SKA archiving

HDB++ Meeting 2019

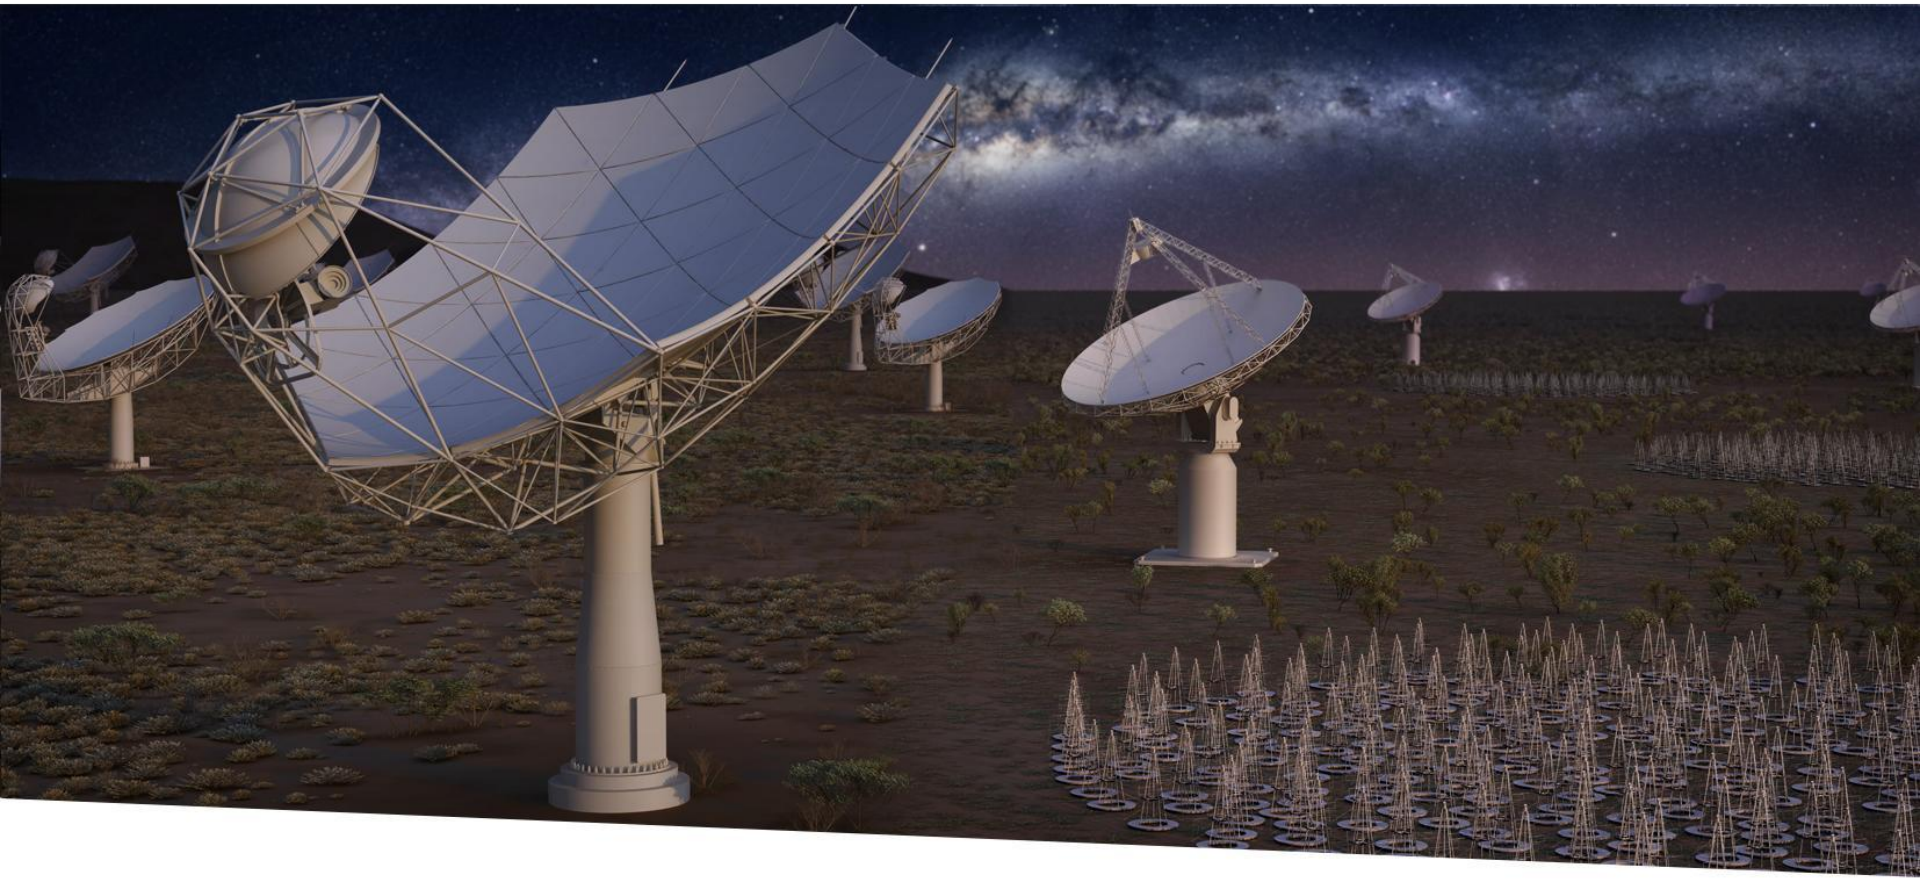

SQUARE KILOMETRE ARRAY

Exploring the Universe with the world's largest radio telescope

Mangesh Patil

18 September 2019

Agenda

- Current Status
- Engineering Data Archive Requirements
- Proposed Architecture
- Prototyping
- Findings
- Concerns
- Next Steps

Current Status

- SKA1-Mid and SKA1-Low - two telescopes, common architecture
- CDR completed for most sub-systems
- System CDR closure ~ mid 2020
- Bridging Activities in Progress - primarily aimed at:
  - Risk reduction
  - Building evolutionary prototype
  - Addressing CDR comments
- Construction
  - Start ~ late 2020 / early 2021
  - Overall 5 years, with early production arrays along the way

Engineering Data Archive Requirements

- Persist the engineering data collected and derived during the Telescope operations.
- Data extraction for performing diagnostics, troubleshooting, system monitoring for health, performance, availability, reliability and management
- Engineering Data comprises of –
  - Monitoring data
  - Alarms/Faults
  - External Data (Weather information, RFI, etc)
- Data retention period - 50 years
- Overall system availability 99.6%
- Common archiver architecture and solution for both telescopes

Data Estimates

- Considerations - Varying sampling frequency for monitoring points
  - 5 % monitoring points @ 250 ms (subsecond rate)
  - 25% monitoring points @ 1 to 5 second
  - 70% monitoring points @ 10 second
  - Data size per sample ~ 100 bytes (includes attribute identifier, value, 2 timestamps, etc).
- SKA1 MID Telescope
  - 64 Meerkat Dishes + 133 Dishes + 10 other Elements (CSP, SDP, etc) = 207 Elements
  - 1000 monitoring points per element
  - Total monitoring points in SKA1 - Mid = 207*1000 = 207000
  - Data Growth: 240 TB/year
- SKA1 LOW Telescope
  - 512 LFAA stations, each with 256 antennas
  - 10 monitoring points per antenna
  - Total number of monitoring points = 10*256*512 = 1320720
  - Data Growth: 1540 TB/year

Architecture

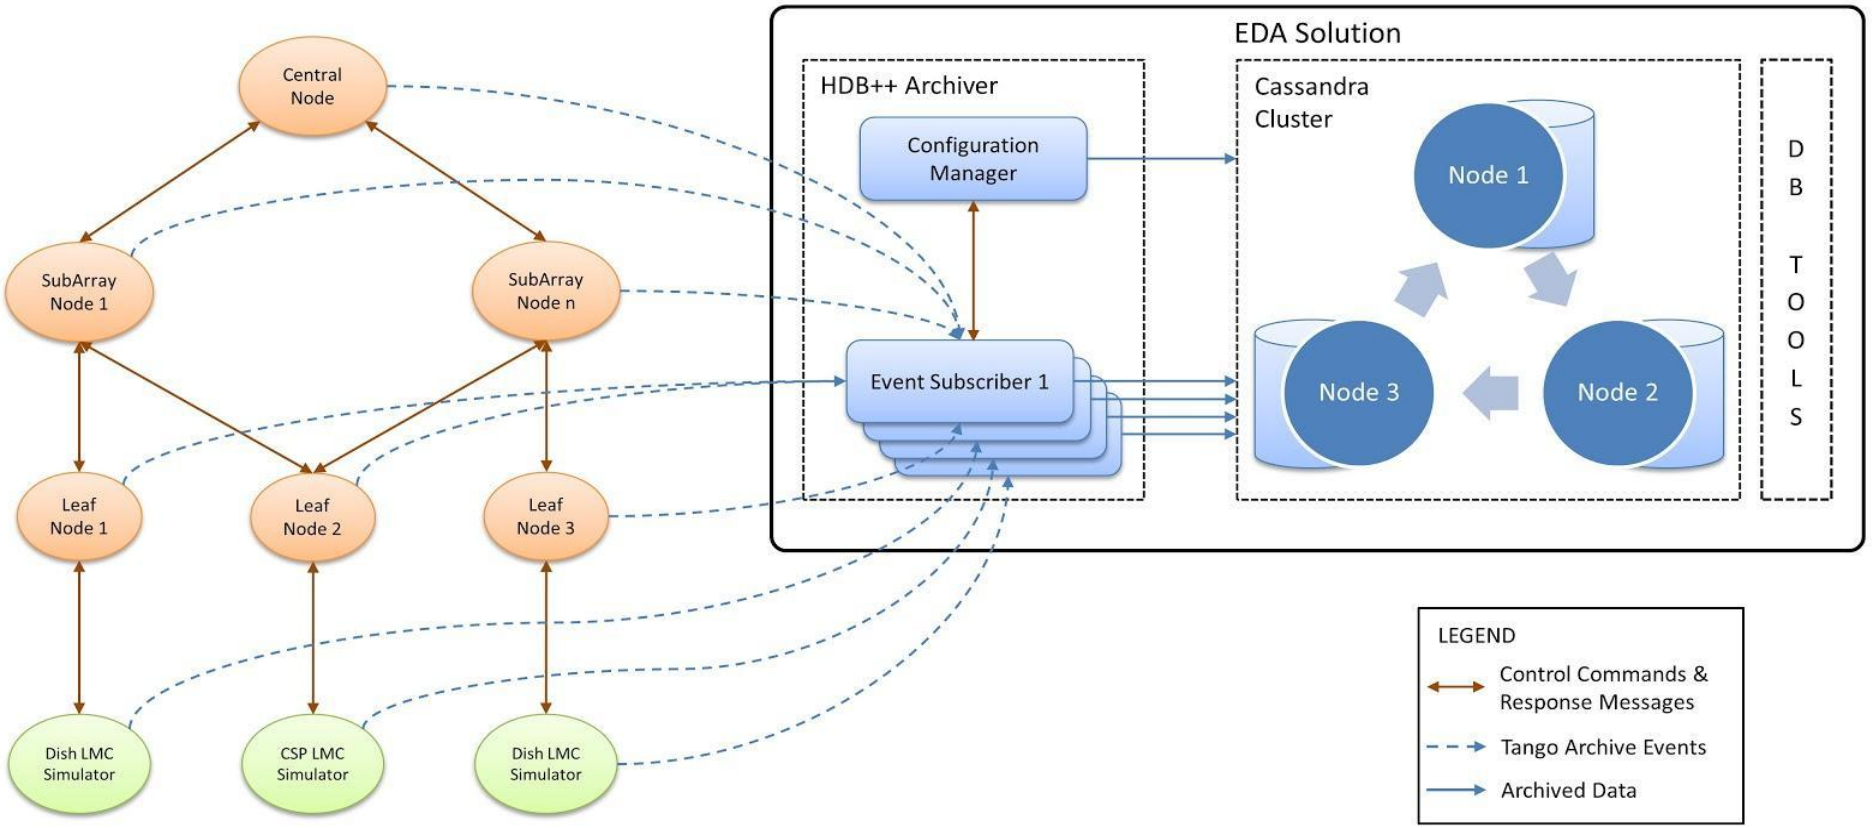

Data Extraction Architecture

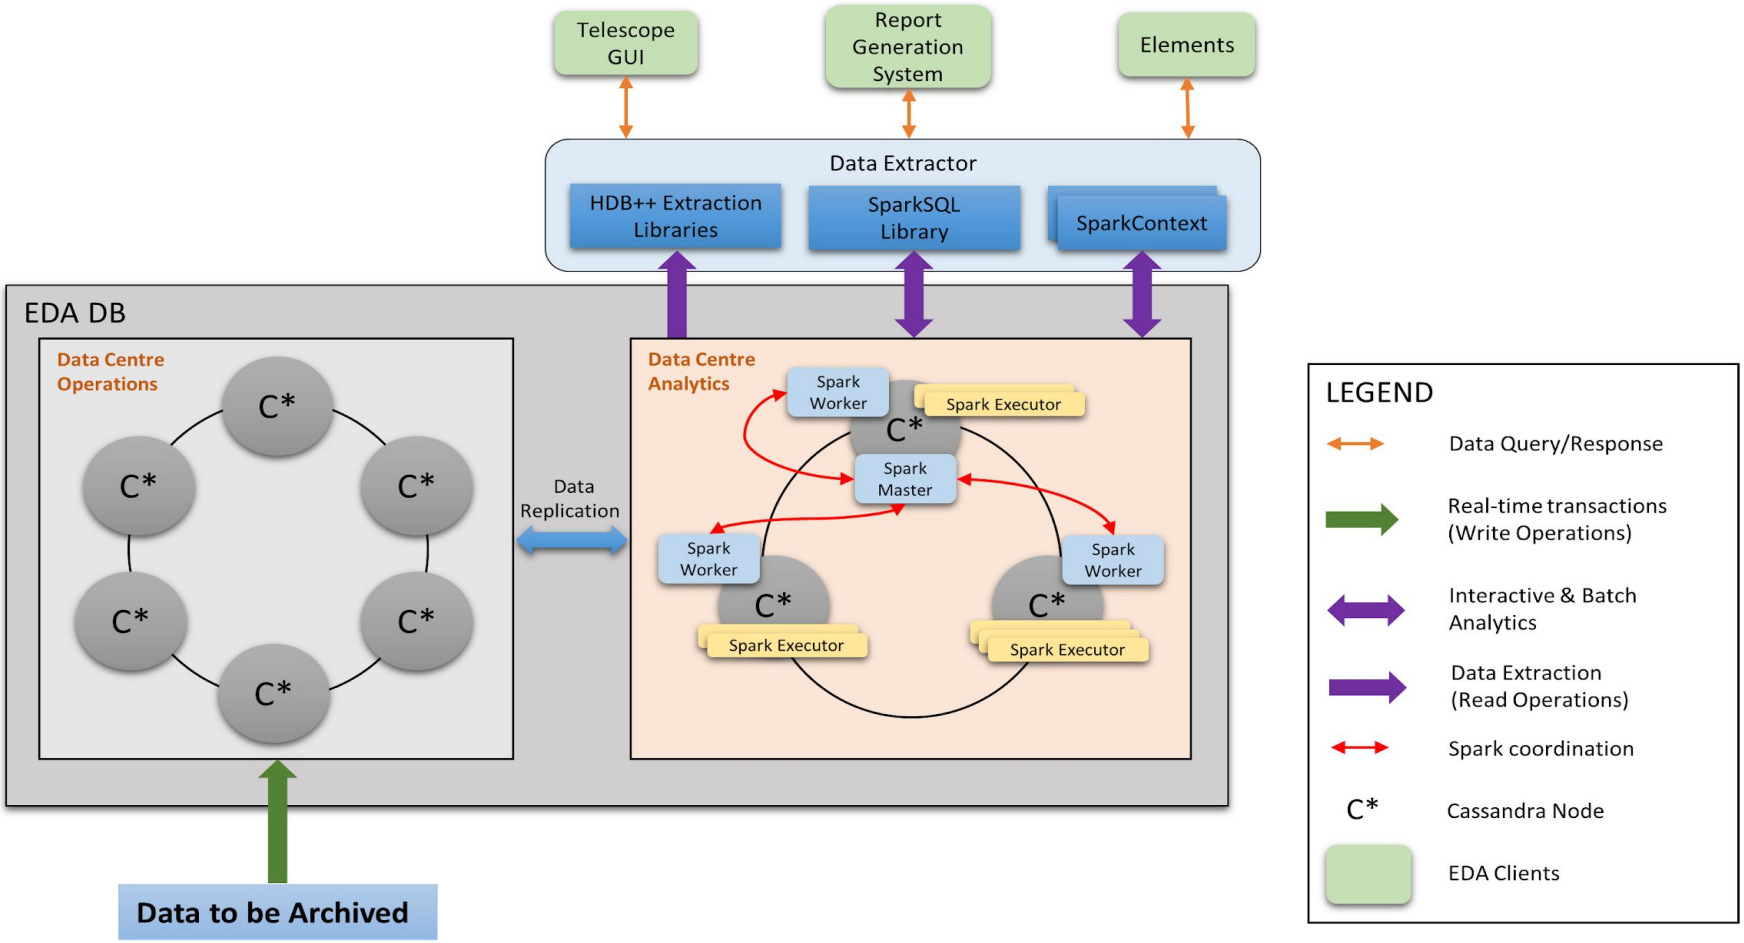

Prototyping

- Performance and Scalability of Archiver and Database
- Throughput handled by a single archiver instance (500 attributes at varying frequencies)
  - 1850 writes per second (year 2017)
  - 2200 writes per second (year 2018)
  - Event Subscribers are able to scale-up linearly to handle 127,000 events/min per Event subscriber (tested with 5 instances of Event Subscribers).

Prototyping….. contd.

- Resilence of archiver - Ability to handle Burst load
  - Gradually step up the load to 60% of ‘threshold’ load
  - Increase load to 2x and maintain for 3-5 minutes
  - Archiver is stable until ‘pending number’ increases

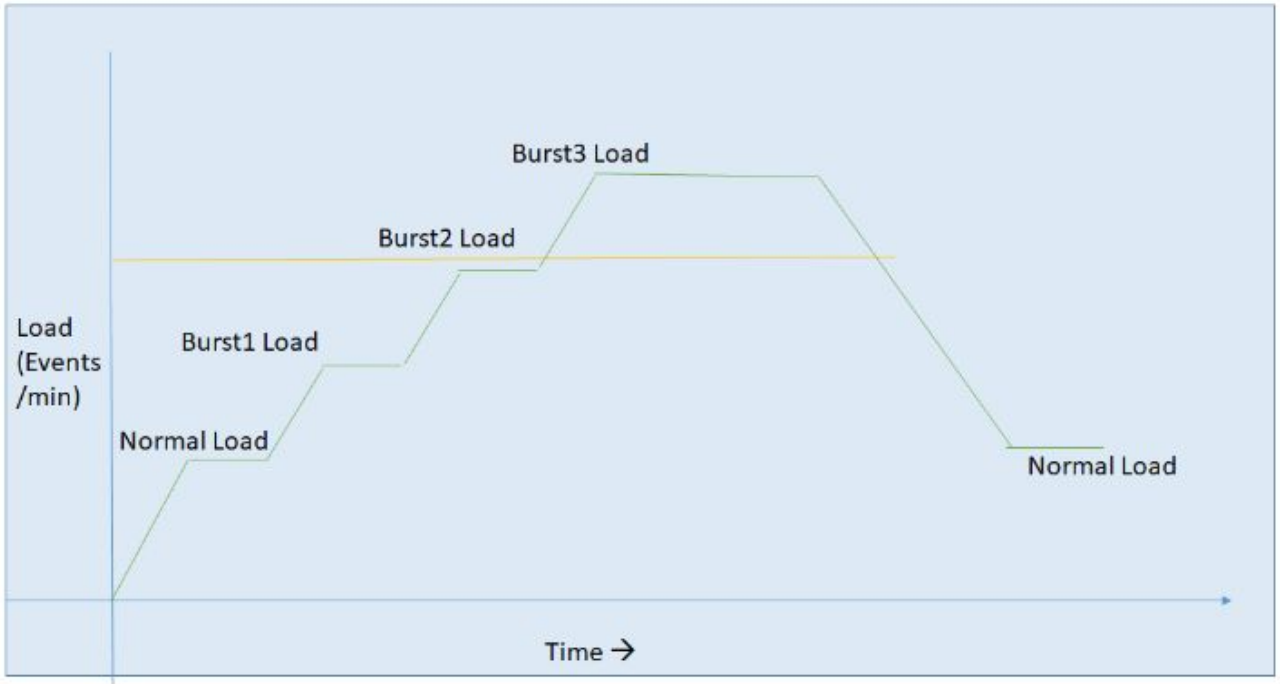

- No significant difference in write performance using ‘light’ schema

CDR Comments

- Power and space constraints limit the number of instance deployed on site
- Data size estimates are approximate. Can be worked through system CDR.
- Risk of over design due to high availability requirements and overestimation of data size/growth
- Stock replication of Cassandra is two way, all data
- Degraded performance of analytics copy of database may affect operational db

Mitigation:

- Start with both clusters (Operations and Analytics) co-located at the CPF. Operations cluster will maintain only few days of data. Analytics cluster maintains all data.
- Consider moderate availability requirement (replication factor =1). This will limit the data storage needs. This should be sufficient for the initial array releases
- Monitor data growth and performance - consider moving Analytics cluster to SOC. Additional computing infrastructure may be added at this stage.
- Adopt custom replication using 3rd party tool for selective, one way replication
- Asynchronous writes possible in replication. This will ensure that Operational db will not be affected by growth of analytics database.

Next Steps

Expecting optimizations to the archiver, which will increase the throughput for ‘write’ operations

Looking at other technologies options to evolve and mature in the SKA construction timeframes (e.g. TimescaleDb)

SQUARE KILOMETRE ARRAY

Exploring the Universe with the world's largest radio telescope