SHARDSECURE

Technical White Paper:

Challenges with Achieving Multi-cloud Data Resilience

Current challenges in the data resilience landscape

The cybersecurity industry generally recognizes two cornerstones of resilience, data integrity and data availability. Both face significant challenges in today's digital landscape. In addition to common cyber risks that threaten data resilience, organizations also face complexity in hybrid- and multi-cloud architectures. Not only must companies choose the most cost-effective and secure storage for their data, but they must also consider different models for data redundancy and weigh the risks of relying solely on data backups to ensure business continuity.

This technical white paper explores these top challenges to data resilience and discusses 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, organizations may employ solutions ranging from basic backup snapshots or a few spare servers to a fully redundant network infrastructure with automatic failure detection and failover.

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 to storage, with no corruption, tampering, or unwanted modification by authorized or unauthorized users.

Data redundancy

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

Today, many organizations are turning to 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 resilience with ShardSecure

The ShardSecure software platform offers efficient multi-cloud resilience for data at rest without the need for multi-cloud redundancy or full data backups. Our technology also offers several features to support data

integrity and availability, including 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 hybrid. 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-premises, private cloud, and third-party public cloud services.

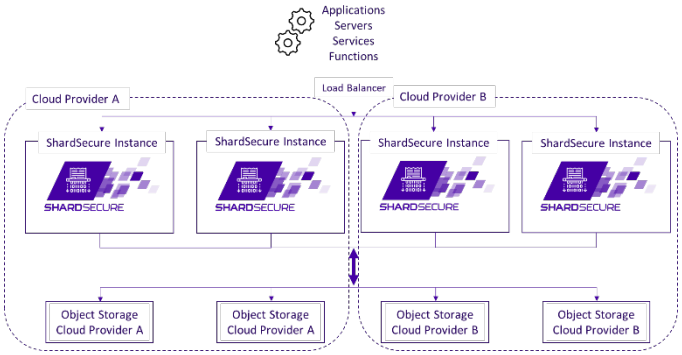

Shift to data integrity

While traditional data resilience solutions focus only on data availability, ShardSecure offers multiple 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 alerts the SOC or security team and begins to reconstruct the affected data. This helps ensure that data remains not only available but also accurate and unaltered.

Unified, multi-protocol across multiple clouds

ShardSecure supports interfaces like S3-compatible API for object storage, iSCSI for block storage, and SMB/NFS for file storage to integrate with any application, server, or other service leveraging cloud storage.

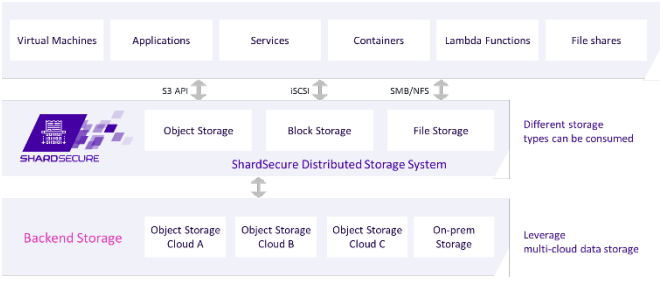

Since the ShardSecure platform acts as an abstraction layer, it can consolidate all storage interfaces into one, even though each cloud provider may support a different data storage interface. This approach reduces the complexity associated with migrating to a multi-cloud architecture. With ShardSecure, data access works the same across all clouds without the need to implement specific APIs or connectors.

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 feature means that organizations can lose up to half of their storage buckets and still access the underlying data.

The self-healing feature works during outages and downtime as well as any data integrity issues like data tampering. If a data storage location or the data within it becomes unhealthy for any reason — a cloud provider issue, a misconfiguration, etc. — the same process reconstructs the unavailable data.

The self-healing process is transparent from the user impact, data workflow, and performance points of view. Although SOC and security teams will be alerted about any data integrity issues, the application or servers reading/writing the data will not experience any data outages or latency from storage.

As the diagram above shows, parity data is stored within primary and secondary storage in case of single bucket outages. Under normal operating conditions, secondary storage only receives write operations, helping customers leverage cost savings for 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 limitations (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.

ShardSecure mitigates risk by keeping data available in multi-cloud and hybrid-cloud configurations. Data resilience is vital for organizations in the face of tampering, deletion, outages, ransomware, and other unexpected events. Everything from accidental misconfigurations to extreme weather can cause companies to lose access to their most important data.

Conclusion

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

The ShardSecure platform offers an innovative approach to multi-cloud data resilience. With its self-healing, high availability, and data integrity capabilities, our technology ensures unstructured data resilience, regardless of storage location.

For more information on how ShardSecure is improving resilience for companies in financial services, manufacturing, tech, and more, take a look at our other data resilience resources or schedule a demo today.

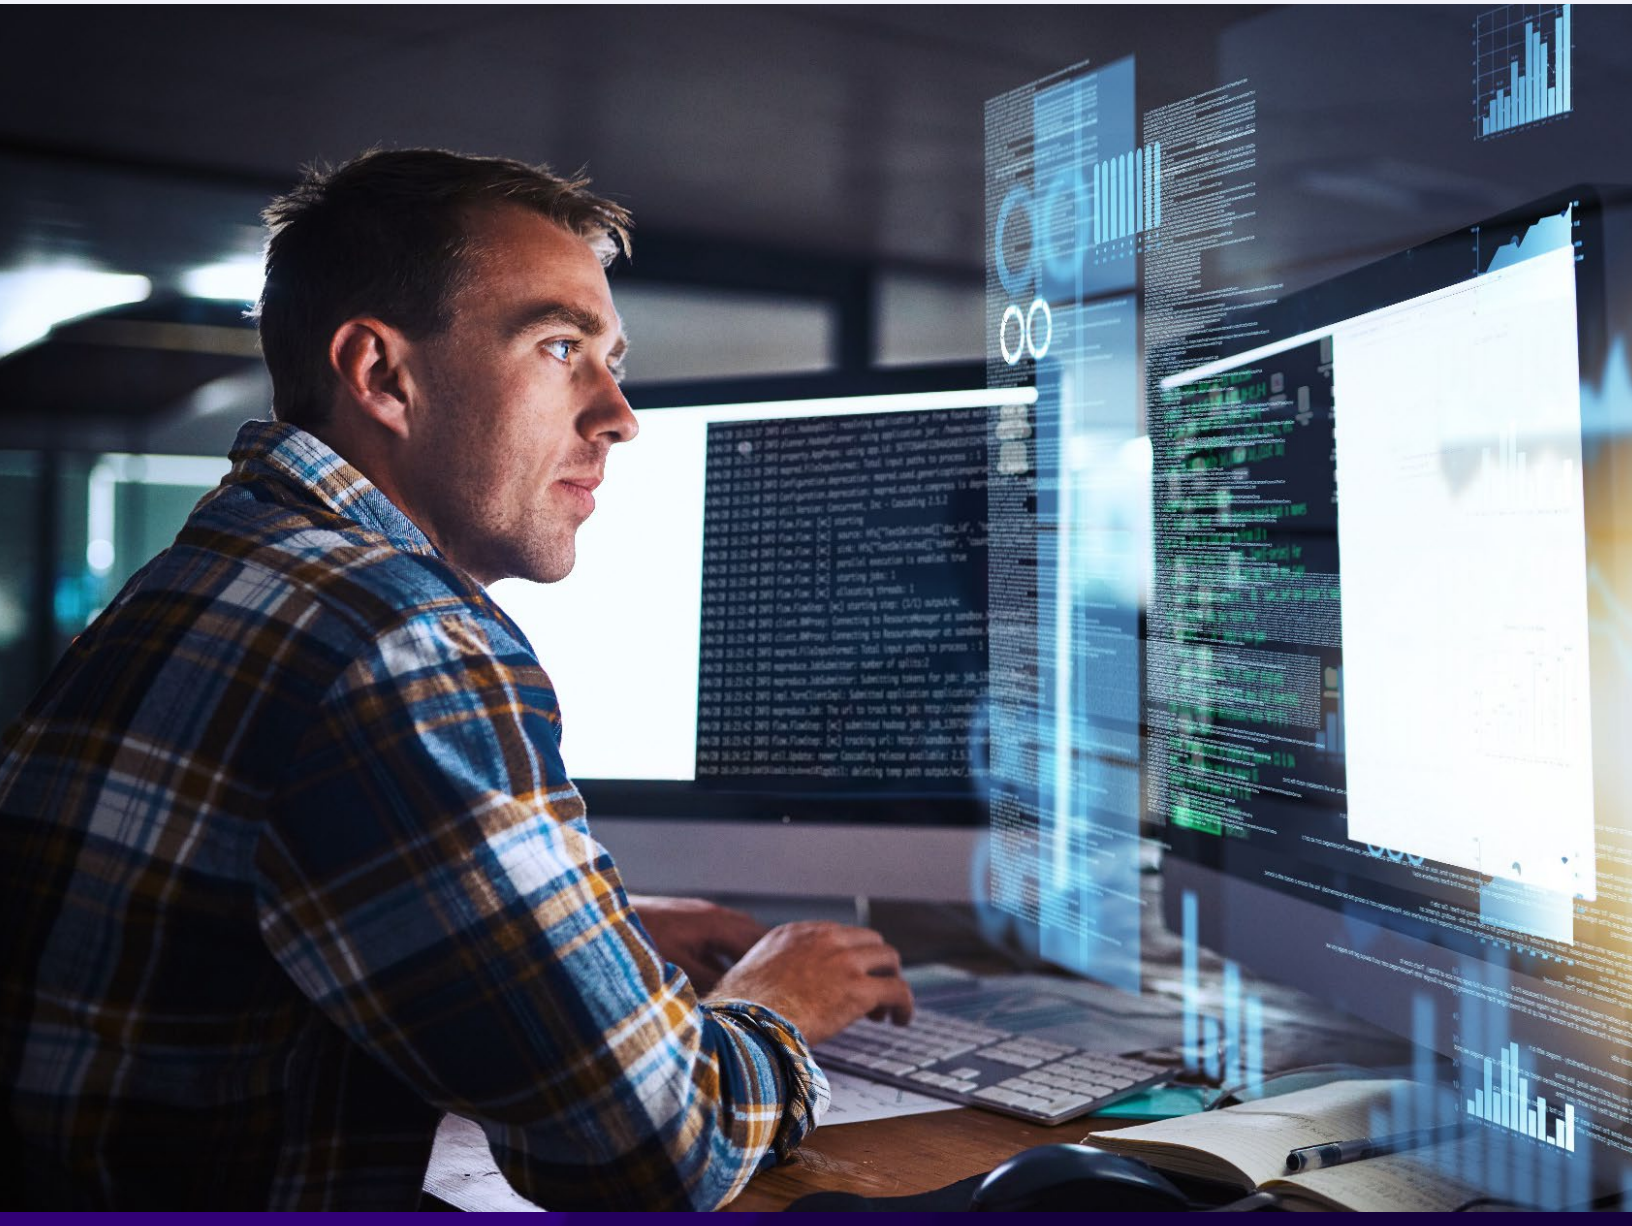

@ShardSecure
@ShardSecure
@ShardSecure

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

info@shardsecure.com

SHARD SECURE