

.00, Jesse J Perez, 355
Reese, 241, Jun-11, 201
745.00, Jeremy P. P
59 SHARDSECURE
594.00, Joe N. Ho

Technical White Paper:

Enabling Resilience and Regulatory Compliance for Multi-Cloud Backups



Current challenges in the data resilience landscape

The cybersecurity industry generally recognizes two cornerstones of resilience, data availability and integrity. Both face significant challenges in today's digital landscape. In addition to the common cyber risks that threaten data resilience, organizations also face the complexity of managing and securing hybrid- and multi-cloud architectures.

Backups are one way that enterprises have traditionally ensured data redundancy and recovered from disasters. But protecting these backups and ensuring business continuity has become another challenging facet of multi-cloud resilience. Backups have historically been kept in a secure and air-gapped location, but a true air-gapped approach is not possible in most modern-day environments (e.g., cloud services). As a result, logical air gapping has become necessary to prevent threat actors from tampering with data backups.

This technical white paper explores the top challenges associated with backup storage and proposes a solution for true multi-cloud resilience.

Data availability

High availability — the ability to operate continuously despite failures — is key to achieving data resilience in multi-cloud architectures. Although it doesn't prevent cyberattacks or other major events, high availability does ensure that business operations can continue in the face of disruptions.

To achieve high availability within the backup storage architectures, single points of failure need to be eliminated. Relying solely on a single cloud provider and its uptime puts organizations at risk, as that cloud provider inherently becomes the single point of failure. Instead, high availability can be achieved by leveraging a true hybrid or multi-cloud approach and having redundant backups available in different storage services.

Data integrity

Data integrity refers to the accuracy, consistency, and reliability of data. It means that every file stored and accessed is byte-for-byte identical to the file that was originally written, with no corruption, tampering, or unwanted modification by authorized or unauthorized users.


Data integrity can be ensured by eliminating any unauthorized changes to the data. Backups should be stored within an immutable system that prevents unauthorized users and threat actors from deleting, tampering with, or otherwise changing the data.

The shift to data redundancy

In the past, organizations achieved resilience by making their systems redundant — typically through using a primary data center and a backup data center with identical infrastructure, devices, and settings. By duplicating each component or system across multiple data centers, companies ensured that their data would remain available and accurate even if one system or storage location went down.

Today, however, many organizations are starting to focus more on data redundancy. This model requires companies to duplicate only their data across multiple cloud storage providers, helping save money on infrastructure costs.

Unfortunately, the data redundancy model still requires organizations to double up on storage fees. This is forcing companies to look for new ways to achieve data resilience without paying for full backups across multiple clouds.

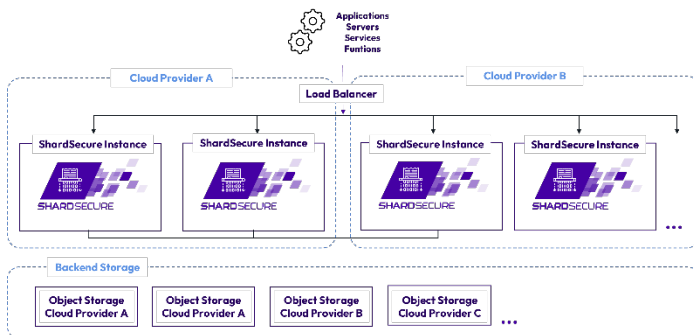


Multi-cloud data resilience with ShardSecure

The ShardSecure platform offers multi-cloud resilience for unstructured data at rest without the need for system or full data redundancy. It also offers several features to support data availability and integrity and provides the ability to reconstruct compromised data in multi-cloud architectures.

Achieve high availability in multi-cloud environments

ShardSecure achieves high availability at multiple levels. First, each instance of ShardSecure is a virtual cluster that can run on-prem, in the cloud, or in hybrid architectures. Second, customers can configure two or more virtual clusters for failover, which provides high availability across multiple clouds as well as in hybrid-cloud environments that use a mix of on-prem, private cloud, and third-party public cloud services.

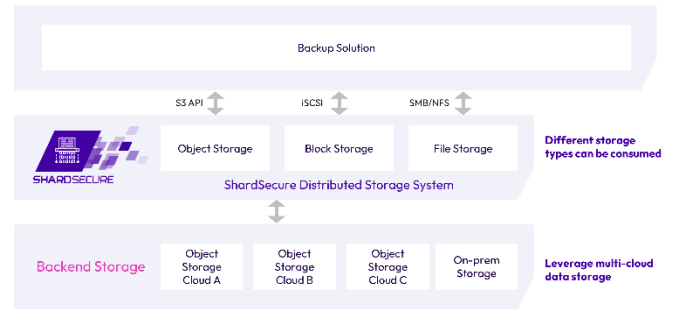


Simple integration across multiple clouds

ShardSecure supports interfaces like S3 for object storage, iSCSI for block storage, and SMB/NFS for file storage to integrate with commercially available backup solutions.

Since the ShardSecure platform acts as an abstraction layer, it can consolidate all storage interfaces into one, even though each cloud provider may support a different data storage interface. This approach reduces the complexity associated with migrating to a multi-cloud

architecture and implementing backup solutions across the enterprise. With ShardSecure, data access operates consistently across all clouds without the need to implement specific APIs or connectors.



Ensure accuracy with multiple data integrity checks


While traditional data resilience solutions focus only on data availability, ShardSecure offers multiple health checks for data integrity as well. These data integrity checks detect unauthorized modifications, including those made by ransomware attacks and malicious tampering.

If a storage location fails a data integrity check, the ShardSecure platform automatically begins to reconstruct the affected data in real time. This ensures that data remains not only available but also accurate and unaltered.



Transparently reconstruct data, including live production data

By creating slight overlaps of distributed data across different storage locations, the ShardSecure platform can reconstruct compromised data in real-time. The platform supports a combination of parity data, redundant storage, and responsive storage to automatically recover any lost data. This self-healing process means that organizations can lose up to half of their storage buckets and still access the underlying data. It also ensures business continuity during outages, downtime and data integrity issues like data tampering.



The self-healing process is transparent to users and existing data workflows and does not impact performance, allowing applications and servers to continue reading/writing production data.

As the diagram above shows, parity data is stored within Tier 1 and Tier 2 storage in case of single bucket outages. Under normal operating conditions, Tier 2 storage only receives write operations, helping customers leverage cool/cold storage. Meanwhile, responsive storage extends both Tier 1 and Tier 2 storage to an additional standby storage bucket. If a single storage bucket experiences an outage, data corruption, or service limitation (e.g., the write command does not work but the read command does), the responsive storage will detect the unhealthy bucket and move its data to the standby bucket until the unhealthy bucket is restored.



Regulatory requirements for secure backup storage

The EU's EBA (European Bank Authority) and Germany's BAFIN (Federal Financial Supervisory Authority) require organizations subject to SERP (Supervisory Review and Educational Process) to ensure that data backups are not affected by the same potentially damaging events as production data. These regulations stipulate that backup data be stored in separate infrastructure from production data to ensure that it will not be affected by

any disaster. This is especially important if an organization relies on a single data center or single cloud provider for infrastructure. In case of a disaster, backup data needs to be accessible at all times.



Privacy of data stored in the cloud

Storing sensitive data with a third party can raise privacy concerns for many organizations. Regulations like the GDPR create hurdles by limiting how personal data can be processed and stored, and the Schrems II ruling further restricts the measures that organizations can use to ensure data privacy for EU personal data. As a result, storing European data with a US cloud provider presents challenges for many multinational and EU-based organizations.

ShardSecure's approach to storing backup data in the cloud has been validated by independent privacy attorneys to meet the requirements of Use Case 5 of the EDPB's recommendations for cross-border data transfers, allowing organizations to store EU personal data within a US cloud provider without violating the GDPR.



Conclusion

Given the continued acceleration in digital transformation, legacy approaches to backup storage are no longer adequate. Companies need to adopt new strategies and technologies to ensure the continued availability and integrity of their critical data in multi-cloud architectures.

The ShardSecure platform offers an innovative approach to multi-cloud data resilience. With its self-

healing feature, high availability, and data integrity checks, the platform ensures unstructured data resilience, regardless of storage location.

For more information on how ShardSecure is improving resilience for companies in the [financial services](#), [manufacturing](#), and [high-tech](#) sectors, take a look at our [other data resilience resources](#) 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

**SHARD
SECURE**