

Database Archiving System for Supervision Systems at CERN: a Successful Upgrade Story

Motivation

WinCC OA File Archiver:

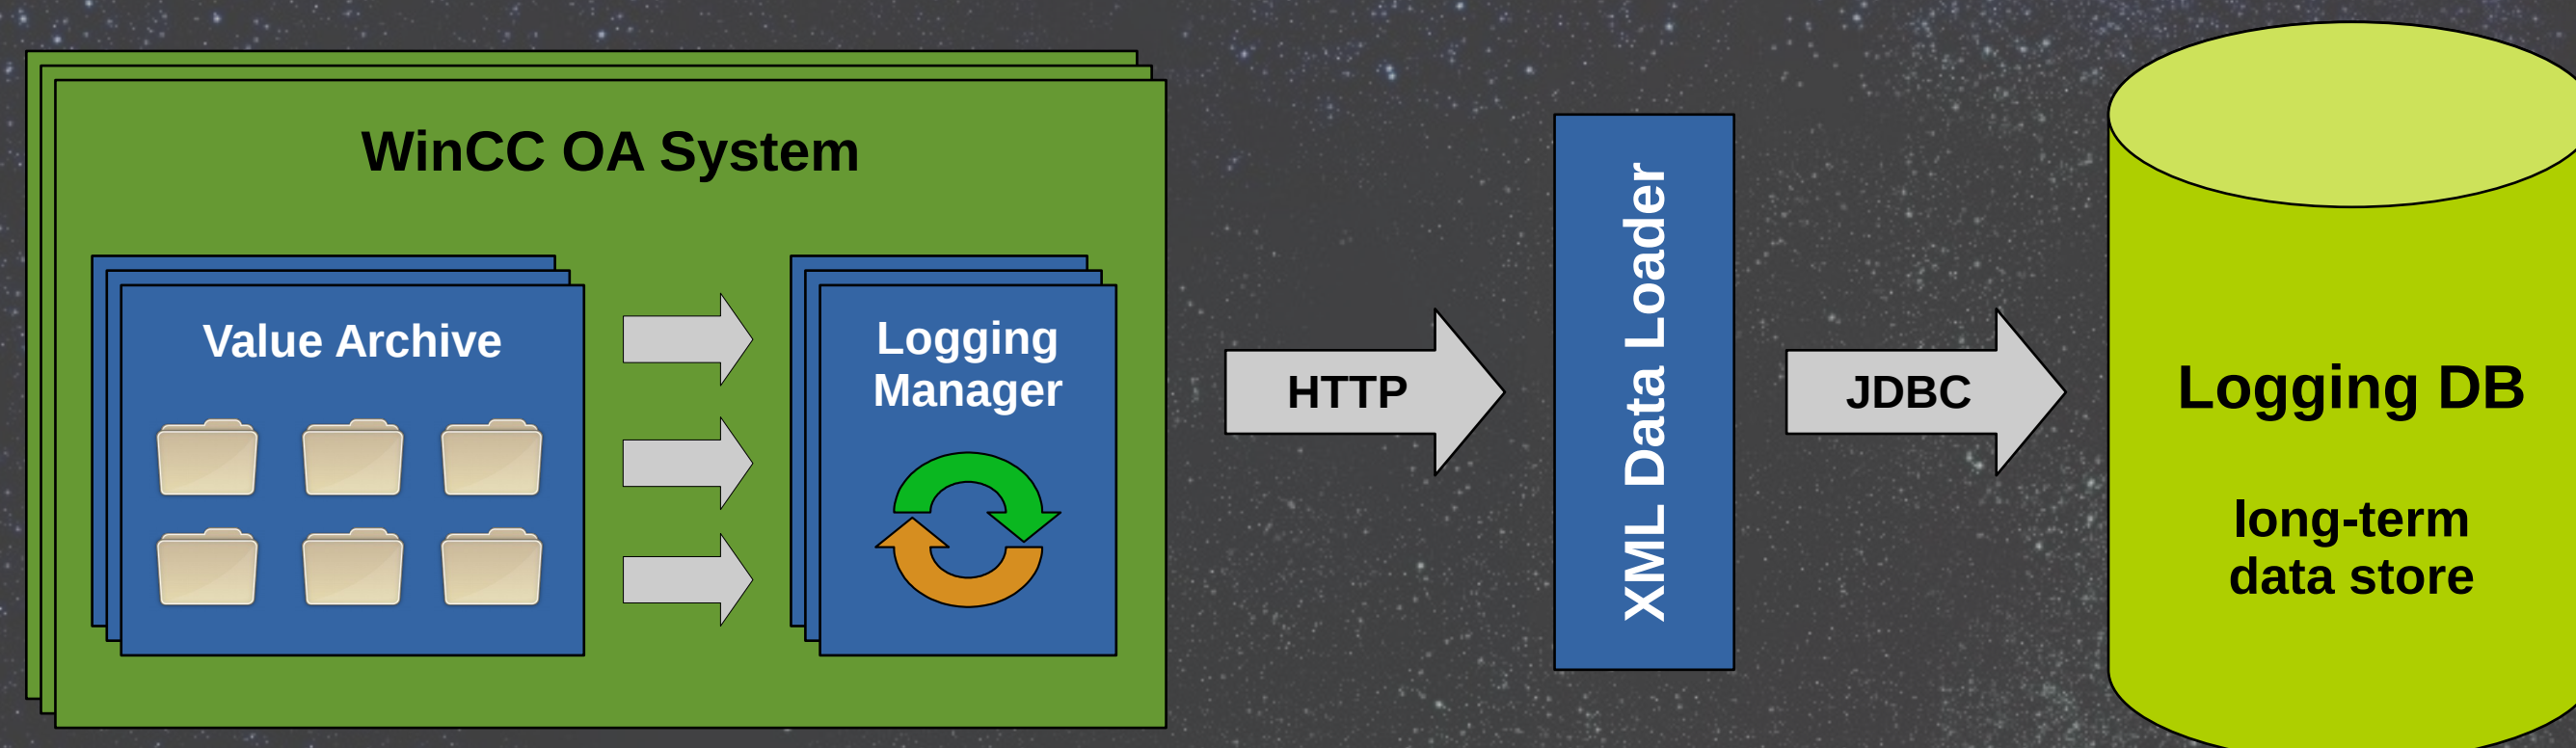
- * Used by 200 WinCC OA systems for LHC and technical infrastructure
- + Stable, proven technology
- Performance scalability issues
- Complicated maintenance on many SCADA nodes

WinCC OA Oracle RDB Archiver:

- + Already used in a few hundred controls systems at CERN
- + Proven scalability and performance for large systems
- + Centralized architecture eases management and handling of the SCADA nodes, provides better tools, and enables data analytics

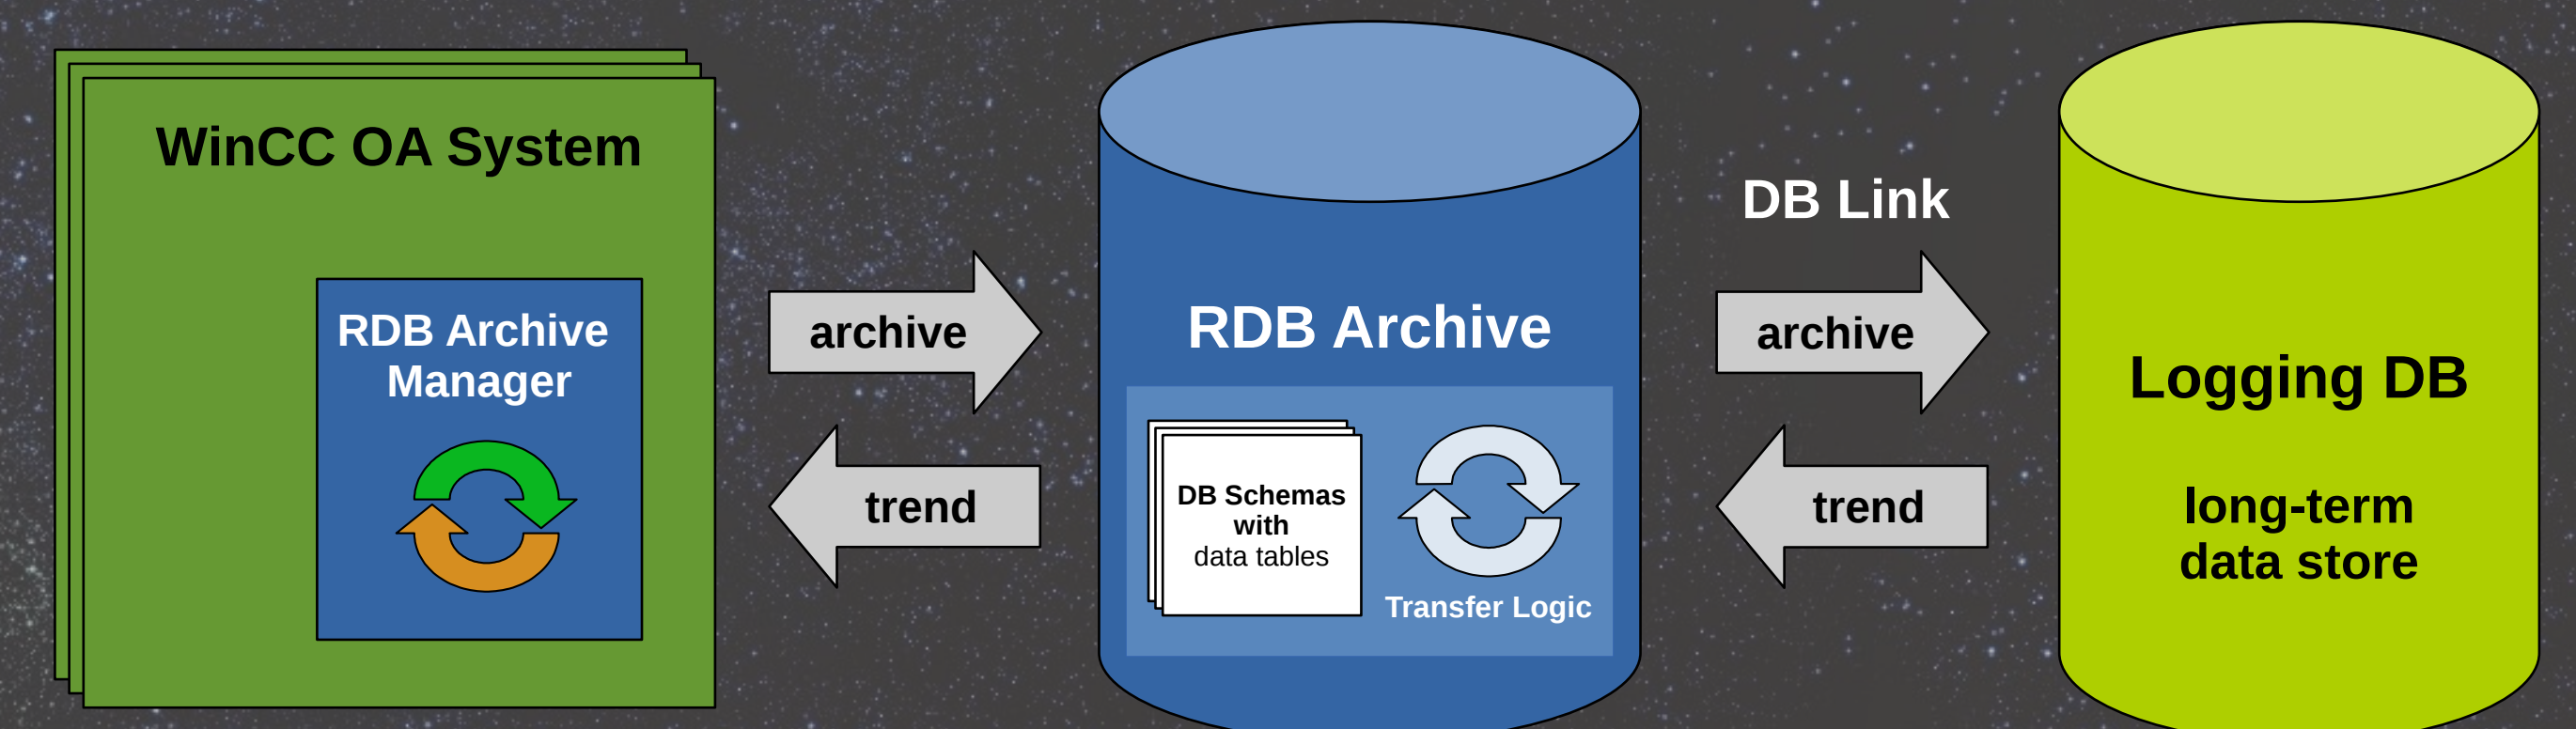
New requirement for the 13 TeV run of the LHC: the magnet protection system (QPS) needs data archived at a rate of 200.000 values/s

Archiving architecture prior to re-design




- Several WinCC OA Value Archive managers write historical data to local disk files
- Custom WinCCOA LoggingManager programs query data, transform to XML and send to a dedicated data loader application of the LoggingDB service
- The architecture would not scale to meet the new requirements for LHC Run II

Archiving architecture after re-design



- Single high-performance RDB Archive Manager writes data to central Oracle DB
- Data requiring long-term storage is transferred over a Dblink to the LoggingDB
- Standard database transfer job mechanism, used by other LoggingDB clients, serves all applications and provides diagnostic and statistical information
- Data stored in the LoggingDB can be made available in WinCC OA transparently

Index organized tables



- Key ingredient for optimization of readout path
- Access by index identified as the primary data-access pattern (transfers, trend plots)
- Query performance increased by a large factor
- Storage-space consumption reduced up to 50%

Row size reduction

Column	Type
ELEMENT_ID	NUMBER (25)
TS	TIMESTAMP(9)
VALUE_NUMBER	BINARY_DOUBLE
STATUS	NUMBER (20)
MANAGER	NUMBER (20)
USER	NUMBER (5)
SYS_ID	NUMBER (20)
BASE	NUMBER (1)
TEXT	VARCHAR2 (4000)
VALUE_TIMESTAMP	TIMESTAMP(9)
CORRVALUE_NUMBER	BINARY_DOUBLE
CORRVALUE_STRING	VARCHAR2 (4000)
CORRVALUE_TIMESTAMP	TIMESTAMP(9)

110 B (90B of data + 20 B indexes)
• Standard used schema: e.g. ATLAS

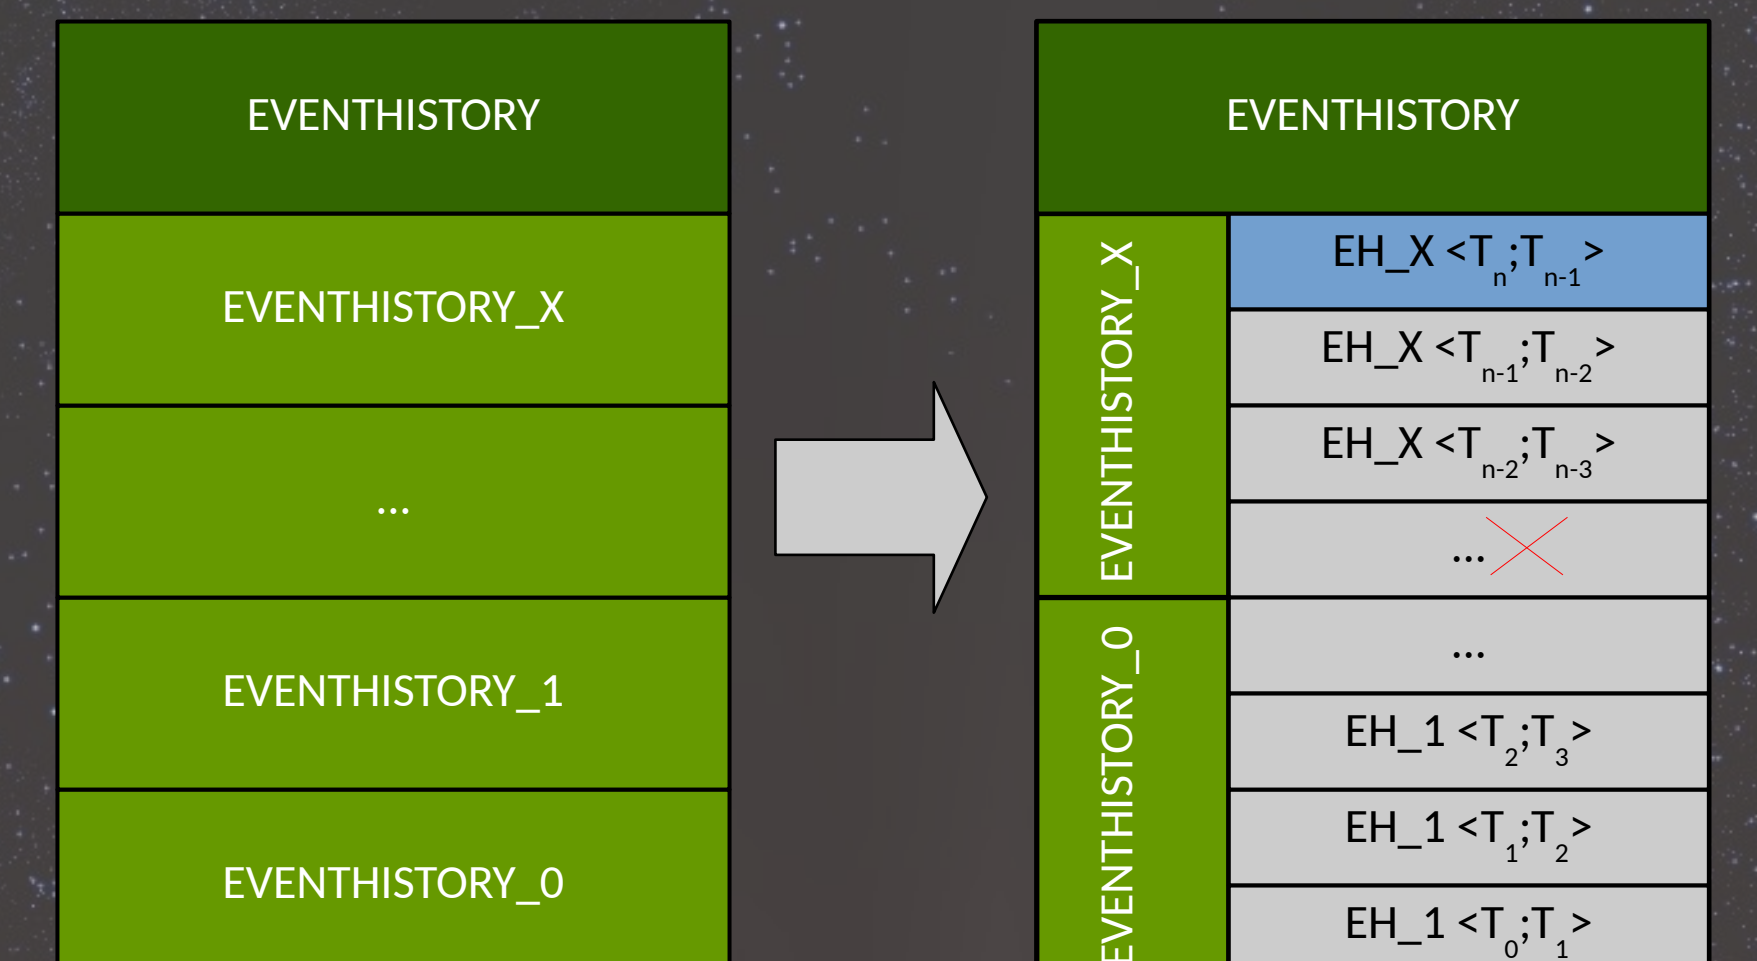
60 B
• Data and indexes merged together due to usage of IOT

30 B
• New schema with reduced column set, IOT and index compression
• Used in QPS

Typically NULL values
Highly redundant

- I/O bottleneck in the database due to Oracle's REDU log size (data-recovery stream)
- Option to eliminate redundant/unnecessary data

Time based partitioning



- Growing IOT leads to degraded query performance
- Time-based partitioning (partition pruning)
- Queries on views with pushed-predicate hints
- Data retention policy (3 months in daily partitions)
- Automated partition-managent job for 200 schemas

Data recovery performance test

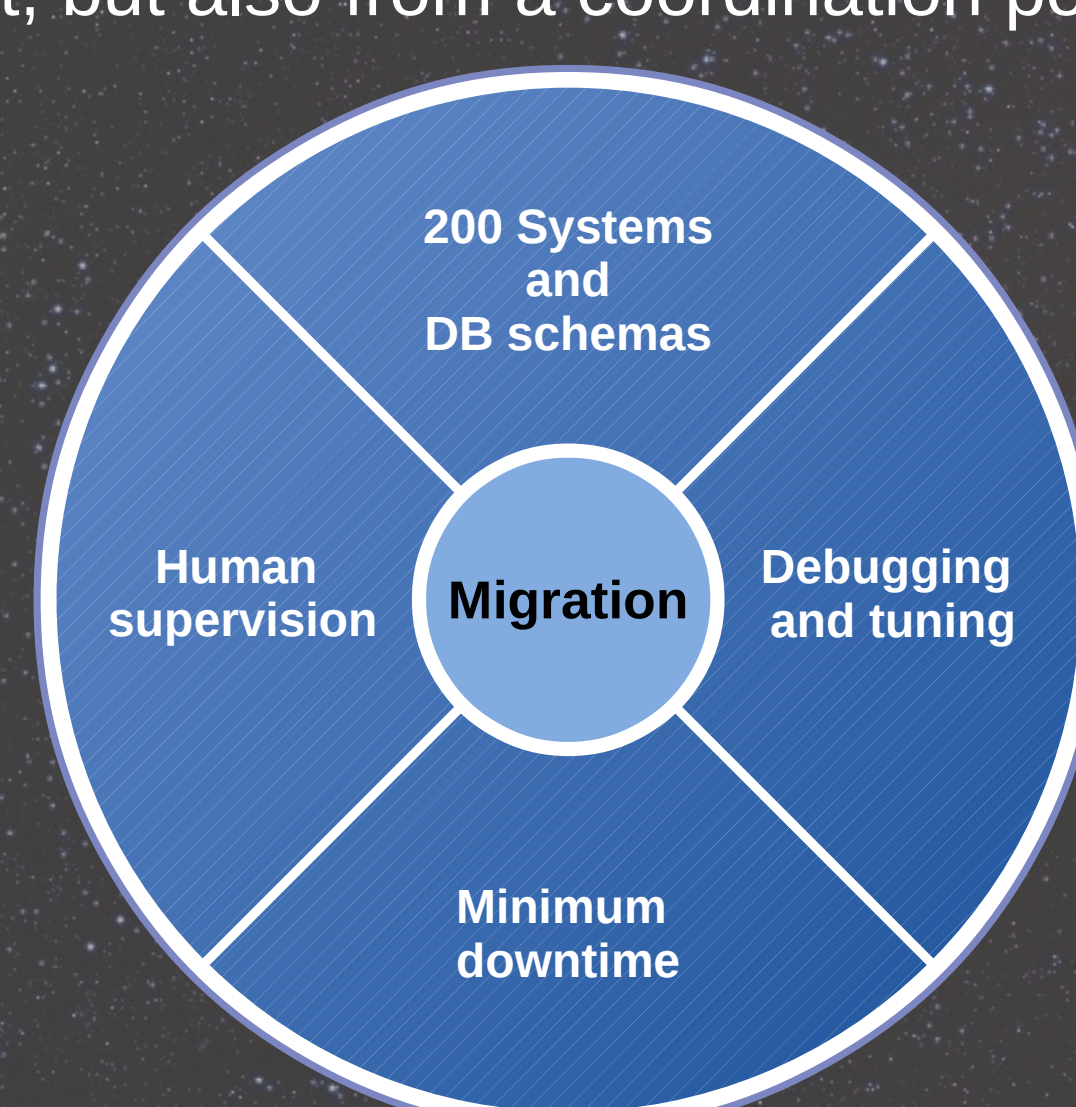
Simulate a database-outage scenario in a set-up mimicking the LHC magnet protection system (QPS)

Nominal Conditions	WinCC OA Projects	50
Database Failure	Archived Signals	150 000
Recovery	Data Input Rate [rows/s]	200 000
Nominal Conditions	Outage duration [h]	8
	Recovery Time [h]	2
	Peak Recovery Rate [rows/s]	1 000 000

- Data buffered by RDB Archivers during DB outage
- Buffers inserted at maximal rate once DB becomes available
- Even in degraded mode (1 DB node) complete recovery in acceptable time.

Migration

A challenging task, not only from a technical standpoint, but also from a coordination point of view.



Conclusion

- Migration completed in time for Run II of the LHC
- Single consistent archiving technology CERN-wide
- Satisfactory performance and reliability
- Optimisations available to all WinCC OA users

