



Research on HALF Historical Data Archiver Technology

XiaoKang Sun, DaDi Zhang, HuHang Chen
University of Science and Technology of China, HeFei, China

Abstract

The Hefei Advanced Light Facility (HALF) is a 2.2GeV fourth-generation synchrotron radiation light source. The HALF control system is a distributed control system based on Experimental Physics and Industrial Control System (EPICS). The Historical Data Archiving System (HDAS) of HALF is responsible to store operational data for the entire facility. Under the EPICS PV data scenario of HALF, a fair database test platform is designed and built to test the read-write performance of databases commonly used in the particle accelerator field. The tested objects include EPICS Archiver Appliance and the five databases MongoDB, HBase, InfluxDB, TimescaleDB, and Cassandra. The test results indicate that TimescaleDB has the fastest read performance, and 1.4×10^6 items of data can be read per second. In the future, a TimescaleDB distributed cluster will be designed and deployed, and an HDAS prototype system will be developed based on this cluster.

Introduction

➤ According to the estimation based on the HALF scale, approximately 25,000 PVs will be stored in HDAS, which will generate tens of TB of data per year.

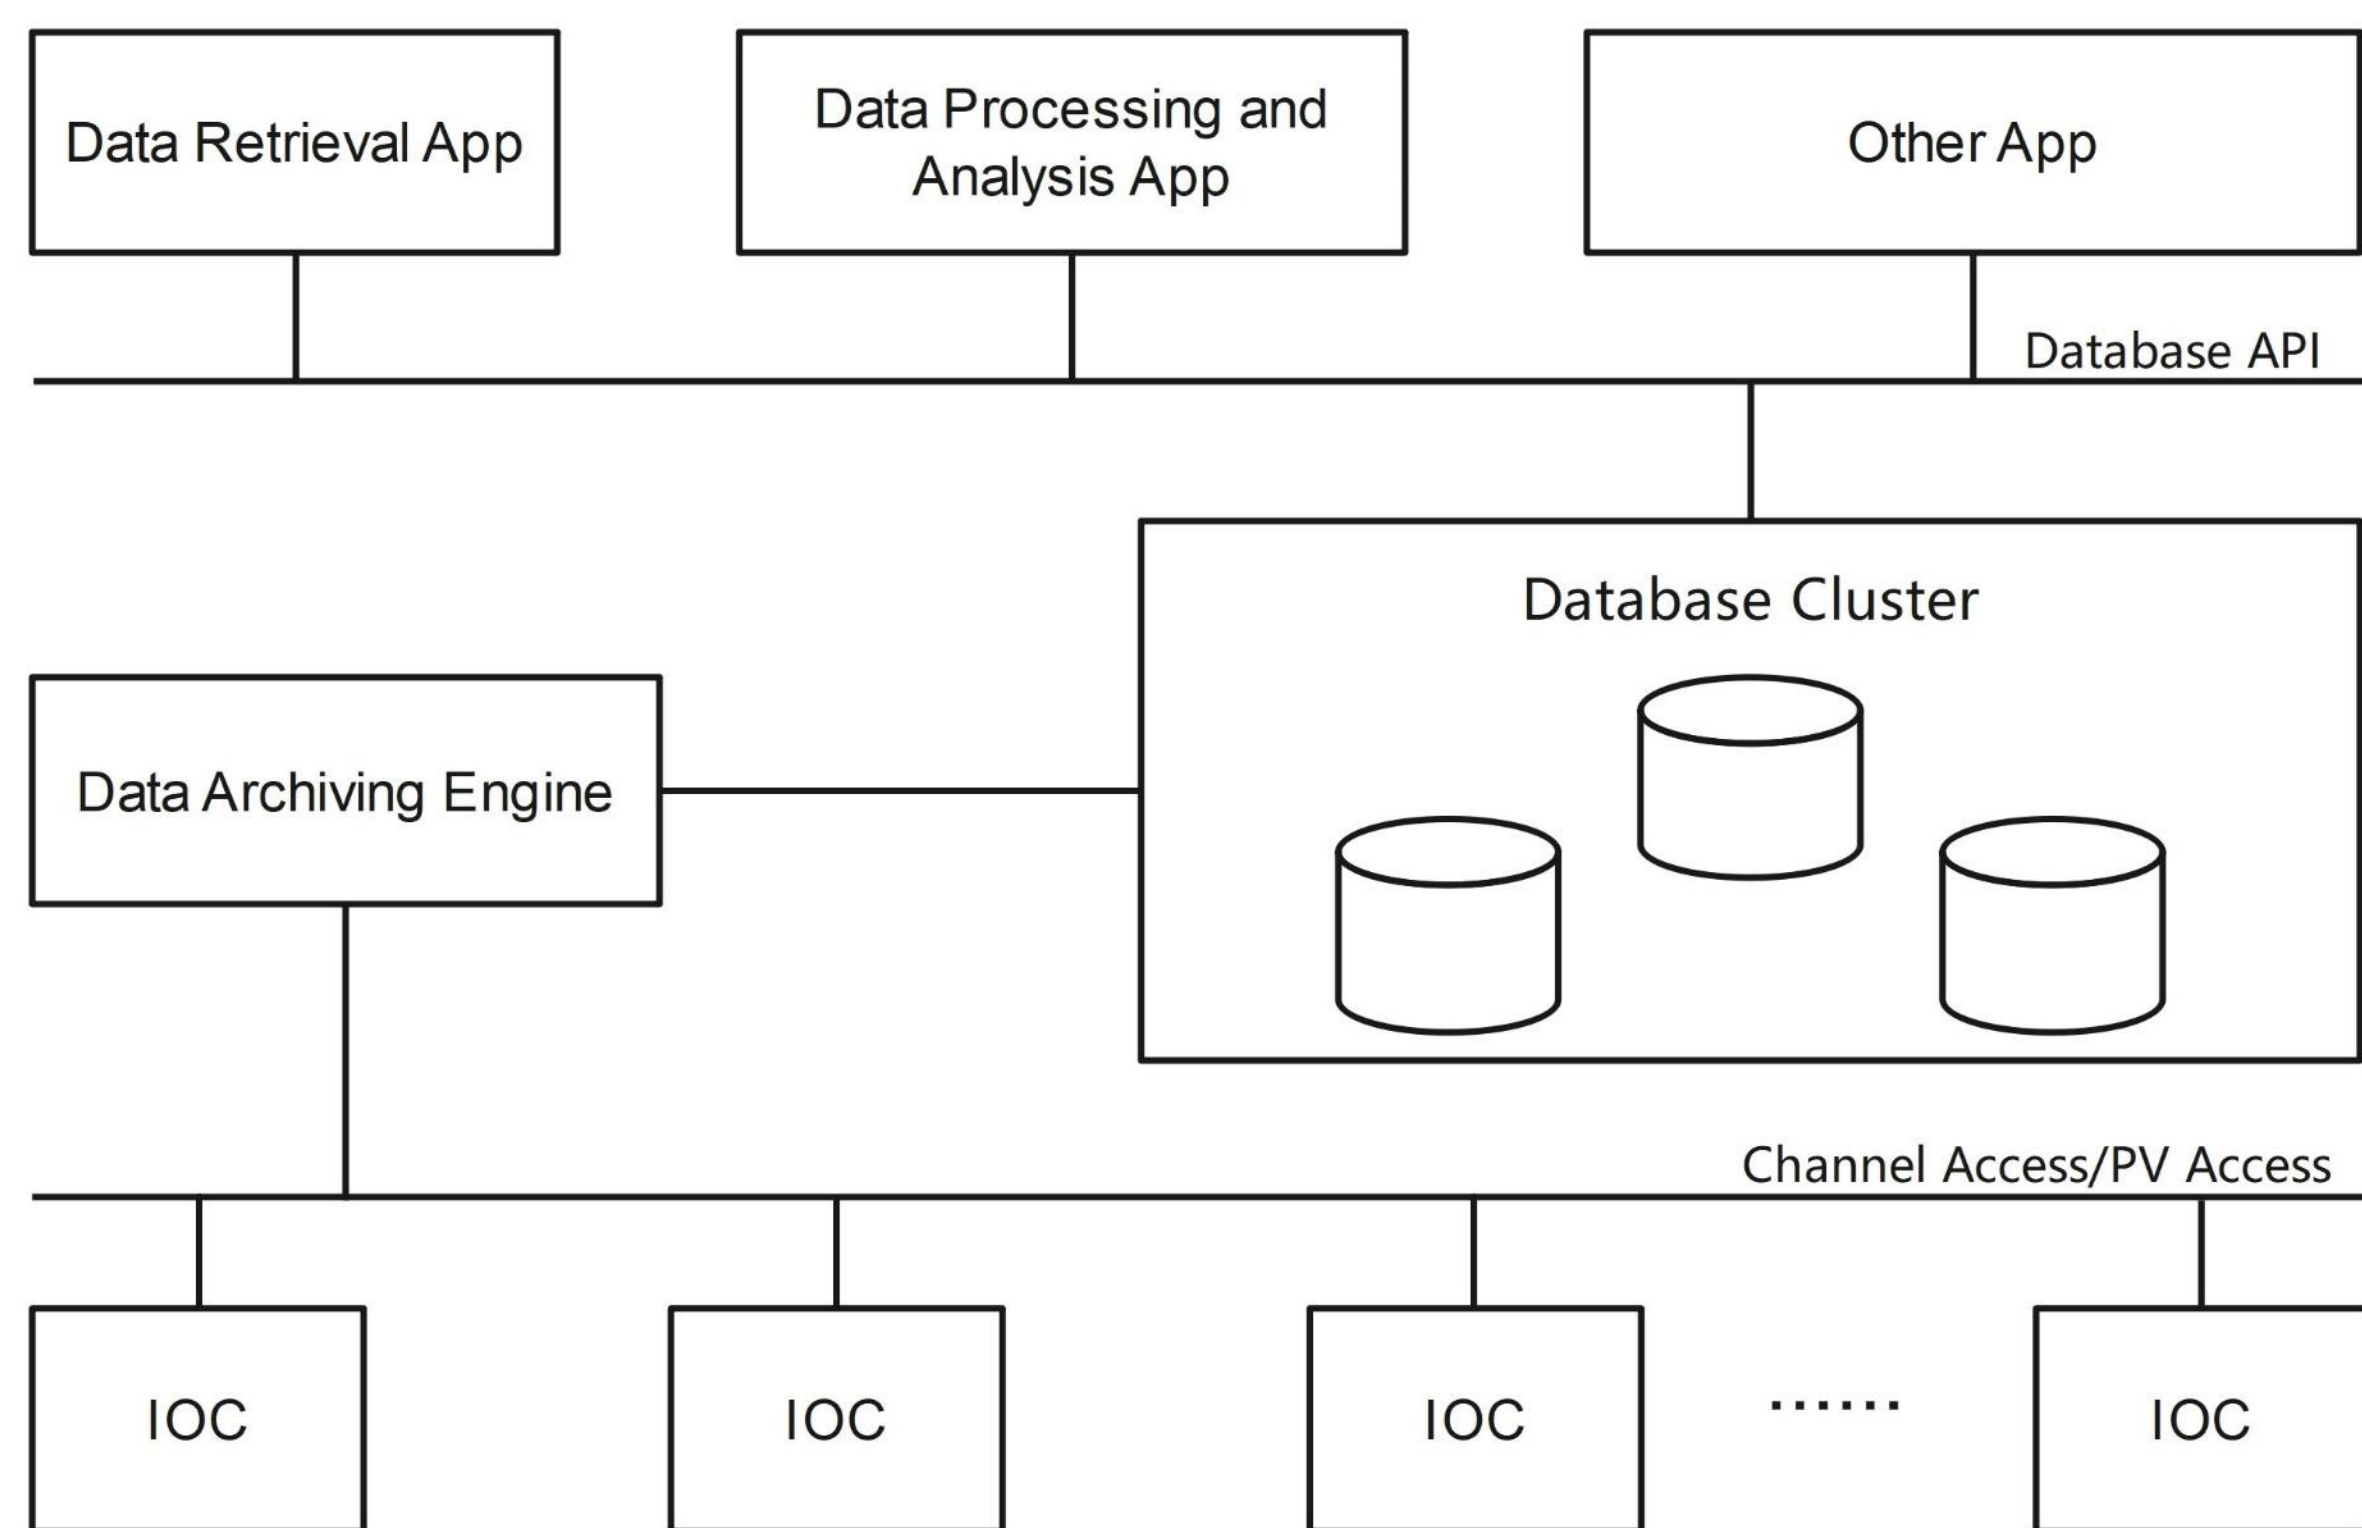


Fig 1: Architecture of the HDAS

- The popular and commonly used databases in the particle accelerator facilities: MongoDB, HBase, InfluxDB, TimescaleDB, and Cassandra.
- There is currently a lack of complete testing and comparison of these databases in the EPICS data archiving scenario
- The read-write performance tests are carried out for the above databases.

Experiment Design and Development

➤ The performance test metric

- QDPS: the volume of data queried per second
- IDPS: Inserted data per second

➤ Hardware platform design

- A fair environment test platform
- The experimental platform consists of eight virtual machines (VM), all deployed within the same local area network

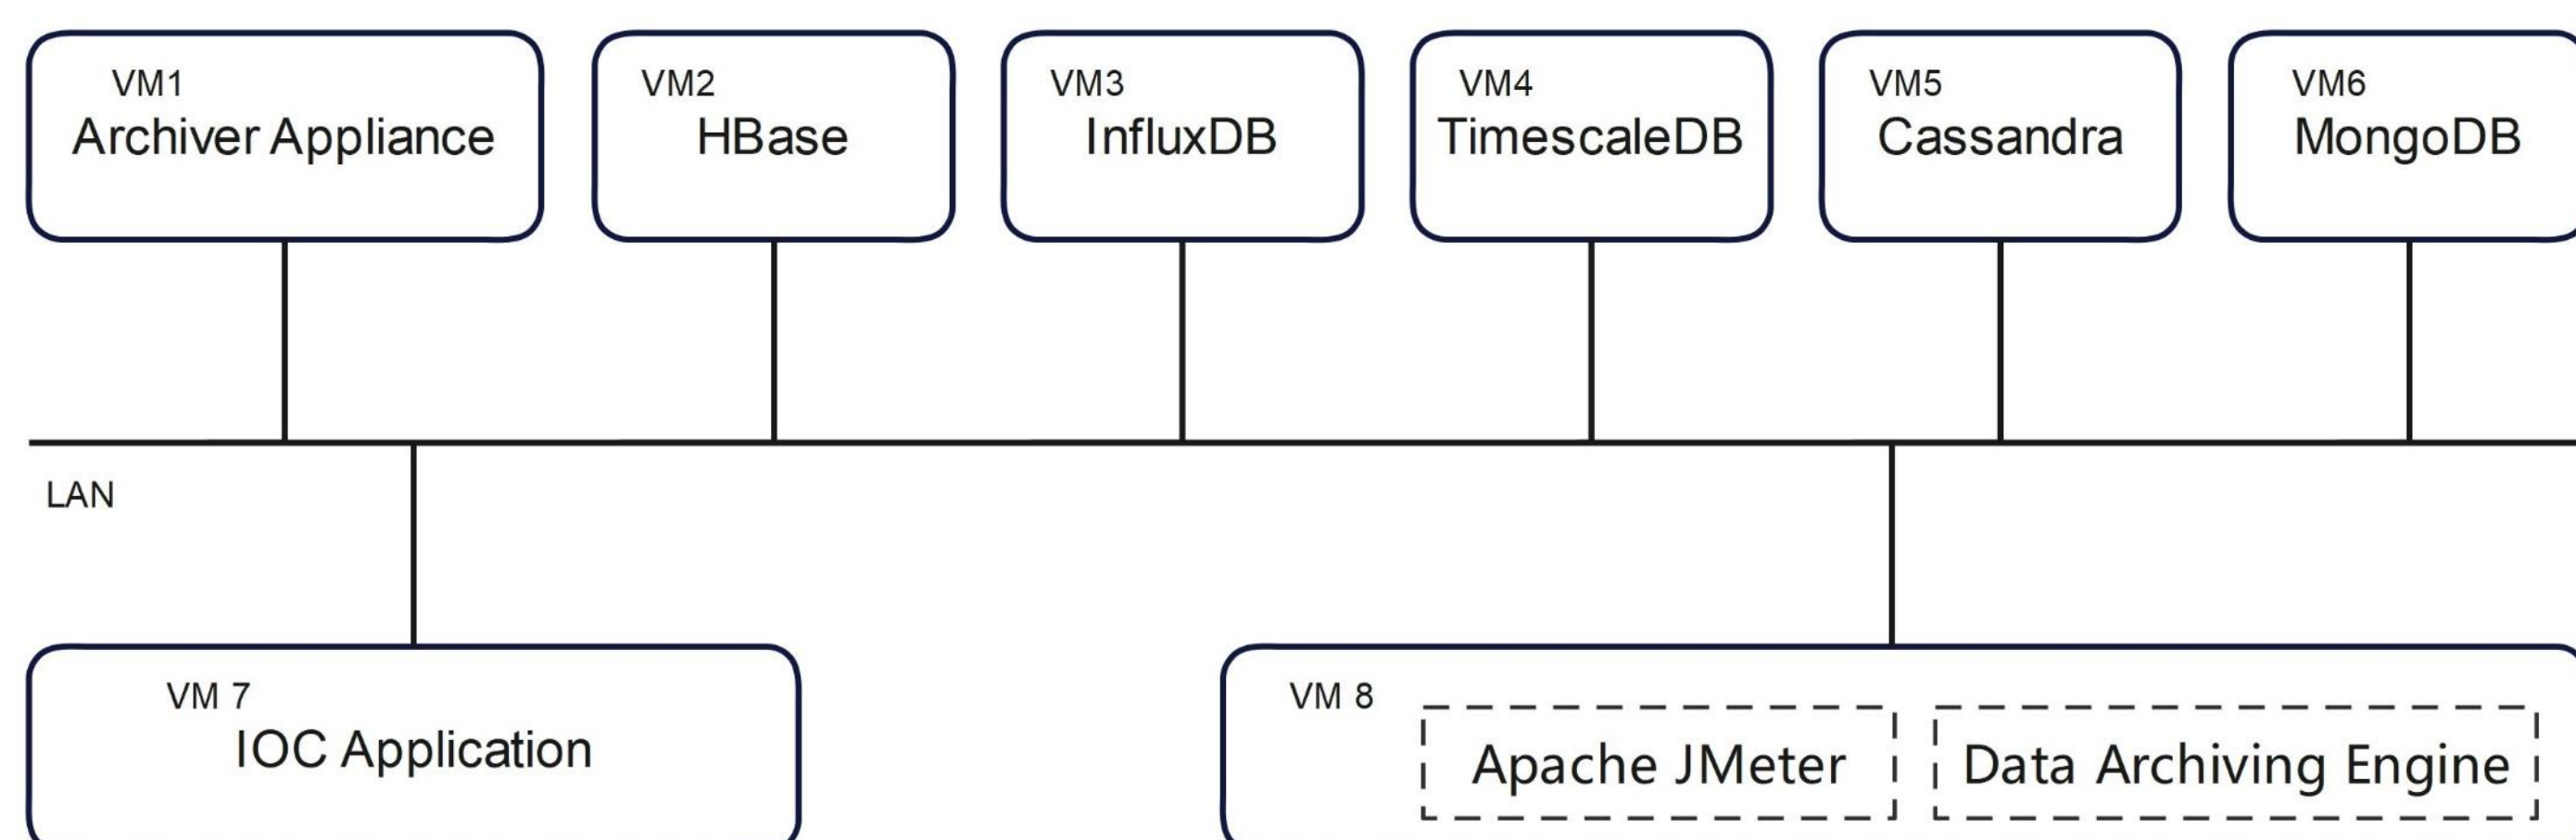


Fig 1: Database performance testing experimental platform

➤ Software development

- IOC Application: a total of 2,000 PVs, each PV is updated at a rate of 10Hz to generate about 4 days of data, with a total of about 7 billion items of PV data
- Databases: PV_id, timestamp(nanos), value, severity, status

Experiments and result analysis

➤ Reading test:

- The data volume generated by a single PV within the time range of 0.4 hours (1.44×10^4 items) to 10 hours (3.6×10^5 items) is designed to do this test.
- The highest QDPS is achieved by TimescaleDB, which can reach 1.4×10^6 .




Fig3: Test results of Reading Test

➤ Writing test:

- The accumulated data volume of 2000 PVs in the range of 0.1 minutes to 1 minute for batch insertion.
- The InfluxDB has the highest insertion speed, whose IDPS can reach 1.5×10^5 . TimescaleDB also has good write performance, and its IDPS can be stabilized at 1.2×10^5 .




Fig4: Test results of Writing Test

Conclusion

- TimescaleDB has best the read performance, and its IDPS can fully meet the HDAS writing requirements. Considering HDAS's "once-write, multiple-reads" scenario, QDPS is a more important parameter than IDPS. Therefore, TimescaleDB emerges as the most suitable database for HDAS.
- Timescale provides the Continuous aggregates function which can further improve the query speed.
- In the future, a TimescaleDB distributed cluster will be designed and deployed, and an HDAS prototype system will be developed based on this cluster