

Economic Audit

# The Economic Value of Hitachi Content Platform Storage

On-premises Versus Public Cloud Object Storage

By Brian Garrett, VP ESG Lab

November 2017

This ESG White Paper was commissioned by Hitachi Vantara and is distributed under license from ESG.



---

## Table of Contents

Table of Contents .....	2
Executive Summary .....	3
Challenges .....	4
Solution .....	5
TCO Analysis .....	7
Value at Scale .....	8
Value over Time .....	8
Other Considerations .....	9
The Bigger Truth .....	11
Appendix .....	12

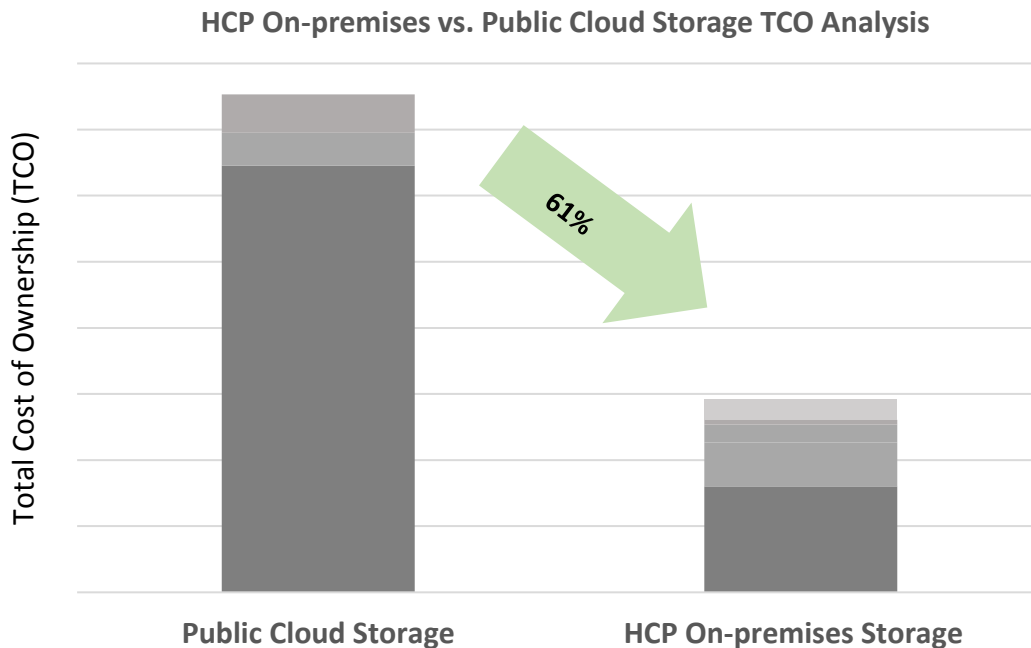
## Executive Summary

The public cloud has an unquestionable place in the IT world. It can be a great solution for transient workloads, test and development, and temporary dynamic workloads. But for *known* workloads with *data-intensive* unstructured storage requirements, the cloud may not live up to its cost-effective reputation. In fact, many IT organizations have been surprised by how their monthly bills grow due to egress fees, the cost of scale, and price premiums for multi-region recoverability. For this reason, many organizations are using the public cloud for some workloads while keeping known workloads or those dealing with sensitive data on-premises. In other words, instead of choosing between public *or* private cloud storage, forward looking IT decision makers are embracing the benefits of a hybrid cloud strategy based on the combination of public *and* private cloud storage.

It's a common misconception that renting public cloud storage is more cost-effective than purchasing private cloud storage. However, as you'll learn in this report, when you amortize all of the costs over a typical period of ownership of three to five years, the total cost of ownership for a private cloud is virtually always lower for most application workloads—especially for workloads that access hundreds to thousands of terabytes of data.

In this report, we'll explore the total cost of purchasing and owning Hitachi Content Platform (HCP) on-premises storage versus the costs associated with renting object storage from a leading public cloud vendor over time. As summarized in Figure 1, ESG has confirmed that the recently refreshed HCP solution, with new features including multi-site erasure coding, can be used to reduce the total cost of owning storage by 61% or more.

**Figure 1. HCP versus Public Cloud TCO**



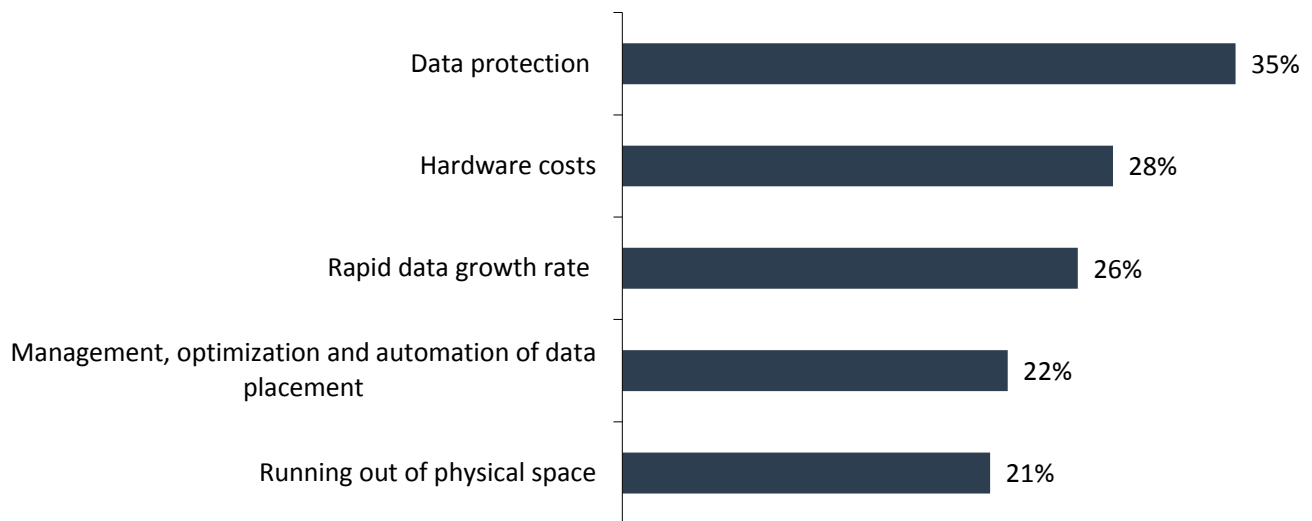
Source: Enterprise Strategy Group, 2017

## Challenges

The top three storage challenges most identified by respondents to a recent ESG research survey—data protection, hardware costs, and rapid data growth rate—have stayed the same from previous ESG research surveys from 2015 and 2016.<sup>1</sup> While 2017 finds data protection atop the list of challenges, the overarching issue that drives data storage concerns is relatively unchanged—data growth is accelerating and the resulting infrastructure required to store and protect that data is costly and complex.

**Figure 2. Top Storage Challenges**

In general, what would you say are your organization's biggest challenges in terms of its storage environment? (Percent of respondents, N=356)



Source: Enterprise Strategy Group, 2017

When storage decision makers were asked about their use of off-premises cloud resources in that same survey, 39% indicated that they had moved at least one workload from a public cloud infrastructure or service provider back to on-premises resources. This emerging trend of workloads being repatriated back from the public cloud to on-premises storage is being driven by concerns about the potential risks and costs associated with public cloud storage.

The public cloud movement is often driven from the top down based on a fundamental belief that the only way to attain agility while driving down costs is to move to the public cloud and rent resources as needed. But more frequently, IT organizations are realizing that tremendous IT agility and cost savings can be gained from building a private cloud in-house.

Public cloud storage is certainly an attractive option. Public cloud storage eliminates the physical administration and ownership of physical resources. It's quick and easy to deploy and is ideally suited for temporary and unpredictable workloads. The initial entry costs are attractive—no large upfront investment is required. However, entry cost is only part of the story: IT professionals need to consider the total cost of ownership. The balance of this report explores how, and when, HCP can be dramatically more affordable than public cloud storage over time. We'll also explore how the economic benefits of HCP can be amplified with a hybrid cloud strategy that provides centralized automation, visibility, and control as data is moved to the public cloud of your choice for the workloads that make sense for your organization.

<sup>1</sup> Source: ESG Brief, [2017 Storage Trends: Challenges and Spending](#), August 2017

## Solution

Hitachi Content Platform is a massively scalable object storage solution designed for secure private and hybrid cloud storage, content distribution and mobility, backup free archiving, and compliance. Designed for the cloud, it enables organizations to store and protect unstructured content such as documents, files, images, and video with massive scalability, and can retain content indefinitely. A single system can store data from different sources such as traditional NAS-based applications, file servers, email, content management repositories such as Microsoft SharePoint, and newer cloud, REST, and S3-based applications simultaneously.

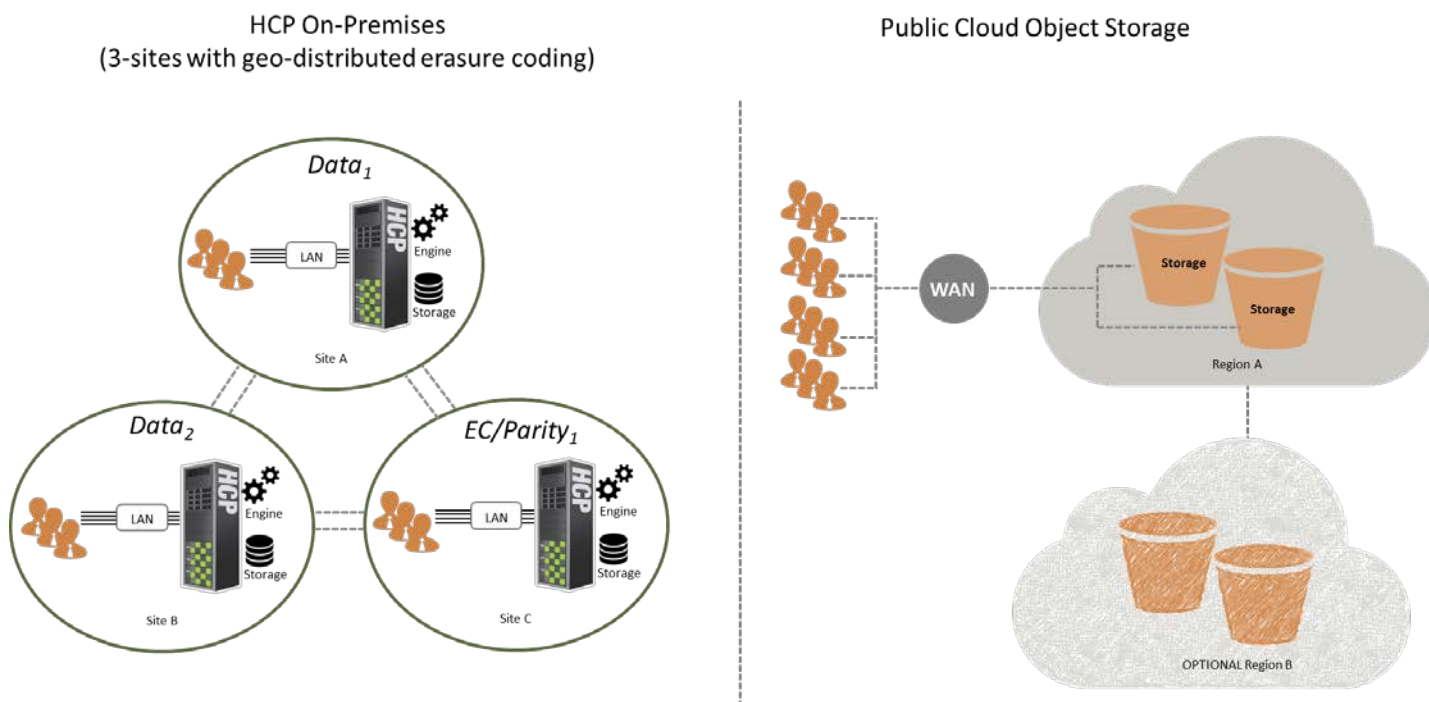
The HCP portfolio eliminates silos; promotes collaboration; enables governance and compliance; provides safeguards for sensitive data; automates management in private, hybrid, and multi-cloud environments; and surfaces insights through sophisticated search and analytics. The HCP portfolio includes Hitachi Content Platform for software-defined object storage; HCP Anywhere for file synchronization and sharing, and data protection capabilities necessary for the digital workplace; Hitachi Data Ingestor, an elastic-scale cloud file gateway; and Hitachi Content Intelligence, where rapid insights emerge from your data.

The latest update to HCP, announced in 2017, includes several enhancements that amplify the cost benefits of providing on-premises storage as a service compared to industry-leading public cloud solutions. Enhancements that deliver the biggest bang for the buck with the lowest total cost of ownership over time include the following:

- A hardware platform refresh including 10TB drive support increases HCP storage node capacity by 67%, the number of supported objects per node by 55%, and usable storage per cluster by 400%.
- Multipart upload support for faster file transfers.
- KVM support for a cost-effective alternative to legacy hypervisors.
- A simplified licensing model, which reduces costs by 20% for customers who only want to leverage foundational HCP capabilities.
- Geo-distributed erasure coding with local reconstruction codes reduces protection capacity overhead and speeds rebuild times after media failures in multi-site configurations.

See Figure 3 for a quick illustration and explanation of the value of geo-distributed erasure coding in a multi-site HCP configuration.

**Figure 3. Reference Architectures**



Source: Enterprise Strategy Group, 2017

The HCP solution on the left distributes parity data among three sites with a goal of surviving a site failure. In this case, erasure coding is spread over three sites, which reduces the amount of capacity, and cost, associated with the traditional disaster recovery method of mirroring data between two sites. Note that erasure coded data can be spread over more than three HCP sites, which reduces the cost of protection even more.

The solution shown on the right illustrates how public cloud object storage can be mirrored across two regions. Most public cloud storage customers don't run with this level of multi-regional recoverability due to its high cost: twice the amount of storage needs to be rented, and replication traffic between regions adds network access fees.

The cost benefits of the recently updated HCP solution are amplified by data reduction capabilities (single-instance storage and compression) that have been features of HCP for years.

HCP not only reduces the cost of ownership compared to the public cloud, it also reduces the risk and compliance concerns associated with public cloud storage. The HCP solution has a rich and growing set of security and compliance features including: encryption of data at rest, tamper-proof write once read many (WORM) support, search and indexing, and data retention and access control policies for auditing and discovery. Public cloud object storage solutions do not support any of these compliance features.

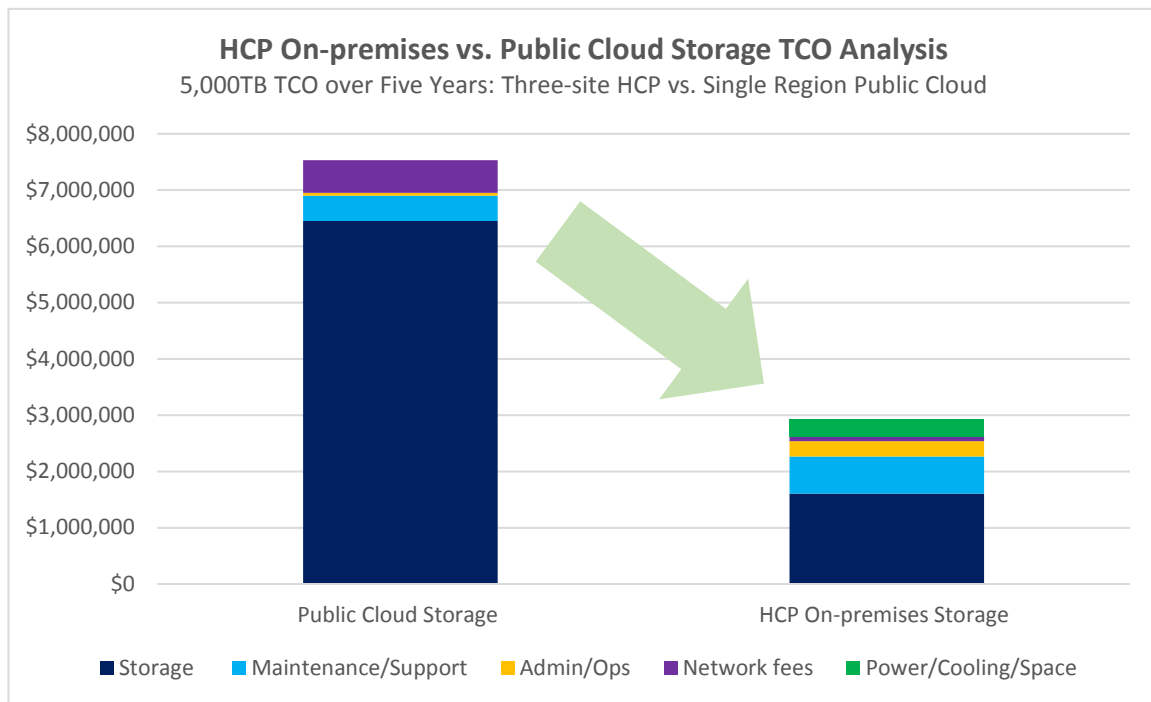
It's important to note that HCP on-premises storage can also be used as a complement to a public cloud strategy. Valuable hybrid-cloud-enabling capabilities include the ability to tier HCP data to the public cloud of your choice (Amazon S3, Google Cloud Storage, Microsoft Azure, and Hitachi Cloud Service Provider Partners, or to any cloud storage service that presents an S3 interface) and non-disruptive migrations between clouds.

## TCO Analysis

ESG used Hitachi pre-sales configuration planning and pricing tools, and publicly available pricing for an industry-leading public cloud vendor to compare the cost of renting public storage with the cost of purchasing and managing HCP on-premises storage over time. Our analysis began with a comparison of the costs associated with 5,000 TB of usable object storage capacity over a five-year period of ownership.<sup>2</sup>

As you can see in Figure 4, the monthly cost of renting public cloud storage capacity dwarfs the cost of purchasing, owning, and managing HCP on-premises storage over five years. In this example, the total cost of owning HCP on-premises is 61% lower than public cloud storage.

**Figure 4. HCP versus Public Cloud TCO**



Source: Enterprise Strategy Group, 2017

The primary cost consideration is the cost of storage capacity over time, as shown in dark blue in Figure 4. Much like the choice between leasing and purchasing a car, the monthly rental fee of public cloud storage seems attractive at first. Yet, over a period of years, those monthly leasing fees add up and dwarf the cost of purchasing HCP on-premises storage.

Note that the cost of power/cooling and on-premises data center space, shown in green, increases the cost of the HCP on-premises solution compared to the public cloud solution, which includes those costs in their infrastructure rental fees. These environmental costs, along with maintenance/support and admin/ops costs are higher for the HCP on-premises solution but relatively low compared to the total cost of ownership over time.

Last, but not least, note that the network fees are higher for the public cloud option due to the monthly fees for network access and egress. Stay tuned for more on this later in this report.

<sup>2</sup> For more details on the assumptions that were used for this analysis, please refer to the Appendix.


## Value at Scale

Next, we analyzed the cost of storage over three years as you increase the amount of capacity required to meet the demands of the business. Capacity requirements were scaled from 1 to 50 petabytes (1,000 to 50,000 TB) and the number of HCP sites was increased from one to three as the size of the deployment grew.

As you can see in Figure 5, The HCP cost advantage increases as you deploy more capacity:

- The 5TB HCP solution, which was mirrored over two sites, costs 60% less than the public cloud.
- The 50TB HCP solution, which uses capacity efficient erasure coding over three sites, costs 70% less.

**Figure 5. Value at Scale**



Source: Enterprise Strategy Group, 2017

## Value over Time

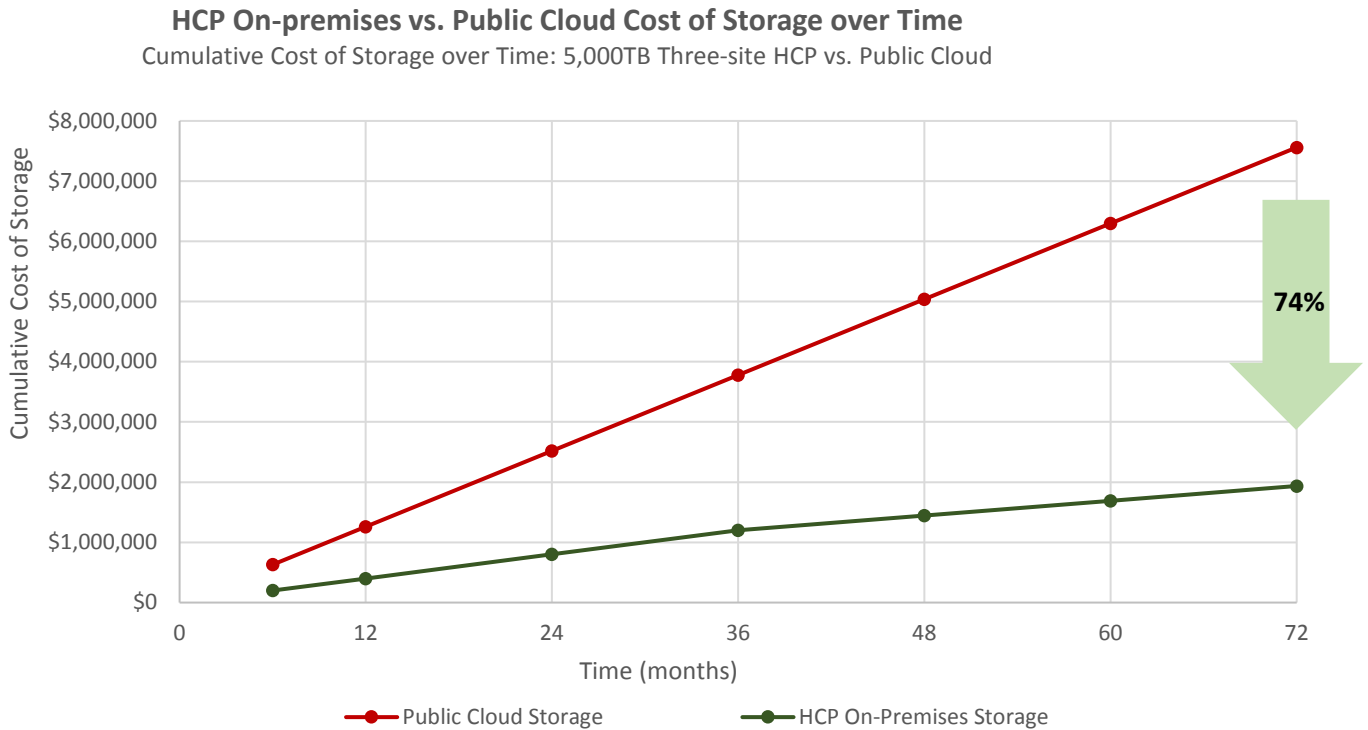
Finally, ESG analyzed the cost implications as the period of ownership increased from three to seven years. As shown in Figure 6, the savings associated with owning an HCP on-premises storage solution increase over time:

- The cost of HCP storage over three years is 68% lower than the public cloud.
- Note how the cost advantage widens due to an Autonomic Technology Refresh (ATR) that is assumed to take place at the beginning of year four. ATR enables seamless upgrading of HCP physical servers and/or storage at any time with no downtime. ATR also provides investment protection over time, with the option of an online refresh of the hardware without having to repurchase software.



- The cost advantage increases to 74% over a seven-year period of ownership.

**Figure 6. Value over Time**



Source: Enterprise Strategy Group, 2017

### Other Considerations

A variety of factors that influence the choice between public and on-premises storage were not included in our cost analysis, but should be considered, including:

**Multi-cloud Availability and Mobility:** It is a common misconception that public cloud storage is backed up and always available. While multi-region disaster avoidance is provided as an option by most public cloud storage providers, it adds complexity and cost that was not included in the modelling and analysis presented in this report.

The de facto standard for providing multi-region public cloud storage availability after a regional failure is to mirror objects between data centers in different regions. This more than doubles the cost of public cloud storage because you need to rent the same amount of storage in each region and pay more for network access as objects are mirrored between regions. This reality became painfully obvious to organizations that lost revenue during an East Coast Virginia AWS S3 outage in 2017 because they didn't architect and pay for multi-region availability.

Multi-cloud recoverability and policy-based mobility is built into HCP as a deployment option. The three-site geo-replication deployment configuration that was used for most of the analysis in this report has built-in multi-site recoverability. If any one of the three sites is impacted by a disaster, then capacity- and cost-efficient erasure coding technology ensures continuous access after a site outage.

**Risk and Compliance:** Moving to the public cloud introduces some level of risk. The organization becomes dependent on the cloud provider to secure and protect data. And while some providers are transparent about how they maintain

security, others are not. Also, the regulatory environment needs to be considered—knowing (and specifying) where data is stored, who can access it, how long it must be retained, and whether it is dispositioned properly is key. This is true for both primary and backup data. Many countries specify that certain types of data cannot be stored outside their geographical boundaries. This can make choosing a cloud provider onerous for organizations with business in multiple geographic regions, and could lead to the need to subscribe to multiple cloud providers, which adds administrative overhead. And still, lack of visibility into the entire data environment makes it hard for IT to parse what is “safe” for the cloud and what should be kept on-premises to ensure regulatory requirements are met. So, although businesses need agility, it must be balanced with risk.

An HCP solution that’s deployed on-premises is inherently more secure and compliant than public cloud storage. It lives behind your firewall and can be deployed and monitored by your staff. Data sovereignty and regulatory compliance are inherently easier when you own the infrastructure. HCP extends these risk reduction benefits with options including encryption at rest at no additional charge, patented “Secret Sharing” technology, and a variety of compliance and recoverability capabilities including write once read many (WORM) support with guaranteed immutability.

**Visibility, Control, Search, and Analytics:** Visibility and control of a scale-out unstructured storage solution is an important consideration, especially for applications and organizational workflows that are subject to compliance and governance initiatives. The HCP family of products delivers centralized visibility and control with integrated search and policy-based analytics. The cost of additional software and services that need to be layered on top of public cloud storage services to provide similar levels of visibility and insight were not considered in the analysis presented in this report.

**Performance and Data Gravity:** Where the applications, users, and data reside should be considered. If the data is born in the cloud and processed in the cloud, then the performance and latency impact of the WAN isn’t a major cost consideration. If significant volumes of data are born in the cloud and processed locally, then meeting performance requirements and the cost of WAN bandwidth needs to be considered. A modest HCP configuration can ingest data at speeds of up to 10 TB/hour; this is practically impossible and cost prohibitive to achieve with a public cloud due to the latency and cost of an outbound WAN connection to the public cloud. The cost and complexity rises even more when public cloud inbound WAN egress fees start to climb due to large volumes of data being retrieved and processed locally (e.g., for data mining or an audit).

**Multi-protocol Mobility:** Public cloud vendors support multiple storage access protocols with different services offerings at different price levels (e.g., AWS S3 for object and EFS for file). The integrated family of HCP products natively supports multiple protocols from a single platform, including the REST and S3-based object protocols of next-generation applications and the traditional file protocols that are used by file servers and content management solutions. HCP Anywhere, an enterprise file sync and share (EFSS) product within the HCP family, extends the value of multi-protocol file support to desktops and mobile devices.

**Price Erosion:** Public cloud prices have declined over the past five years and should continue doing so over the next five years. On-premises storage prices have declined as well, but for a fair comparison of the costs, it should be noted that the cost of an on-premises solution remains fixed over the cost-of-ownership period.

While ESG used an upfront fixed capacity acquisition cost for the analysis presented in this report, a pay-as-you-grow strategy with an on-premises solution that takes advantage of the latest storage media prices can be used to offset the effects of public cloud price erosion. The fact that cloud providers need to amortize the cost of the storage media that they’re buying today as they deliver storage as a service over the useful life of that media offsets the impact of public cloud storage price erosion as well. As an illustration of this fact of life for public cloud storage providers, the monthly \$/GB price for public cloud object storage capacity has been pretty much constant over the past 12 months while disk and solid-state media prices have dropped by 0.46% and 1.6% per month over the same period.

## The Bigger Truth

Utilizing public cloud is a tempting “quick fix” for building an agile IT environment, reining in IT costs, or starting up a new service. However, it may not be a cost-effective or long-term solution for many applications. ESG’s in-depth economic modelling and analysis has proven that the recently updated Hitachi Content Platform has a dramatically lower total cost of ownership compared to an industry-leading public cloud provider: 61% lower for 5,000 terabytes of object storage over five years.

ESG believes that forward looking IT decision makers would be wise to steer clear of the public *or* private cloud storage debate and instead consider the benefits of a hybrid cloud strategy based on the combination of public *and* private cloud storage. The HCP portfolio of products turbocharges your hybrid cloud strategy with valuable capabilities including centralized visibility and control, multi-protocol and mobile device support, integrated search and analytics, encryption at rest at no additional charge, and adaptive cloud tiering, which enables HCP to tier content to one or more public cloud providers. While public cloud storage services can be deployed with additional services and third-party software to deliver similar levels of functionality, they can’t match the levels of visibility, control, compliance, and risk mitigation of the HCP portfolio without adding significant cost and complexity.

Your mileage will vary depending on a variety of factors, including the nature of your applications, which public cloud provider you’re considering, and your compliance and security requirements. The level of savings will vary, but if your organization needs hundreds to thousands of terabytes, or more, of storage for unstructured data, then we’re confident that the economic benefits will hold true.

If your organization is using public cloud object storage as a service, check your bill and consider the savings you’ll achieve with HCP. If you’re considering the benefits of a scale-out object storage repository, use the lessons learned here to compare the costs of owning Hitachi Content Platform with those associated with renting a public cloud. If you’d like to learn more about the HCP family of solutions, click [here](#).

## Appendix

The cost factors and assumptions that were used for this TCO analysis include:

**Storage:** The publicly available pricing of an industry-leading public cloud vendor in \$/GB/month was compared to the cost of purchasing the same amount of usable storage for an HCP solution with a prototypical discount of 20% for hardware and 50% for software. The cost of usable HCP storage capacity included the benefits of compression, deduplication, and multi-site data protection, including erasure coding. Deduplication and compression rates of 18% and 19% were assumed based on the typical levels of capacity savings that have been achieved by HCP customers in production deployments.

**Maintenance and Support:** The monthly cost of business class support<sup>3</sup> from an industry-leading public cloud provider was compared with the cost of hardware maintenance and support for the HCP solution. An HCP Autonomic Technology Refresh (ATR), which provides an online hardware refresh without having to repurchase software licenses, was assumed at the end of the third year of ownership.

**Administration and Operations:** The on-premises manpower costs associated with provisioning and managing public cloud storage services were compared to the cost of deploying and managing an on-premises HCP solution. The analysis assumed that a fully burdened full-time engineer (FTE) with an annual cost of \$100,000 can manage five times more public storage capacity per year (50,000 TB of public cloud capacity versus 10,000 TB of on-premises HCP).

**Network Fees:** Public cloud network access and egress fees were calculated based on monthly access and egress rate assumptions of 7% and 3% respectively. HCP on-premises network fees were based on publicly available pricing for a pair of T1 WAN connections per HCP site.

**Power, Cooling, and Space:** The average US cost of electricity as of August 2017 (\$0.1053 per kWh)<sup>4</sup> and a typical co-location fee for data center floor space (\$9/month/40U) were used to calculate the cost of power, cooling, and space for the on-premises HCP solution. These costs don't apply to the public cloud storage solution.

<sup>3</sup> 24x7 electronic access to cloud support engineers and 1 to 24 hour telephone response time guarantees, depending on severity.

<sup>4</sup> Source: [Electric Power Monthly, U.S. Energy Information Administration](#), August 2017.

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 from time to time. 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.



**Enterprise Strategy Group** is an integrated IT research, analysis, and strategy firm that is world renowned for providing actionable insight and intelligence to the global IT community.

© 2017 by The Enterprise Strategy Group, Inc. All Rights Reserved.

