditional model. This represents a cost savings of \$571,920, an 88% cost reduction. For an organization with dozens or hundreds of production databases, this represents significant annual savings and a simplified model of DR testing and execution.

i Why This Matters

Restarting a SQL Server environment after a failure is a daunting and expensive task. It can require days of detailed work to put each element back in working order, and a backup production server instance and storage environment must be maintained. The downtime required is costly to the business in terms of production operations, and performance issues when failed over to the DR environment can damage reputations.

ESG validated that Actifio was simple to deploy and use. Creating a tailored recovery plan took only a few clicks, and the ability to use Nearline storage for copies that can be used and reused for a variety of cases in addition to DR, saves significant cost and effort. Actifio combined with Google Cloud Storage (GCS) is a powerful combination of high performance at a low cost point. Our testing shows that Actifio provides the performance needed to run DR copies of databases with performance equivalent to SSD-class storage in the production environment, while using much less costly Nearline Disk.

ESG calculated a three-year infrastructure cost model for creating a DR copy of a 1TB SQL Server database and found that Actifio with GCS is 88% lower than a traditional backup/replicate/reconstitute model. This significant savings is amplified for an organization with many databases to protect.

We looked further into the analysis to identify some of the soft benefits, which include self-service provisioning, compliance-driven data masking, role-based access controls, and real-time updates from the database to DevOps copies. Combined, these benefits help an organization to be more agile, deliver solutions to market faster, and tighten recovery point objectives (RPOs) since the DR copies are always current.

The Bigger Truth

ESG asked IT and data protection professionals what risks their organizations face due to their inability to get a holistic view of all of their data because they have too many data silos. Wasted IT budget (52%), focus taken away from strategic tasks (44%), and difficulty meeting SLAs (29%) all figured prominently.⁴ Considering that nearly a quarter of organizations (23%) reported a problematic shortage of data protection skills in 2020,⁵ leveraging Actifio with GCP is a strong combination that reduces cost and complexity while delivering on the promise to reduce the time to DR readiness.

ESG validation testing has confirmed that:

- Actifio was simple to deploy and use to protect a SQL Server database—creating a tailored recovery plan took only a moment and a few clicks.
- Actifio provided the ability to use Nearline object storage for DR copies with equivalent performance to the production environment, a significant capability for performance-sensitive applications like SQL Server. In addition, other groups like developers, testers, data analysts, and security specialists can leverage additional copies with the same characteristics.
- Automated policy-based backup, replication, and DR resulted in a three-year infrastructure cost reduction of 88%.

The results that are presented in this document are based on testing in a controlled environment. Due to the many variables in each production data center, it is important to perform planning and testing in your own environment to validate the viability and efficacy of any solution. ESG is confident that the conservative methodology and lack of locality of reference of the workload generators mean that organizations should see slightly better results than were presented in this report.

If your organization is looking to streamline data protection for your SQL Server databases and optimize disaster recovery, then ESG believes that you should consider the technical and economic advantages of Actifio and Google Cloud Storage.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.

The goal of ESG Validation reports is to educate IT professionals about information technology solutions for companies of all types and sizes. ESG Validation reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objectives are to explore some of the more valuable features and functions of IT solutions, show how they can be used to solve real customer problems, and identify any areas needing improvement. The ESG Validation Team's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments.



© 2020 by The Enterprise Strategy Group, Inc. All Rights Reserved.



www.esg-global.com



contact@esg-global.com



508.482.0188

⁴ Source: ESG Research Report, [The Evolution from Data Backup to Data Intelligence](#), February 2020.

⁵ Source: ESG Research Report, [2020 Technology Spending Intentions Survey](#), February 2020.