



2.3 mSAN Deep Dive

Micron21 - 2025-03-18 - Section 2: Storage

Introduction

In today's fast-paced digital landscape, businesses require storage solutions that are not only high-performing but also incredibly reliable and resilient. Micron21's mCloud platform leverages advanced technologies to provide a 3N clustered Ceph storage-as-a-service solution. This innovative architecture ensures that data is replicated in real-time with minimal latency, creating three copies of your data across three Melbourne-based data centres. Designed for clients who demand real-time replication across different physical geographical regions, our Ceph storage solution offers unparalleled performance, reliability, and scalability.

Ceph Storage: The Backbone of Our Storage Solution

Ceph is an open-source, distributed storage system that provides excellent performance, reliability, and scalability. It unifies object, block, and file storage into a single platform, making it ideal for cloud environments. Ceph's architecture is built on the Reliable Autonomic Distributed Object Store (RADOS), which ensures data is automatically replicated, balanced, and recovered across the storage cluster.

Key Features of Ceph:

Scalability: Seamlessly scales from a few nodes to thousands.

High Availability: Eliminates single points of failure through replication and self-healing capabilities.

Flexibility: Supports object, block, and file storage in one unified system.

Performance: Optimized for both high-throughput and low-latency workloads.

3N Replication: Triple Redundancy for Maximum Data Protection

Our Ceph storage solution employs a 3N replication strategy, where N represents the number of copies of data stored. In this case, data is replicated three times across three independent storage nodes located in separate data centres:

Primary Data Centre: Micron21 Kilsyth

Secondary Data Centre: Melbourne CBD

Tertiary Data Centre: Port Melbourne

Benefits of 3N Replication:

Data Durability: Protects against data loss even in the event of multiple node failures.

High Availability: Ensures continuous access to data during maintenance or unexpected outages.

Consistency: Provides strong data consistency across all replicas.

When data is written to the cluster, Ceph synchronously replicates it to all three nodes before acknowledging the write operation. This guarantees that all copies are up-to-date, and clients always read the most recent data.

High-Performance NVMe Storage

Each of our Ceph nodes is equipped with enterprise-grade NVMe SSDs, delivering exceptional performance:

Low Latency: NVMe drives communicate directly over the PCIe bus, reducing latency compared to traditional SATA or SAS SSDs.

High Throughput: Capable of handling intensive read/write operations, making them ideal for high-performance applications.

Enhanced Endurance: Designed for enterprise workloads with high write-intensive operations.

By utilizing NVMe storage, we ensure that data operations are executed swiftly, which is crucial for applications that require real-time data access.

Dedicated High-Speed Networking Infrastructure

Our Ceph storage nodes are interconnected via a dedicated, low-latency 100 Gbps Cisco Nexus 9000 (9k) leaf and spine network. This high-performance network is exclusively used for Ceph storage traffic, separate from our transit and management networks.

Network Features:

Leaf and Spine Architecture: Provides multiple pathways for data, reducing latency and eliminating bottlenecks.

100 Gbps Connectivity: Ensures high bandwidth for data replication and client access.

Low Latency: Critical for synchronous replication across data centres.

Isolation: Dedicated network prevents interference from other traffic types.

By isolating the storage network, we maintain optimal performance and security for data operations.

Geographically Dispersed Data Centres

Our three data centres are geographically separated by at least 45 km, providing true physical isolation:

Micron21 Kilsyth Data Centre: Primary site with state-of-the-art infrastructure.

Melbourne CBD Data Centre: Secondary site in the heart of the city.

Port Melbourne Data Centre: Tertiary site offering additional redundancy.

Advantages of Geographical Separation:

Disaster Recovery: Protects against regional disasters like floods, fires, or power outages.

Independent Power Grids: Each data centre is connected to different power sources, reducing the risk of simultaneous outages.

Separate Cooling and Network Systems: Minimizes the risk of a single point of failure affecting all sites.

This setup ensures that even in the unlikely event of two data centres failing, your data remains safe and accessible from the third.

Dark Fibre Connectivity: Spanning Data Centres

To achieve real-time replication with minimal latency, our data centres are interconnected using dark fibre technology:

Dedicated Fibre Optic Cables: Provides exclusive use of the fibre strands for our network traffic.

High Bandwidth and Low Latency: Supports 100 Gbps speeds necessary for synchronous data replication.

Secure Communication: Reduces the risk of interception or eavesdropping, enhancing data security.

Dark fibre allows us to control the network infrastructure completely, optimizing it for Ceph's replication requirements.

Seamless Integration with OpenStack mCloud Platform

Our Ceph storage solution is fully integrated with the OpenStack-based mCloud platform, enabling advanced cloud functionalities:

Virtual Machine Mobility: Deploy or migrate virtual machines across any of the three data centres instantly.

High Availability (HA): With real-time data replication, virtual machines can failover seamlessly between sites.

Scalable Resources: Dynamically allocate compute and storage resources as your needs evolve.

Benefits of Integration:

Unified Management: Manage compute, storage, and networking resources from a single interface.

Automation: Utilize OpenStack's orchestration tools for automated deployment and scaling.

Flexibility: Support for various workloads, including cloud-native applications and legacy systems.

Ceph Clustering Mechanics: How It Works

Understanding the technical aspects of Ceph's clustering provides insight into its robustness:

RADOS: The Core of Ceph

Object Storage Daemons (OSDs): Handle data storage, replication, recovery, and rebalancing.

Monitors (MONs): Maintain cluster state, configuration, and consensus for distributed decision-making.

Placement Groups (PGs): Logical partitions that distribute data across OSDs

for load balancing.

CRUSH Algorithm

Ceph uses the Controlled Replication Under Scalable Hashing (CRUSH) algorithm to determine data placement:

Data Distribution: CRUSH calculates where data should reside, eliminating the need for a central lookup table.

Scalability: Supports cluster growth without significant performance impact.

Fault Tolerance: Automatically adjusts data placement in response to node failures or additions.

High Availability Through Real-Time Replication

The combination of Ceph's architecture and our networking infrastructure delivers unparalleled HA:

Automatic Failover: If an OSD or entire data centre becomes unavailable, Ceph reroutes operations to available replicas without manual intervention.

Consistent Performance: Load is evenly distributed, ensuring that the failure of a node doesn't degrade performance.

Self-Healing: Ceph automatically detects and repairs inconsistencies, maintaining data integrity.

Ideal for Mission-Critical, Always-On Services

Our 3N clustered Ceph storage is designed for organizations that cannot afford downtime:

Use Cases:

Financial Institutions: Require zero data loss and uninterrupted access for transactions.

Healthcare Providers: Need secure, always-available access to patient records and imaging.

E-commerce Platforms: Downtime directly impacts revenue and customer trust.

Government Agencies: Demand high security and resilience for critical operations.

Conclusion

Micron21's 3N clustered Ceph storage-as-a-service offers an exceptional combination of performance, reliability, and resilience. By leveraging cutting-edge technologies like NVMe storage, a dedicated 100 Gbps Cisco 9k leaf and spine network, and dark fibre connectivity, we provide a storage solution that meets the most demanding requirements.

Our integration with the OpenStack mCloud platform further enhances flexibility and scalability, allowing clients to deploy and manage resources efficiently across multiple data centres. With real-time replication and geographical redundancy, your data is protected against virtually any scenario, ensuring that your mission-critical services remain always-on.

Experience unparalleled storage performance and resilience with Micron21's mCloud Ceph storage solution. Contact us today to learn how we can support your business's critical infrastructure needs.