Challenges of Real-time Processing in HPC Environments – the ASKAP experience

Eric Bastholm | Team Leader

19 October 2017

CSIRO ASTRONOMY AND SPACE SCIENCE
www.csiro.au

Outline

- Australian SKA Pathfinder (ASKAP)
- Some specific computing challenges encountered
  - Disk Performance
  - Application Performance
  - Process Isolation
- Lessons learned
- Takeaways

ASKAP:
Australian Square Kilometre Array Pathfinder

ASKAP

- 36-antenna interferometer
- Located in radio quiet zone “outback”
- Supercomputer for data processing
- 10,000 cores, 4 GB/core, 200 TFLOPs Peak
- Data Ingest ~ 2.8 GB/s = ~ 10 TB/h
- Cyclic disk buffer, deleting data post processing
- 5 PB data products per year

image credit: Alex Cherney / terrastro.com


St George Ranges
Port Hedland
Karratha
Tanami
Marble Bar
Onslow
Exmouth
Telfer
Karijini
Karlamilyi National Park
Coral Bay
Paraburdoo
Newman
Minilya
Kintore
Carnarvon
Gibson Desert Nature Reserve
Ngaanyatjarra-Giles
Yulara
WESTERN AUSTRALIA
Warburton
Geraldton
Plumridge Lakes
Mamungari
Great Victoria Desert Nature Reserve
Cervantes
Forrest
Nullarbor
Perth
Caiguna
Madura
Eucla
Mandurah
Bunbury
Esperance
Albany


PAWSEY
supercomputing centre


PAWSEY
20

ASKAP

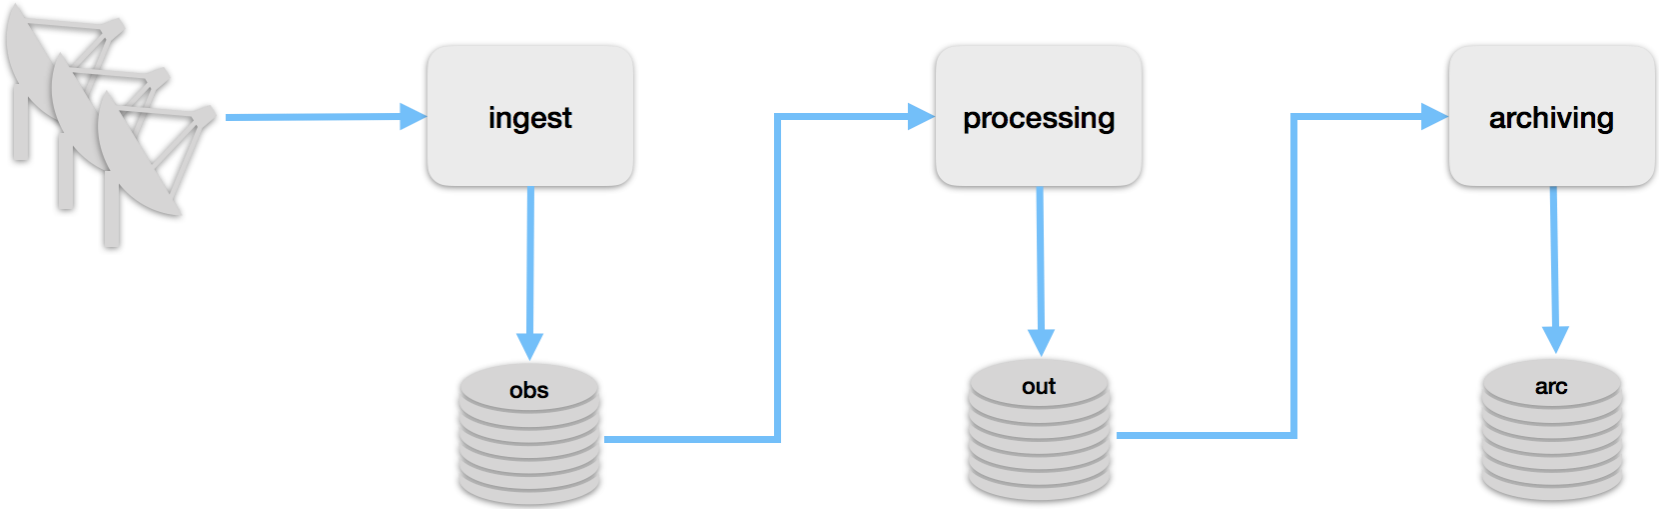

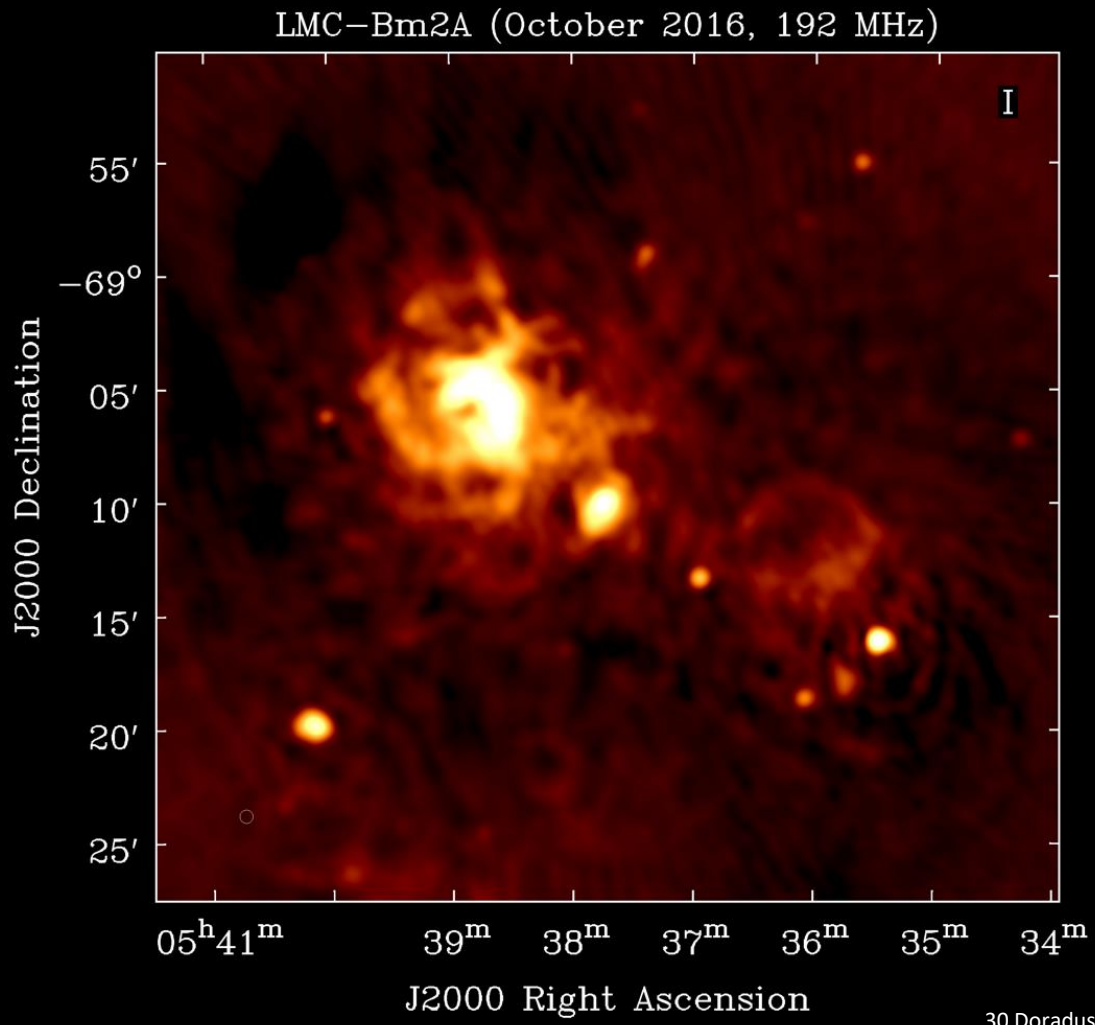
LMC−Bm2A (October 2016, 192 MHz)
I
55′
−69°
05′
10′
15′
20′
25′
J2000 Declination
05h41m
39m
38m
37m
36m
35m
34m
J2000 Right Ascension

Computing Challenges:

Computing Challenges

- Disk performance
- Application performance (memory and messaging)
- Process isolation

Computing Challenges

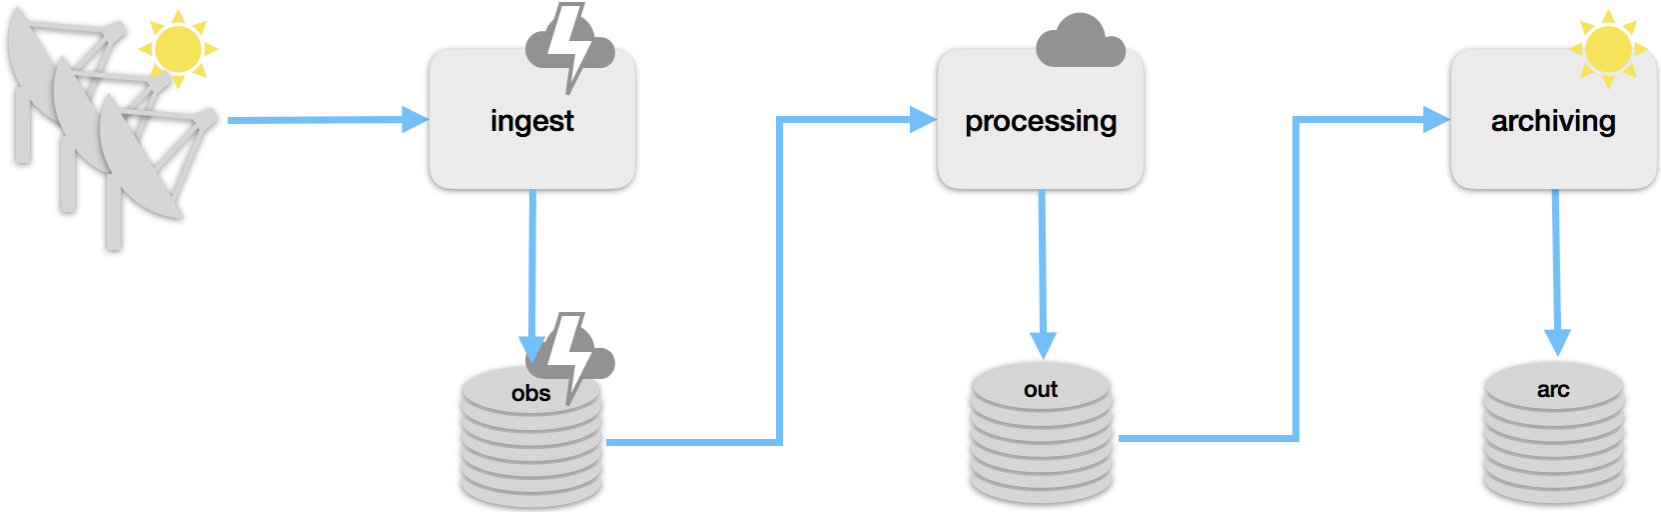

Computing Challenges:
Disk Performance

Disk Performance

- High performance parallel Lustre file system
- 1 PB designed for high throughput (10 GB/s)
- Single thread performance – not so much
- Performance can vary a lot depending on
  - Other users
  - System parameters
  - Legacy code (non-parallel)

Disk Performance

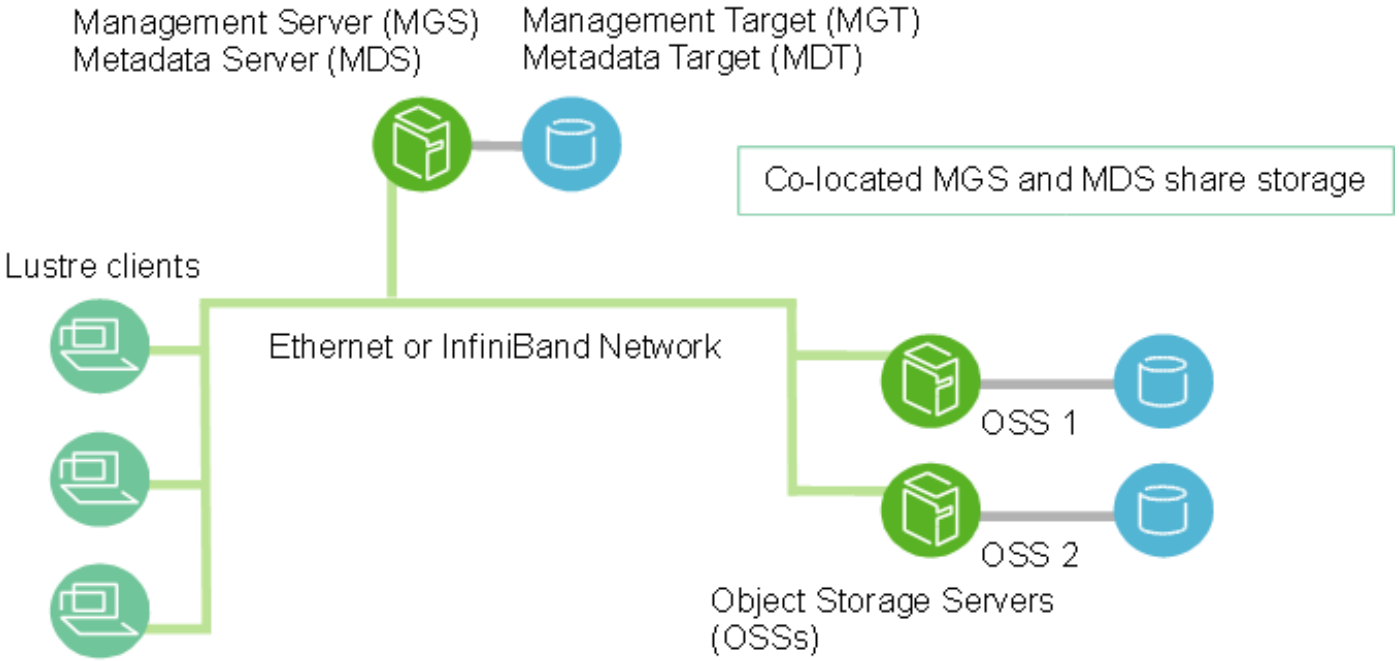

doc.lustre.org

Disk Performance

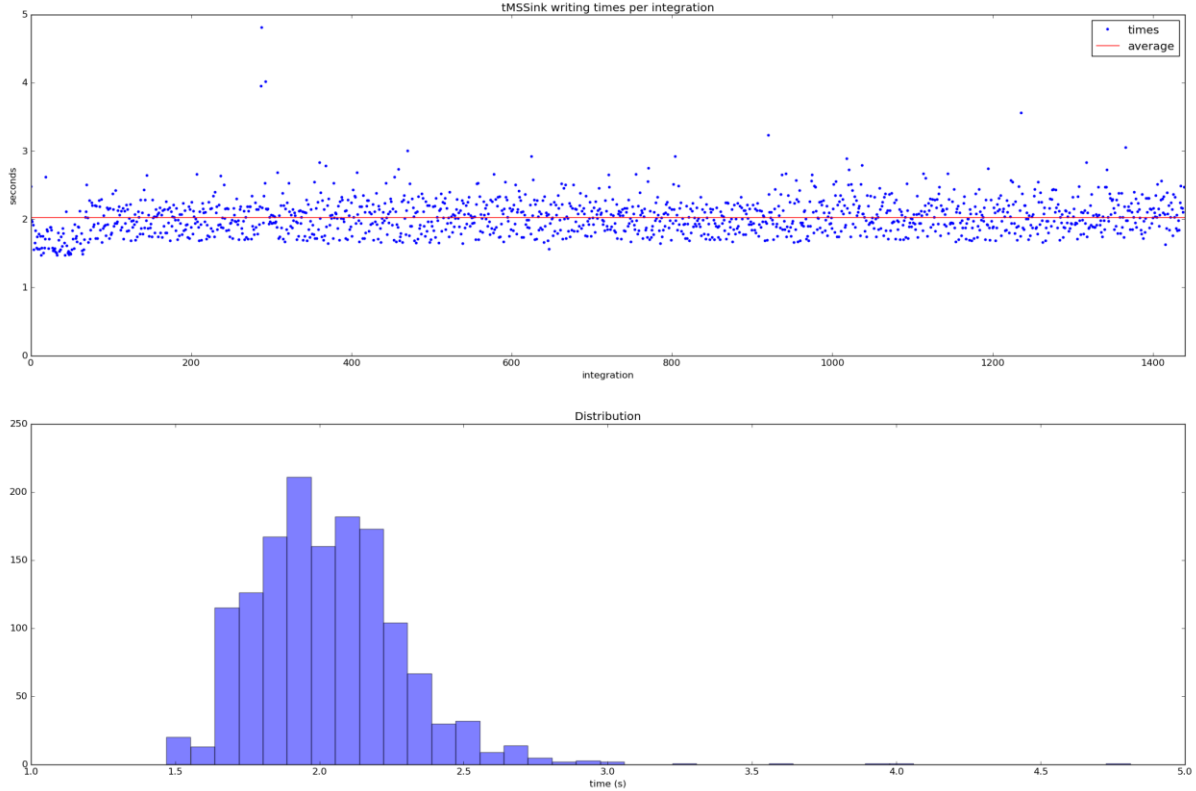

Disk Performance

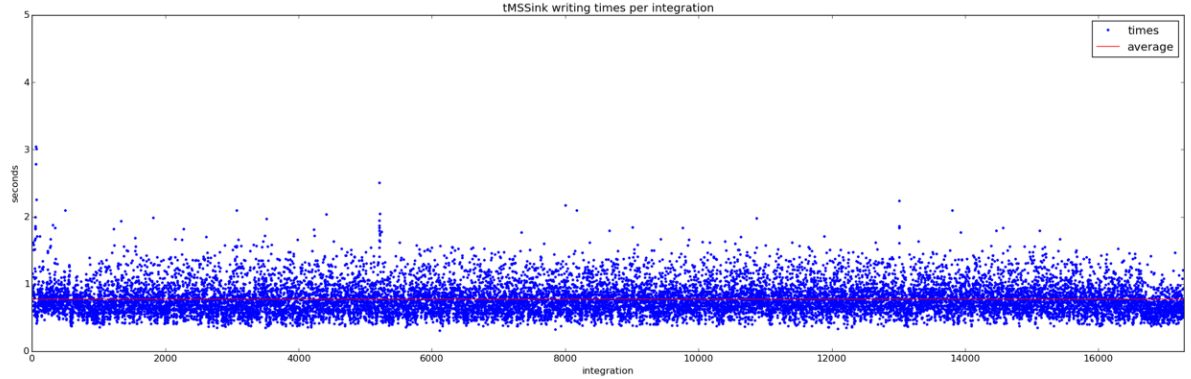

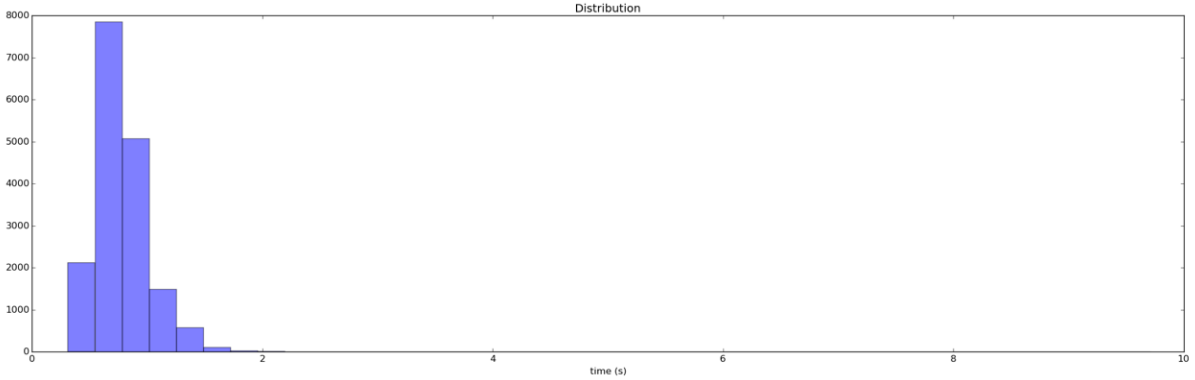

Disk Performance

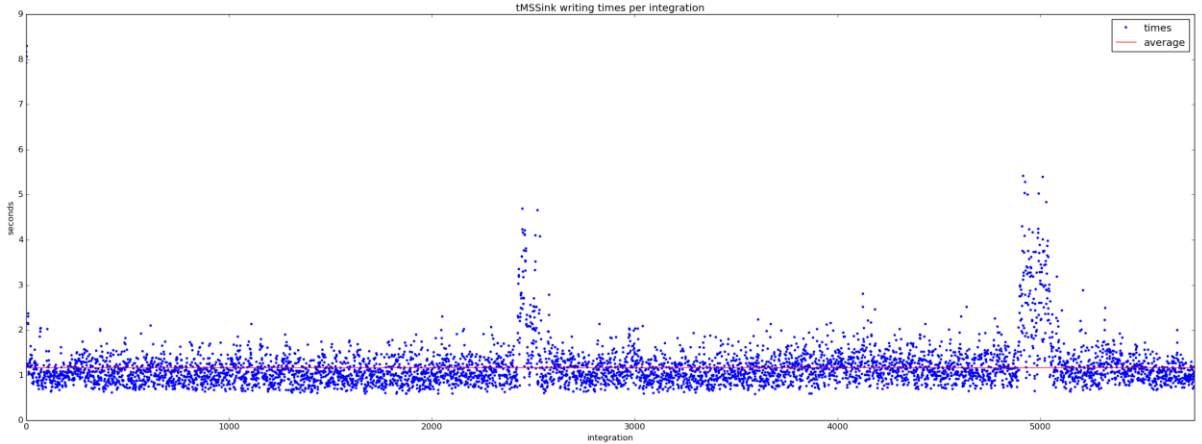

Disk Performance

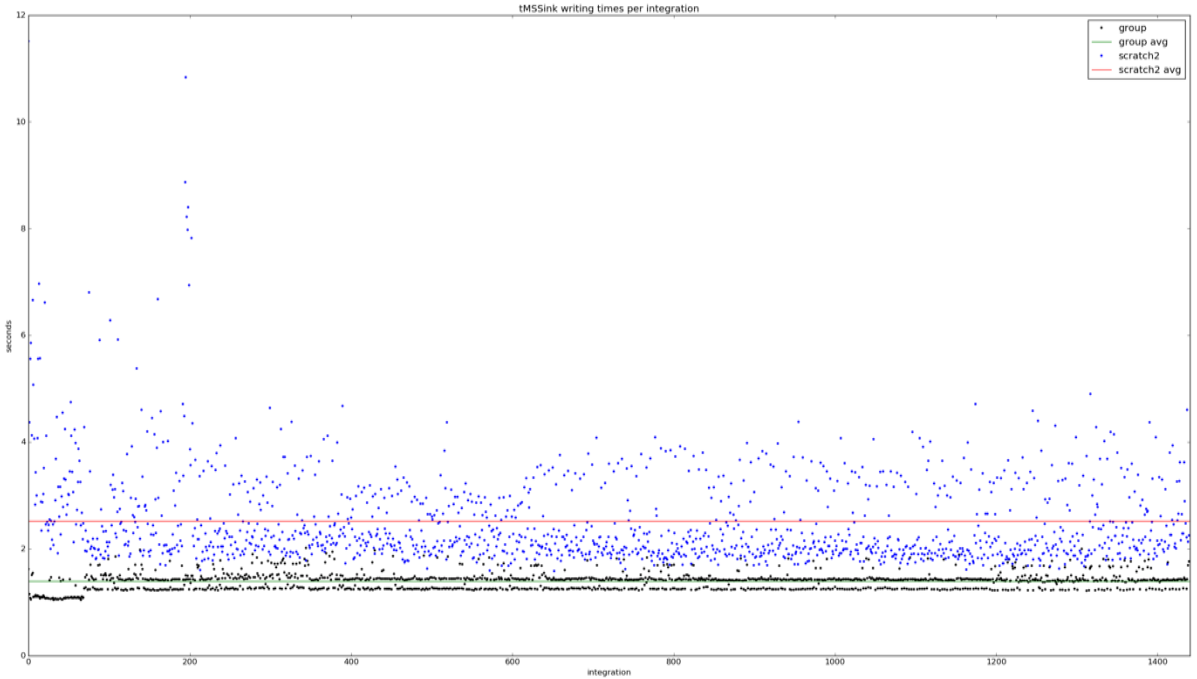

Computing Challenges:
Application Performance

Application Performance

- Locking in messaging (MPI) code between parallel processes
- Low memory transfer rate
- Causes buffer overruns
- Lost data
- Write more files from separate processes

Application Performance

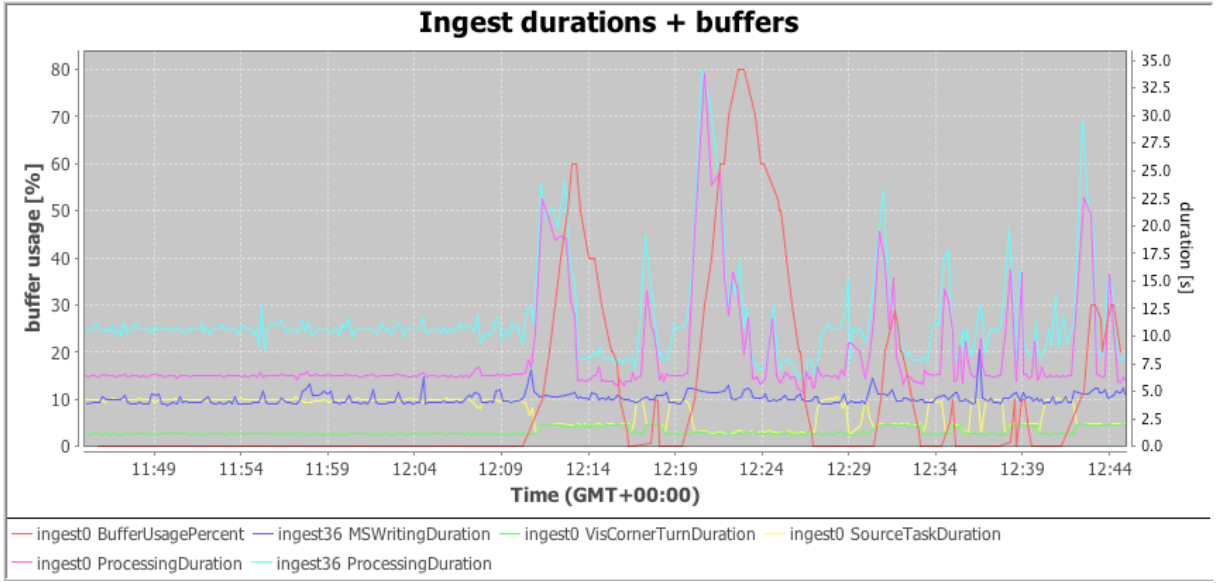

Application Performance

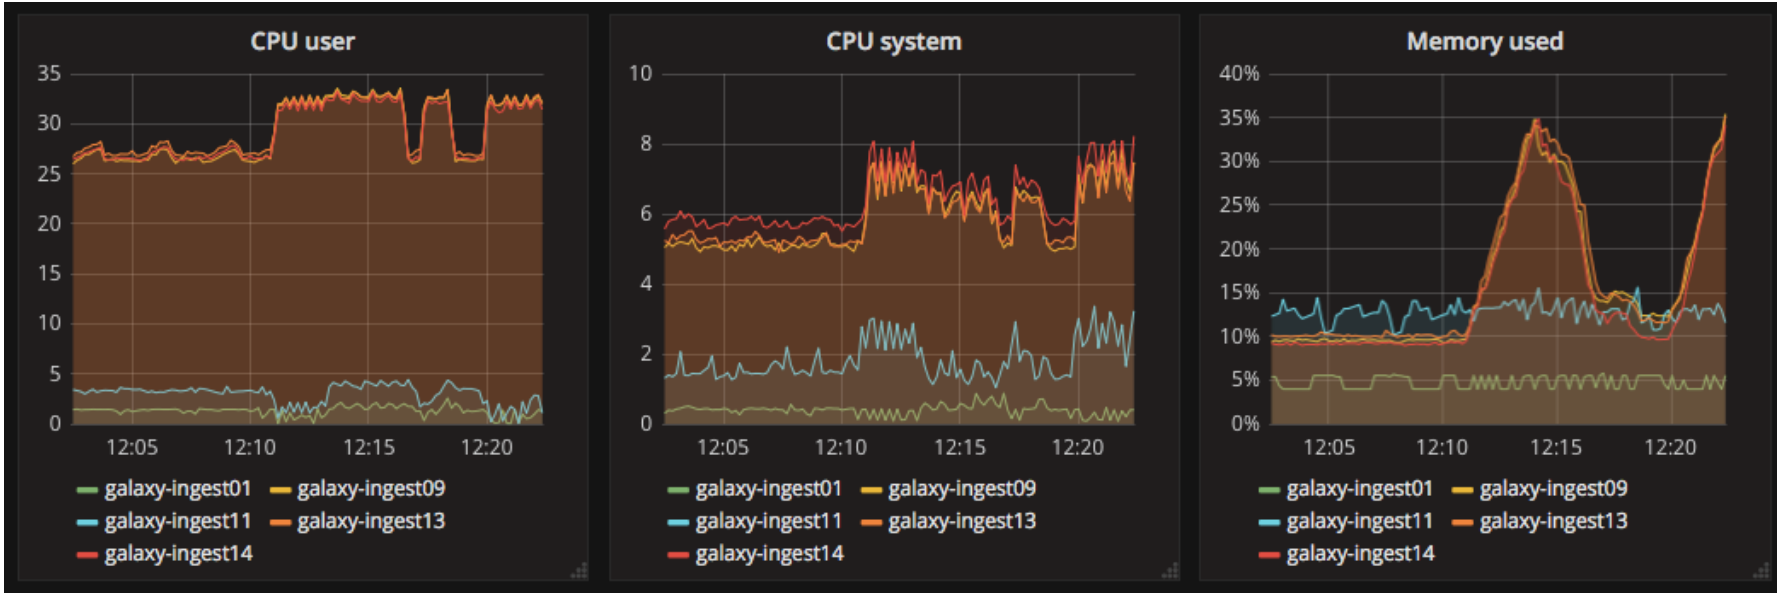

Computing Challenges:
Process Isolation

Process Isolation

- Isolate the critical processes as much as possible
- Dedicate resources they need
- Control read/write access to disk
- Reduce OS jitter by minimising installed image
- Pure isolation not possible in a shared environment

Process Isolation

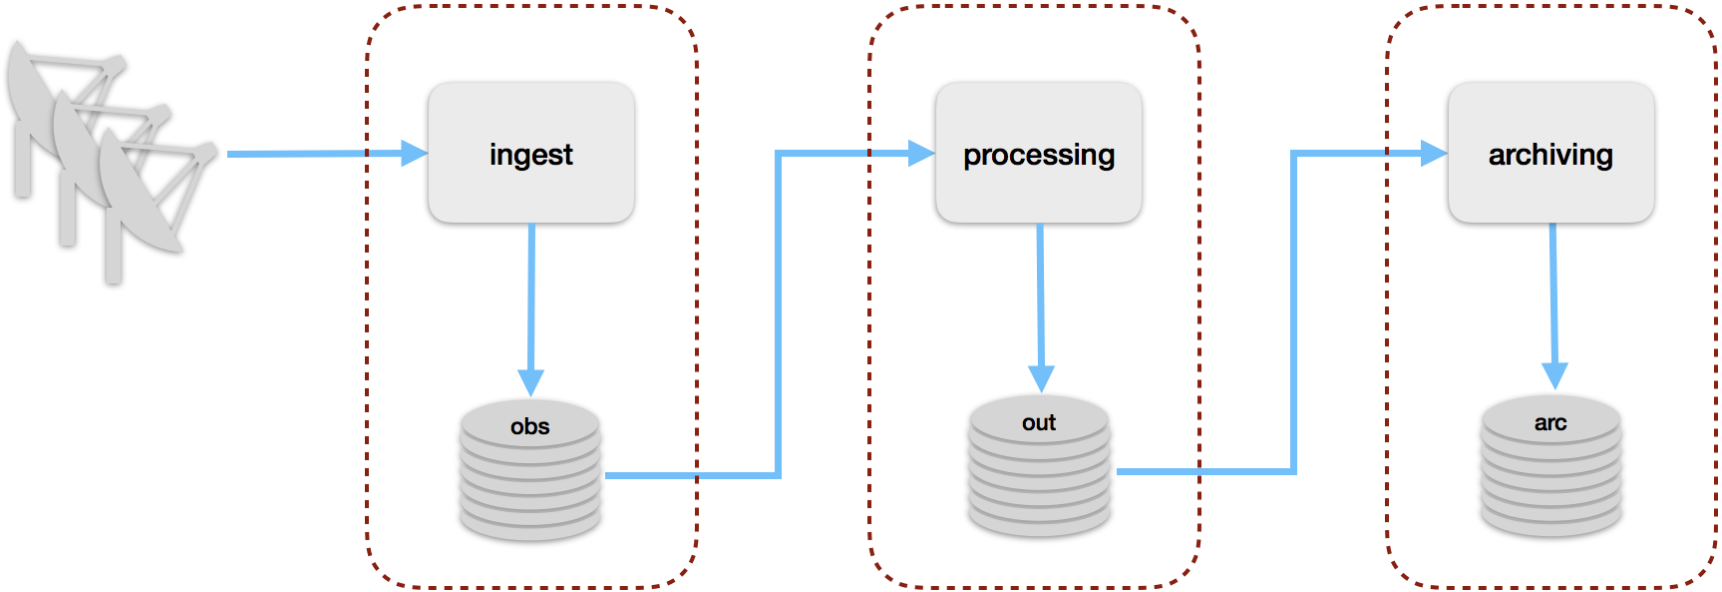

Lessons Learned:

Lessons Learned

- Don’t assume that the platform is infallible just because it’s high tech, BIG, and *FAST*
- Identify possible performance issues
- Prototype at a granularity necessary to enable the important design decisions – walk before you try to run
- Use realistic data sets
- Establish good working relationship with the platform providers
- Develop monitoring and reporting early
- Develop good testing early

Takeaways:

3 Takeaways

- Prototype and integrate often, don’t implement final solution in one leap
- Monitor everything
- Work with platform provider – make them part of the team

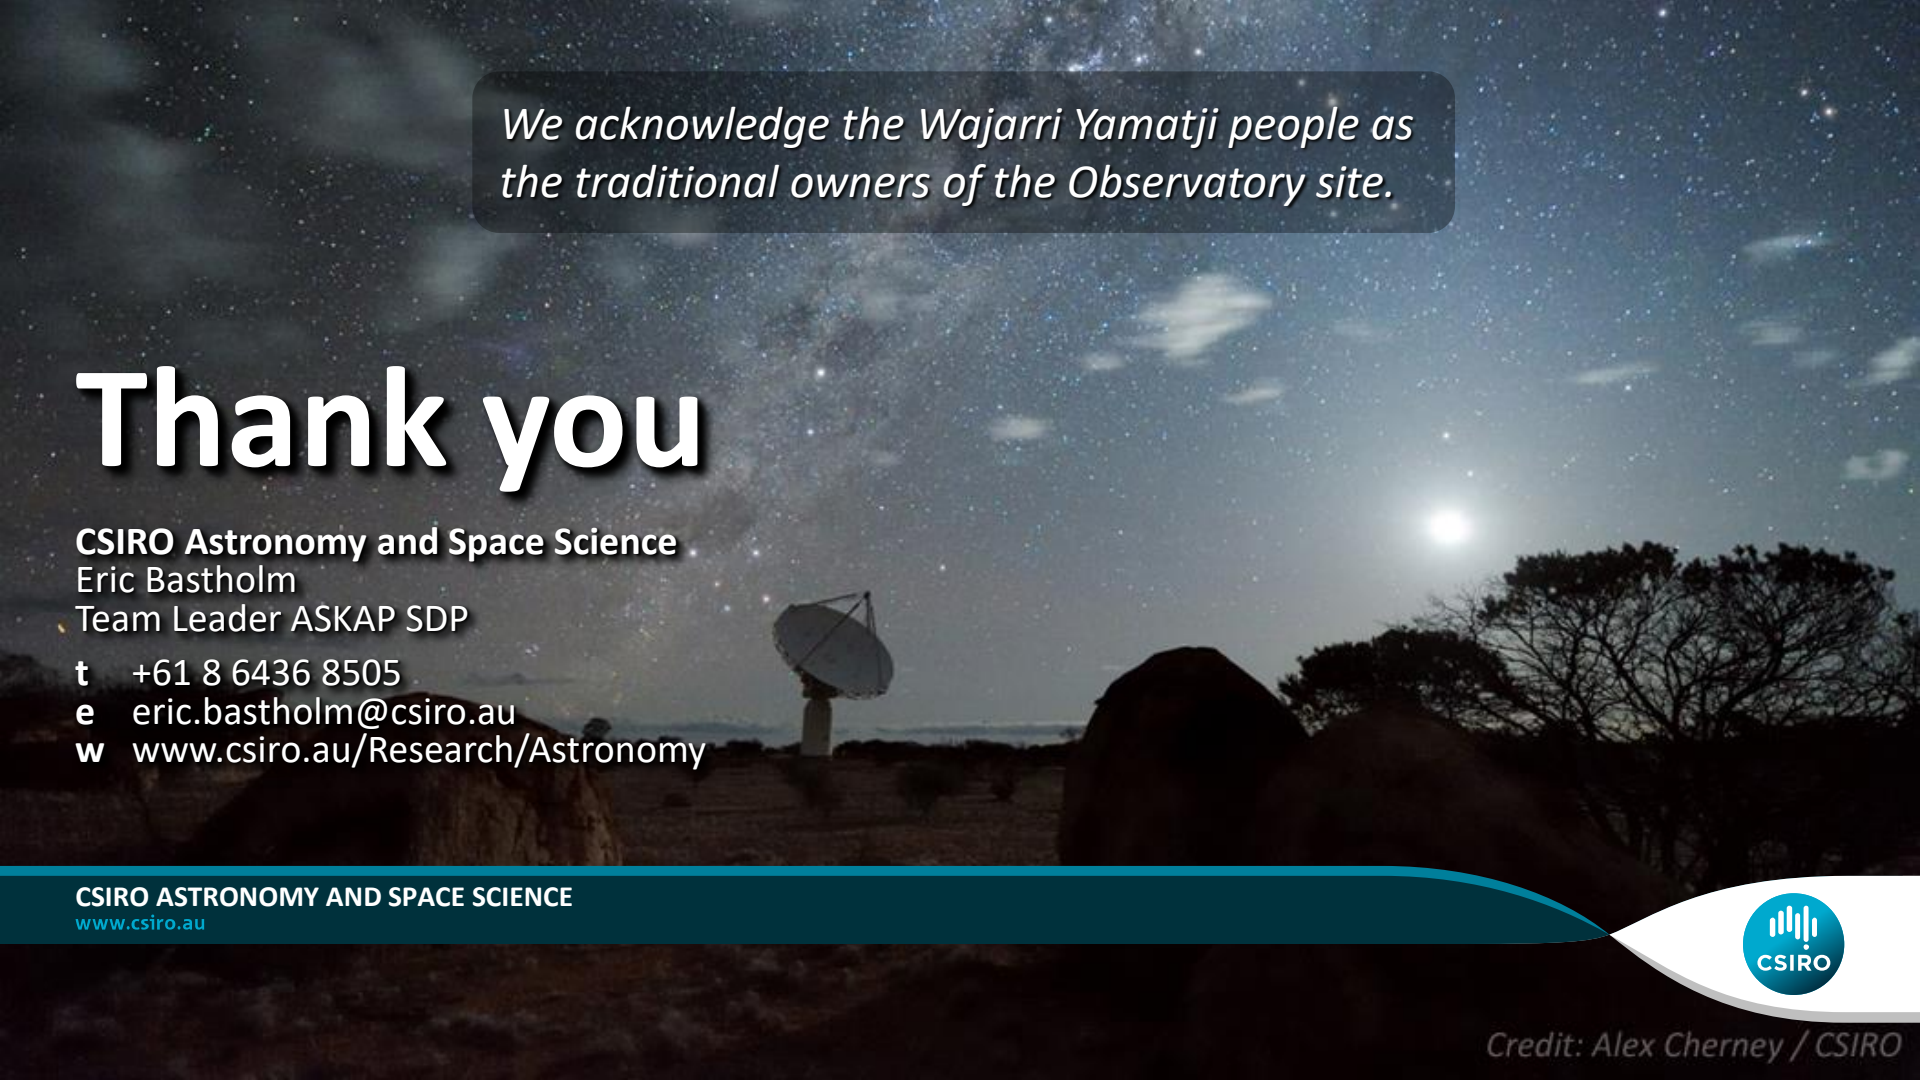

We acknowledge the Wajarri Yamatji people as the traditional owners of the Observatory site.

Thank you

CSIRO Astronomy and Space Science
Eric Bastholm
Team Leader ASKAP SDP

t +61 8 6436 8505
e eric.bastholm@csiro.au
w www.csiro.au/Research/Astronomy

Credit: Alex Cherney / CSIRO