HAMMERSPACE

Protecting Data – and Metadata – in the Hammerspace Data Platform

V. 1.1 - December 2025

Introduction

In the Hammerspace Data Platform, metadata doesn't just describe your data—it actively governs and protects it. This metadata-centric approach leverages Service Level Objectives (SLOs) to automate data protection policies, ensuring that data is managed efficiently and securely across the entire environment.

This whitepaper details how metadata is both in charge of protecting data and is itself protected through multiple robust methods that span various RPO and RTO requirements.

Why Put Metadata in Charge of Data Protection?

Hammerspace's unique architecture separates the metadata layer from the data layer, then uses declarative policies known as Service-level Objectives (SLOs) to automate the protection and movement of data.

Using metadata to automate data protection has these advantages:

1. **Centralized Control Through Metadata:** Metadata serves as the intelligent core of the Hammerspace data platform. It records the location, access patterns, and policies for every piece of data, allowing for efficient management even at massive scales involving billions of files. This centralization simplifies operations like data replication, migration, and recovery.
2. **Automated Enforcement with SLOs:** By defining SLOs, organizations set the desired outcomes for data availability, durability, and performance. The metadata database enforces these SLOs automatically, orchestrating data placement, replication, and access controls without manual intervention.
3. **Efficiency and Scalability:** Managing billions of individual files and objects is complex and resourceintensive. In contrast, a metadata database can handle billions of records efficiently, utilizing advanced database technologies for indexing, transaction management, and optimized storage. This makes backup, replication, and recovery operations faster and more reliable.
4. **Metadata-Driven Data Protection:** With metadata at the helm, the system can make informed decisions quickly, optimizing resource utilization and minimizing downtime. This ensures that data is protected according to organizational policies and adapts dynamically to changes in the environment.
5. **Using Checksums to Guarantee Data Integrity:** The Hammerspace metadata database maintains checksums of the data as a way to guarantee the data remains correct over time.

Protecting a database with billions of records is significantly more efficient than safeguarding a file system containing billions of individual files. The reasons are shown in Figure 1 below.

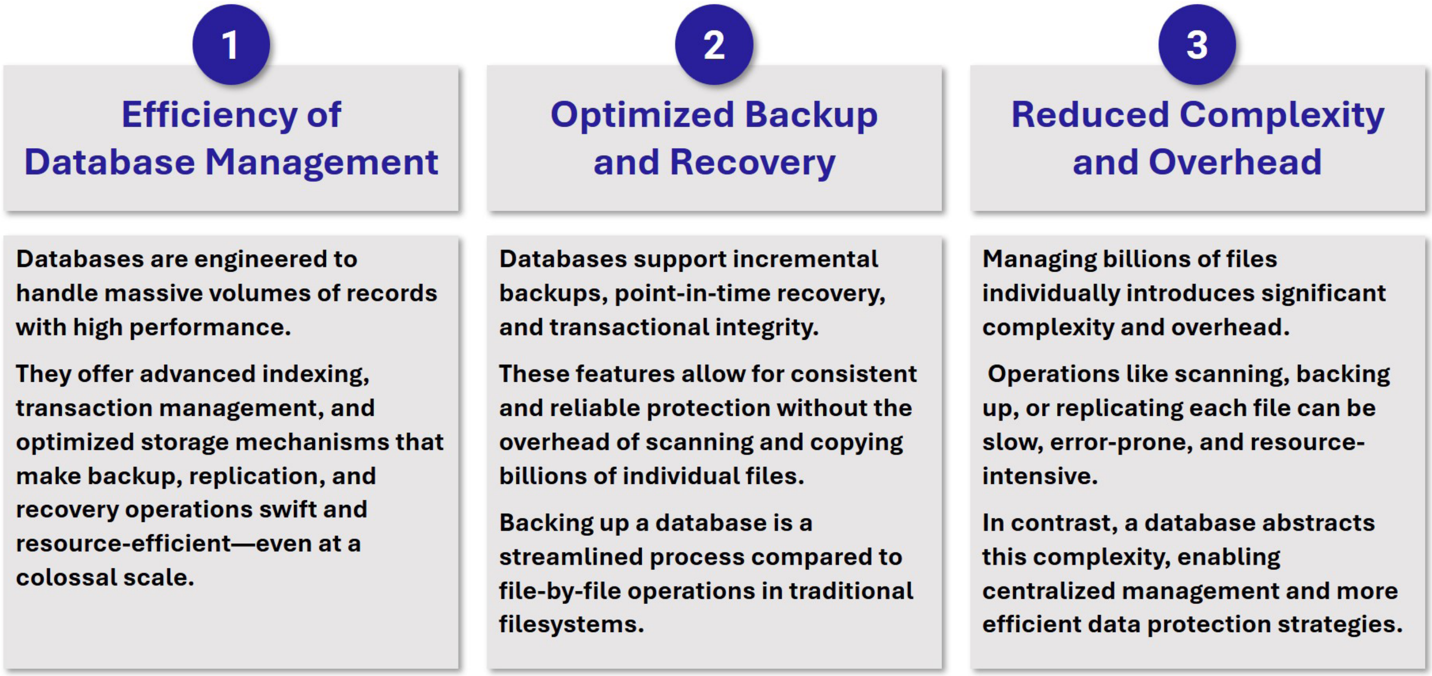

Figure 1 – Why Protecting a Metadata Database is Significantly More Efficient

Comprehensive, Multi-Layered Approach for Protecting Metadata in Hammerspace

The Hammerspace Data Platform employs a comprehensive, multi-layered approach to ensure reliable retention and protection of metadata. This approach includes five distinct types of redundancy, each capable of multiple levels of duplication—2-way, 3-way, or even N-way redundancy – to enhance data protection.

In most Hammerspace environments it is common to employ all of these methods to achieve optimal resilience and availability.

The methods are summarized in Figure 2 below, then discussed in more detail.

	Method 1 Mirroring within Anvil using block storage	**Method 2** Synchronous mirroring between Anvils in a cluster	**Method 3** Async replication between Hammerspace clusters	**Method 4** Periodic checkpoints of metadata database	**Method 5** Export of metadata to JSON for reconstitution
RPO	**Zero**	**Zero**	**Seconds**	**Configurable** (depends on checkpoint frequency)	**Variable** (depends on export frequency)
RTO	**Zero**	**Zero**	**Zero**	**Minutes to Hours** (more data=longer)	**Minutes to Hours** (more data=longer)
Overhead	**Minimal Impact**	**Minimal Impact ***	**Minimal Impact ***	**Minimal Impact**	**Minimal Impact**
Distance	**Single Site / Single Availability Zone**	**Across Availability Zones** (stretch clusters, metro clusters)	**Across Distinct Data Centers Globally**	**Not Distance-Dependent**	**Not Distance-Dependent**
Redundancy Level	**2-way, 3-way, N-way**	**2-way, 3-way**	**2-way, 3-way, N-way** (dozens of sites)	**Multiple Checkpoints** (N-way over time)	**Muiltiple Exports** (N-way over time)
Notes	Mirroring enhances reliability of both storage types	Automated failover with Pacemaker & Corosync	Replication interval typically in seconds	Restore time increases with data volume	Enables independent data recovery

**Possible performance reduction witih high latency or large N*

Figure 2 – Hammerspace Multi-Layered Approach to Metadata Protection

Definitions:

- **RPO (Recovery Point Objective):** Maximum targeted period in which data might be lost due to a major incident.
- **RTO (Recovery Time Objective):** Duration within which a business process must be restored after a disruption to avoid unacceptable consequences.
- **Overhead:** Impact of the redundancy method on system performance.
- **Distance:** Geographical scope over which the redundancy method can be applied.
- **Redundancy Level:** Number of copies or replicas maintained to protect metadata (e.g., 2-way, 3-way, N-way).

METHOD 1

Mirroring of Storage Within the Metadata Server (Anvil) or Use of Reliable Block Storage

Description:

The metadata server, known as Anvil, can utilize either direct-attached storage (DAS) like NVMe drives or reliable shared block storage.

Implementation:

- **Direct-Attached Storage (DAS):** Utilizes local SSDs or NVMe drives within the Anvil server, with mirroring to protect against drive failures.
- **Reliable Block Storage (Optional):** Employs highly reliable shared block storage solutions like Storage Area Networks (SAN) or networked Elastic Block Store (EBS), which provide inherent redundancy.

Benefits:

- **Zero RPO (Recovery Point Objective):** No loss of committed data.
- **Zero RTO (Recovery Time Objective):** Uninterrupted data access with only a brief metadata stall.
- **No Overhead:** Minimal impact on the running system.

Use Cases: Ideal for environments within a single availability zone where immediate failover capabilities are required.

Redundancy Levels: Supports multiple redundancy configurations (e.g., 2-way, 3-way, N-way mirroring) within the same availability zone.

Note: Starting with reliable block storage is optional; many environments prefer DAS NVMe becaues it is faster.

METHOD 2

Synchronous Mirroring Between Anvil Nodes

Description:

The metadata database is automatically synchronously mirrored across the two or three Anvil nodes in a Hammerspace cluster. If possible, Anvils should be separated to reduce the chances of them all going down at once, which would impact file system availability. Separation may be as simple as powering them from different power rails and housing them in different rack rows, but stretch clusters are an option where low-latency connectivity is available. With stretch clusters, Anvils may be located in different buildings on a campus or within a metro area.

Implementation:

- **Synchronous Replication:** Metadata is instantly replicated between Anvil nodes over low-latency connections.
- **Automated Failover:** Utilizes tools including Pacemaker and Corosync for automated hot failover between nodes.

Benefits:

- **Zero RPO and RTO:** Ensures no data loss and uninterrupted access.
- **High Availability:** Provides resilience against node failures across availability zones.

Use Cases: Recommended for all Hammerspace deployments, and definitely a best practice.

Redundancy Levels: Can be configured for multiple redundancy levels (e.g., 2-way, 3-way), depending on the number of Anvil nodes.

Performance Consideration: Low-latency links are required between Anvils to ensure metadata replication doesn't impact system performance.

METHOD 3

Asynchronous Replication of Metadata Between Anvils in Separate Hammerspace Clusters

Description:

Metadata is asynchronously replicated between Anvils in completely separate Hammerspace clusters using a shared object store or cloud bucket. This can involve dozens of sites replicating across entirely distinct data centers with higher latency connections.

Implementation:

- **Asynchronous Replication:** Metadata is replicated at configurable intervals, typically measured in seconds.
- **Active-Active Global Namespace:** All sites maintain an active namespace for continuous availability.

Benefits:

- **Minimize Risk:** RPO measured in seconds.
- **Zero RTO:** Shared namespace remains available even during replication.
- **Site Independence:** Works over high-latency or even intermittent connections. Sites that become disconnected (by design or by accident) continue to function normally. Metadata synchronization "catches up" when site connectivity returns. Any file change collisions are disambiguated using versions.

Use Cases: Ideal for geographically distributed environments requiring high availability across different data centers or regions, including global operations.

Redundancy Levels: Supports multiple replication targets, allowing for 2-way, 3-way, or N-way redundancy across sites.

Performance Consideration: While overhead is minimal, high latency or a large number of replication targets can reduce performance. The main impact to clients is when they are accessing data that lives on storage at another site, and then only the first time.

METHOD 4

Periodic Checkpoints of the Metadata Database

Description:

The system creates periodic checkpoints of the metadata database, which are stored as backups for restoration purposes. Multiple checkpoints can be maintained to provide multiple recovery points.

Implementation:

- **Checkpointing:** Metadata is periodically saved at configurable intervals (e.g., hourly, daily).
- **Storage:** Checkpoints are stored on NAS or object storage, potentially offsite for added security.
- **Anvil Pre-Installation:** Anvil can be pre-installed and ready to resume from a checkpoint, reducing recovery time.

Benefits:

- **Disaster Recovery:** Provides a fallback option in case of catastrophic failures.
- **Minimal Overhead:** Negligible impact on system performance during checkpointing.

Use Cases: Recommended for all installations as a secondary layer of protection against catastrophic failures.

Redundancy Levels: Multiple checkpoints (e.g., hourly, daily, weekly) can be maintained for increased protection.

Performance Consideration: Restoring from a backup will take longer with larger datasets, potentially multiple hours depending on the volume of metadata.

METHOD 5

Export of Metadata to JSON for Filesystem Reconstitution

Description:

Metadata can be exported into a JSON format, enabling the filesystem to be reconstituted independently of the Hammerspace platform. Multiple exports can be maintained for N-way redundancy over time.

Implementation:

- **Metadata Export:** Dumps metadata into an independent format like JSON or a tar file containing symbolic links referencing data instances.
- **Reconstitution:** Tools like Python scripts can reconstruct the filesystem from the exported metadata.

Benefits:

- **Data Portability:** Allows data to be moved or recovered without relying on the original metadata servers.
- **Disaster Recovery:** Provides an additional layer of protection through direct data de-assimilation.

Use Cases: Useful for long-term archiving, migration, or scenarios where data needs to be recovered without Hammerspace.

Redundancy Levels: Multiple exports can be scheduled and maintained, offering N-way redundancy over time.

Performance Consideration: The process of reconstituting the filesystem will take longer with larger datasets.

Conclusion: A Paradigm Shift in Data Protection

Just as the advent of cloud computing initially stirred concerns over security – only to later prove significantly more secure than traditional corporate IT infrastructure – Hammerspace is redefining the standards of data reliability and protection. At first glance, data platforms like Hammerspace might seem riskier than traditional Network Attached Storage (NAS) solutions due to their distributed architecture and abstracted layers. However, this perception couldn't be further from the truth.

Traditional NAS systems, while familiar, are limited by their hardware-centric design and lack the flexibility to adapt to modern data demands. They require extensive management of individual files and storage devices, making scalability and robust protection challenging. In contrast, the Hammerspace metadata-centric approach puts intelligent metadata at the helm of data management.

This shift transforms how data is protected:

- **Metadata-Driven Governance:** Metadata acts as the brain of the operation, governing data protection through automated policies and Service Level Objectives (SLOs). It orchestrates data placement, access controls, and replication – even across heterogeneous storage systems, multiple sites, and clouds - with precision and efficiency unattainable in traditional NAS systems.
- **Efficient Scaling:** Managing billions of files individually is resource-intensive and prone to error. Data Orchestration abstracts this complexity by handling billions of metadata records within a database optimized for performance and scalability. This ensures that data protection mechanisms remain robust, even as data volumes grow exponentially.
- **Advanced Protection Mechanisms:** With metadata in control, the platform can employ sophisticated data protection strategies, such as global namespace management, automated failover, and intelligent data tiering, which are difficult to implement effectively in traditional NAS environments.

By combining these advanced metadata protection strategies with the intelligent governance of data through metadata, the Hammerspace Data Platform achieves a level of reliability that traditional NAS solutions simply cannot match.