SHARDSECURE

White Paper

Cloud Resource Optimization and Cost Savings in AWS

Introduction: the cloud storage landscape

Now that they've migrated to the cloud, companies are asking what's next. For many enterprises, the answer is to strengthen their security and minimize their costs within their existing cloud storage. A January 2023 Financial Times article notes the growing burden to reduce cloud storage expenditures: "There's huge pressure from customers to reduce their costs," said Barry Briggs, a former Microsoft executive and current analyst at independent research firm Directions on Microsoft.

One of the top opportunities for cost-cutting is data storage. Semi-structured and unstructured data sets in particular are growing exponentially, driving higher monthly cloud subscription costs and massive resource consumption. Unstructured data as a whole is growing four times faster than structured data, with 55% to 65% annual growth, and it now makes up at least 80% of all enterprise data.

With all this data to pay for, organizations are increasingly turning to cloud resource optimization: rethinking their cloud storage, reducing their risk, and minimizing unnecessary spend. In this white paper, we'll cover one method of cloud resource optimization: achieving stronger security and higher cost savings in AWS without rewriting legacy systems.

Costs by Storage Type in AWS

Storage	Per GB per Month	250 TB per Year
S3 Infrequent Access	$ 0.013	$ 38,400
S3 Standard	$ 0.021	$ 64,512
EFS IA	$ 0.025	$ 76,800
EBS gp3	$ 0.080	$ 245,760
EBS gp2	$ 0.125	$ 384,000
EFS One Zone	$ 0.160	$ 491,520
EFS Standard	$ 0.300	$ 921,600

AWS storage services

As the world's largest cloud service provider, AWS offers many different solutions to meet a myriad of storage needs for unstructured data. Choosing the right solution can help companies cut costs, increase agility, reduce time to market, and eliminate infrastructure maintenance.

There are several major types of data storage in AWS, including:

- **Amazon Elastic Block Store (EBS):** High-performance block storage service designed for transaction-intensive workloads at any scale. EBS comes in the form of a virtual hard disk attached to Amazon Elastic Compute Cloud (EC2) instances in a way that's similar to the local disk drive on a physical machine.

- **Amazon Elastic File System (EFS):** Highly elastic, scalable set-and-forget file system that allows users to share file data without provisioning or managing storage. It can be used with both cloud and on-premises data, allowing organizations to grow and shrink their file systems automatically as they add and remove files.

- **Amazon Simple Storage Service (S3):** Highly affordable object storage. S3 is cost-efficient, scalable, and ideal for cloud storage. It can provide significant cost savings compared to not only other AWS storage types but also traditional storage like SSDs, NAS and SAN.

These types of AWS storage meet different operational needs, but they also have drastically different costs. For instance, the price of S3 Infrequent Access is $0.013 per GB per month, while the price of EFS Standard is $0.3 per GB per month. At those prices, 250 TB of data per year would cost $38,400 in S3 versus $921,600 in EFS — a factor of 24.

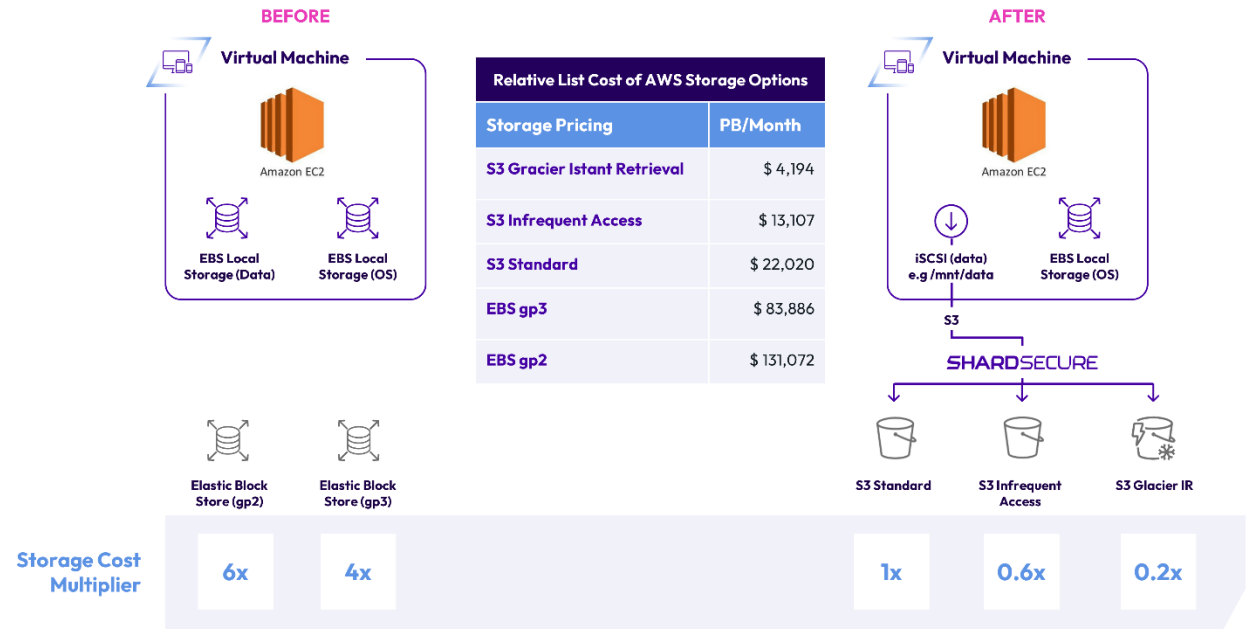

Cost impact of moving from EBS to S3 storage

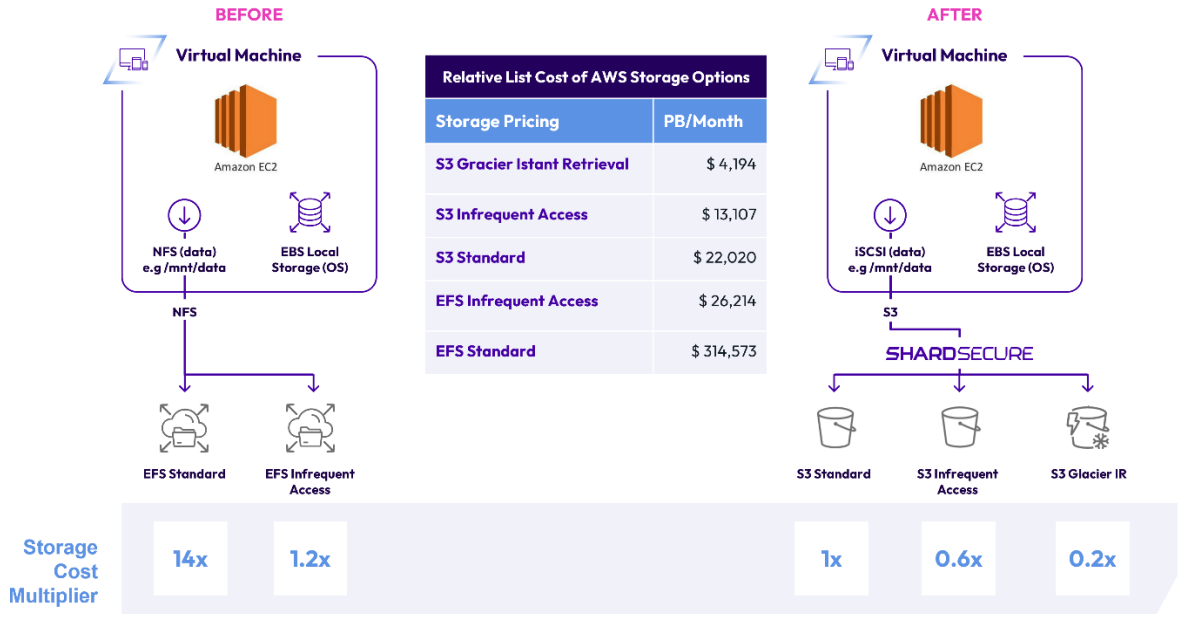

Cost impact of moving from EFS to S3 storage

Legacy storage systems

Companies want to leverage the most affordable storage services they can. For AWS, that means object storage like S3. But legacy applications usually rely on block devices like hard drives or network attached storage (like EFS) to provide traditional file system access to files. That means they need to be rewritten to support object storage.

So, if an organization's application relies on legacy storage like EBS or EFS, they will have to rewrite their entire application to take advantage of the cost savings of S3.

Misconfigurations in AWS

Choosing the most cost-effective AWS storage is crucial — but so is protecting sensitive data once it's moved to that storage. While AWS offers strong durability and some security features of its own, misconfigurations and breaches are still all too common.

The shared responsibility model means that companies are responsible for the security of their data in the AWS cloud, while AWS is responsible for the security of the AWS cloud. Essentially, this means that businesses bear the responsibility of avoiding mistakes that publicly expose their own data.

Unfortunately, common misconfigurations like public access to S3 buckets, outdated IAM policies, key rotation problems, and unsecured backup storage leave data exposed. In 2021, more than 1.6 million files containing personal identifiable information (PII) from 80 US municipalities were left exposed. In 2022, 6.5 TB of airport data was exposed in AWS, including navigation information, proprietary software, and airline crew PII. Other recent AWS misconfiguration incidents include the exposure of the personal data for 3 million senior citizens via the website SeniorAdvisor, 3 million people via the online booking website FlexBooker, and 50,000 patients via the Utah-based Covid testing company Premier Diagnostics.

The bottom line? Moving to a more cost-effective storage tier in AWS is only successful if it doesn't lead to a data breach. Fortunately, strong data protection software can mitigate the impact of not only misconfigurations but also AWS outages, ransomware attacks, and more.

Greater cost savings and flexibility with ShardSecure

ShardSecure allows organizations to optimize their cloud usage and protect their data in AWS, including in multi- and hybrid-cloud infrastructures, without ever rewriting applications, redesigning data flows, or changing their user experience. With ShardSecure's transparent plug-and-play technology, companies can easily leverage object storage like AWS S3 and enjoy vastly improved data security and resilience.

The ShardSecure solution leverages iSCSI, a traditional SAN protocol, to resemble EBS or EFS storage to the server. The technology's transparency means that data can be moved to S3 on the backend for maximum cost savings with virtually no changes to a company's applications.

The solution is ideal for enterprises with large amounts of data — that is, 250TB or more of EFS storage. ShardSecure's performance is very similar to EFS's, so data access remains fast and easy for users.

Advanced data security in AWS

ShardSecure's solution protects unstructured data and metadata in specific files, folders, and storage locations. By splitting data into very small pieces (microshards) and then distributing those containers to multiple customer-owned storage locations, we ensure that data is unintelligible to everyone from cloud providers to cyberattackers.

Companies can use multiple AWS buckets, a mix of AWS and other storage providers, or even AWS and on-prem data centers to store their material. Regardless of the configuration they choose, the data will remain secure from internal and external threats.

Even in the highly unlikely scenario that someone is able to gain access to all the microshards in AWS for a given data set, ShardSecure's technology ensures that those microshards cannot be reconstructed.

- The solution strips file content, filenames, file extensions, and all other metadata, meaning that there is not enough identifying information for reassembly.
- The solution allows organizations to add a configurable amount of poison data to their real data.
- The technology also requires multiple components to be used in concert with both each other and the complete data set for reassembly, meaning that it's not possible for an unauthorized user to deploy their own instance of
- ShardSecure to reconstruct data.

Native ransomware protection

In addition to ensuring data privacy, ShardSecure also offers strong protection in the face of ransomware and other cyberattacks. Our self-healing data feature can detect when data is deleted, tampered with, or encrypted in ransomware attacks. Reconstruction of affected data then begins automatically and transparently, without costly downtime or disruption to users.

This process also serves as an early detection method, with automatic alerts to security teams for fast incident response. And, with strong data confidentiality, it mitigates the impact of double extortion: Any data exfiltrated by an attacker will remain unintelligible and useless to that attacker.

As a result, organizations can maintain their critical operations and prevent any loss of data or data access in the face of serious ransomware attacks.

Highly available file system with redundancy and failover

ShardSecure achieves high availability at multiple levels. First, each instance of ShardSecure is a virtual cluster that can be run on-premises or in the cloud. Second, customers can configure two or more virtual clusters for failover.

The same self-healing data feature that neutralizes ransomware also works to keep data accurate and available during cloud outages, server crashes, severe weather events, physical storage device failures, and more.

Even in the event of human error — the single largest cause of cloud network security breaches — or those all-too-common cloud misconfigurations, ShardSecure's strong data resilience prevents the loss of business continuity.

Easy integration, migration, and access

ShardSecure is easy to manage and has a low impact on operations teams. A plug-and-play approach, it provides an easy and transparent implementation with no need to change user behaviors or data flows.

ShardSecure's technology also allows for instant data access and fast data migration with just a few clicks. It is quick and seamless to integrate, with only one line of code change needed to get started.

Conclusion

The cloud has been profoundly transformative for many organizations over the last five years. But many are still stuck using yesterday's storage for today's challenges.

ShardSecure helps companies leverage significant cost savings for large quantities of data in AWS. With our strong data privacy technology, native ransomware protection, high availability, and ease of integration, we enable organizations to optimize their cloud usage and reach for tomorrow's successes.

For more information on ShardSecure and cost savings in the cloud, visit us online or schedule a demo today.

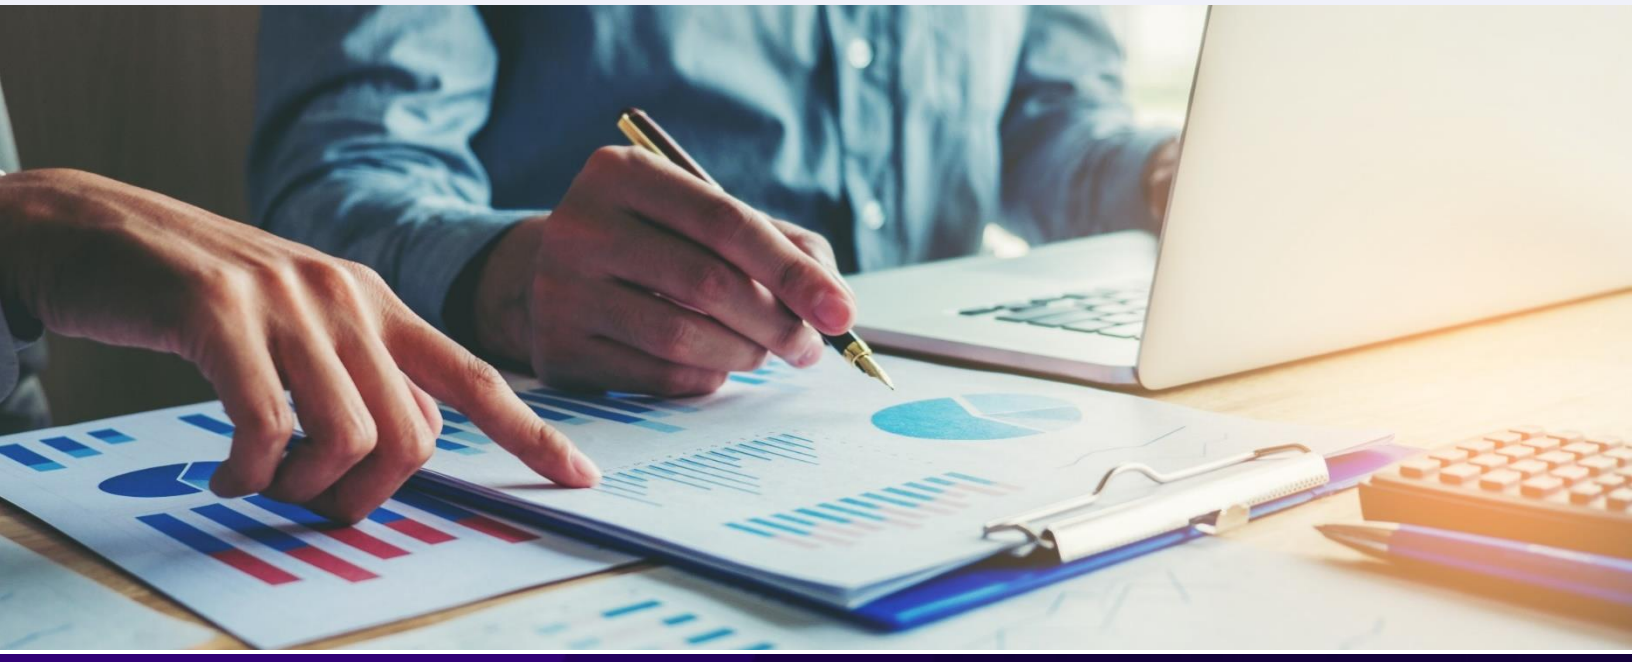

@ShardSecure
@ShardSecure
@ShardSecure

101 Avenue of the Americas
9th Floor, New York, NY 10013
United States of America

info@shardsecure.com