Transactions of the Korean Nuclear Society Autumn Meeting
Changwon, Korea, October 30-31, 2025

Enhancement to the Redundancy Design for High Availability in the KOMAC Archiving System

Jae-Ha Kim *, Young-Gi Song, Sung-yun Cho, Jung Hae-Seong,
Hyeok-Jung Kwon
Korea Multi-purpose Accelerator Complex, Korea Atomic Energy Research Institute, Gyeongju 38180, Korea
*Corresponding author:jhkim1@kaeri.re.

Keywords : Archiving system, redundancy

1. Introduction

KOMAC operates a 100 MeV proton linac providing proton beams to several target rooms and controls it with an EPICS-based distributed control system. On this system, an operational archiving system is established by the Archiver Appliance and engineered for high availability: two archiver nodes operate in a primary/secondary topology coordinated by the Master Control and Coordination System (MCCS). MCCS monitors the primary via heartbeats and health checks and, on fault, transfers the virtual IP and promotes the secondary. To keep state consistent at switchover, both nodes write to a network file store exported via NFS, and all archiver metadata, including PV and channel definitions and engine configuration, is stored in a single shared MariaDB referenced by both nodes. This consolidation removes duplicated configuration effort, prevents file-handle or schema drift, and preserves a continuous time series through failover. This paper describes the implementation of the MCCS-based redundant archiver system with NFS-backed storage and centralized MariaDB, and its integration into the KOMAC EPICS control environment.

2. KOMAC Data Archiving System

At KOMAC, an EPICS-based data archiving system built on the Archiver Appliance collects PVs from IOCs over Channel Access. The deployment uses clustered engine and management nodes with load balancing to distribute PVs and improve scalability and stability. Operators use Phoebus CSS and an internal web application for data search, monitoring, and plotting. To minimize data loss during faults, the service runs in an active–standby configuration coordinated by MCCS. It monitors health and, on failure, automatically reassigns a virtual IP and performs ordered service restarts. Both nodes write to a shared NFS file store and reference a centralized MariaDB, ensuring consistent metadata and failover with minimal interruption, suitable for accelerator operations. Figure 1 shows the architecture of the KOMAC archiving system.

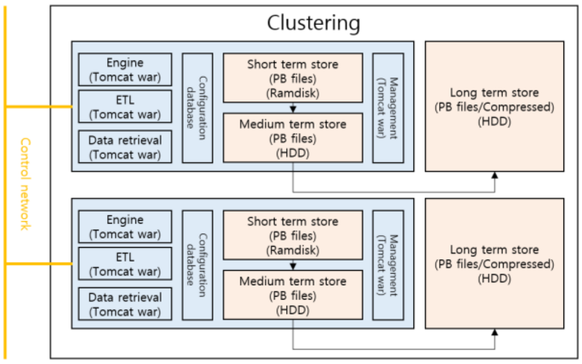

Fig. 1. The architecture of the KOMAC archiving system.

2.1 Redundant deployment

High availability is implemented with MCCS, with the archiving service running on primary and secondary nodes. MCCS continuously monitors network links, processes, services, application health, the operating system, and disk status. When a fault is detected, it transfers the virtual IP, runs the registered start/stop/status scripts, and promotes the secondary node. The application layer hosts the Archiver Appliance. The database layer is centralized on the data storage server using MariaDB. The data area is isolated on the storage server with a rebuilt RAID layout, and both nodes mount the same path via NFS. The inactive node does not write to the archive directory until it is promoted, which preserves data consistency and enables failover with minimal interruption. Figure 2 shows the redundant configuration of the KOMAC data archiving system.

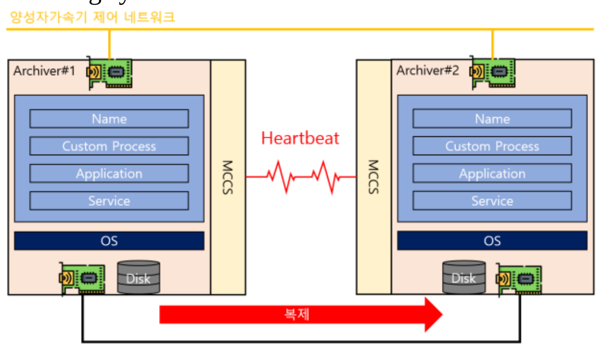

Fig. 2. The redundant configuration of the KOMAC data archiving system.

2.2 Failover Control and Operations

Cluster operations are managed through the MCCS GUI, which lists the managed resources such as the VIP, NFS mounts, archiver engines, database connections, and monitoring services, and shows their Active–Standby status, as shown in Figure 3.

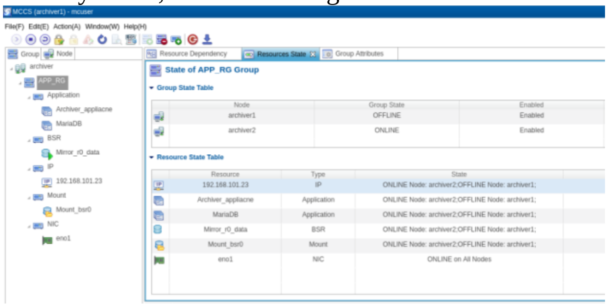

Fig. 3. MCCS GUI for redundant archiving.

Clients access the archiving service via the VIP, so connectivity is unchanged regardless of whether Archiver1 or Archiver2 is active. Using the GUI, a manual failover can be triggered to validate behavior under controlled conditions. After switchover, VIP-based access is checked to confirm normal operation. We also perform fault injection by forcibly terminating the archiver process to verify service switchover and examine the recorded data for gaps in Δt to assess potential loss, as shown in Figure 4.

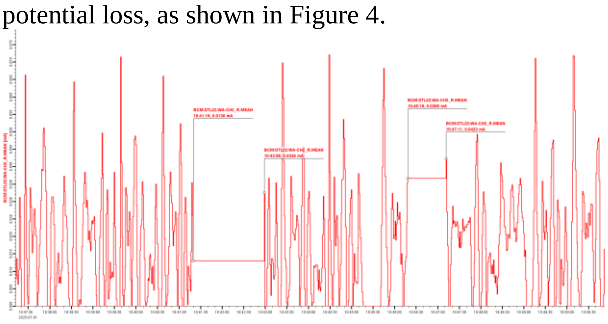

Fig. 4. Failover test procedure and data-gap analysis.

3. Conclusions

We implemented a redundant data archiving system coordinated by MCCS and confirmed correct failover through manual switchover and fault-condition tests. VIP-based access remained available after switchover, and data recording continuity was maintained. For future work, we will move from local storage to an external SAN and update mount and path policies, access control, monitoring metrics, and switchover scripts. The goal is to support storage high availability and site-level disaster recovery..

REFERENCES

[1] EPICS Archiver Appliance Documentation (Read the Docs), accessed Aug. 21, 2025.
[2] Mantech, “MCCS (Mantech Continuous Cluster Server),” product/overview page, accessed Aug. 21, 2025