



Ilifu

**Rob Simmonds**

CTO, Ilifu

Acting Director, IDIA

Professor, Dept. Computer Science, UCT









Inter-University Institute  
for Data Intensive Astronomy

## My Cloudy History

- CTO within Compute Canada organisation
- Started in this area before Cloud rolled in
- Deployed Nexus and all Globus environments
- On initial board of Grid Canada CA
- Worked with HP labs (now HPE labs) on utility computing
- Worked on setting up Cloud environments for Compute Canada
  - Initially deployed Nimbus
  - Later deployed OpenStack (once it was created) environments East and West





# IDIA

Inter-University Institute  
for Data Intensive Astronomy

## IDIA

- Partnership of SA universities
- Major focus on MeerKat data processing on the path to the SKA
- Director one of the original SKA proposers
- Simmonds led the data delivery design for the SKA and is consulting to the SKA project now
- Associate Director Brad Frank represents SA on SKA Regional Centre planning committee





Inter-University Institute  
for Data Intensive Astronomy

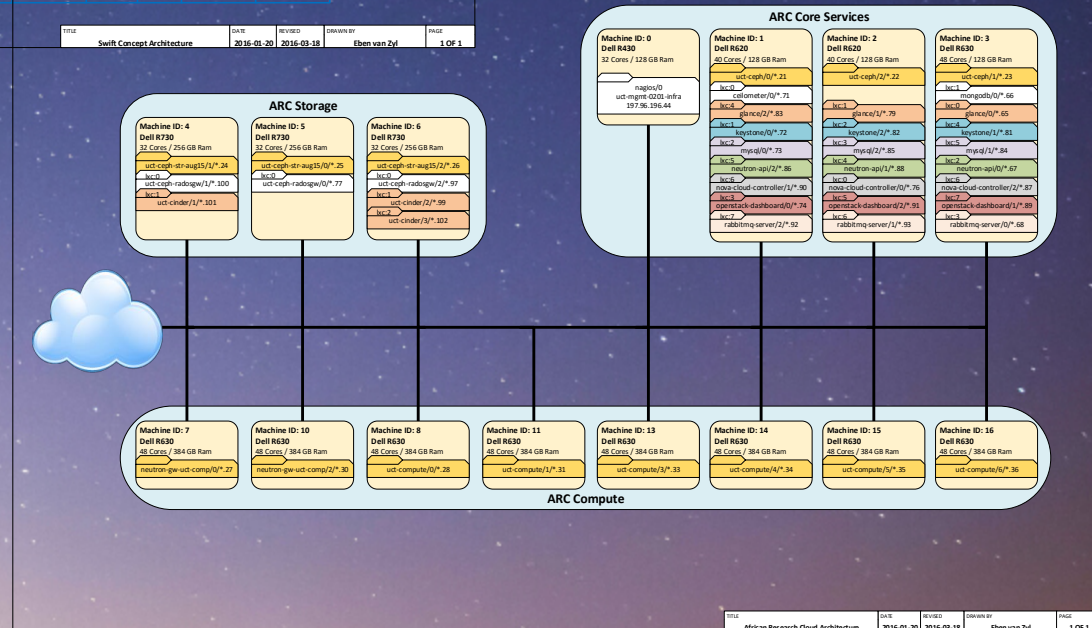
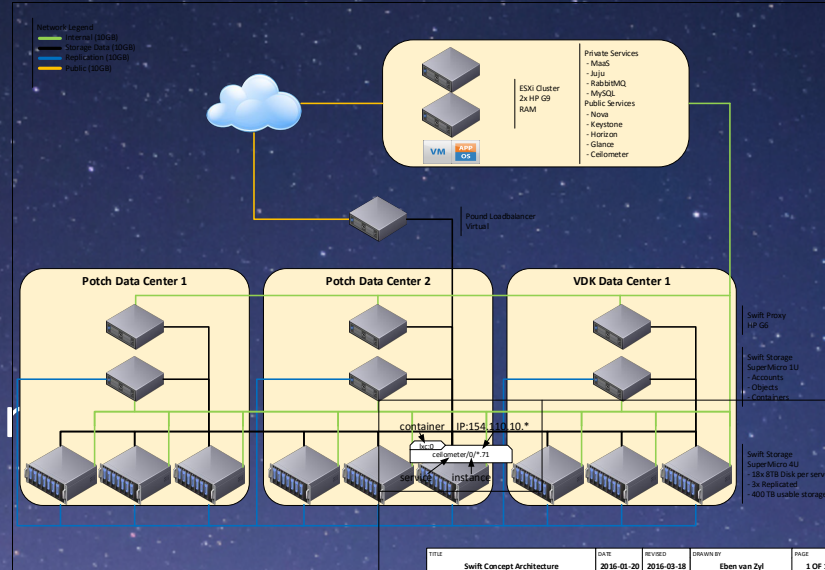
## Iifu cluster

- Iifu is “Cloud” in the Xhosa language
- Started at ARC at UCT and NWU
- Funded at a Tier 2 site by DIRISA in 2017
- Built using funding from DIRISA, IDIA and CBio
- Primarily supporting Astronomy and Bioinformatics
- Astro support for MeerKat LSPs and MeerLicht



# ARC

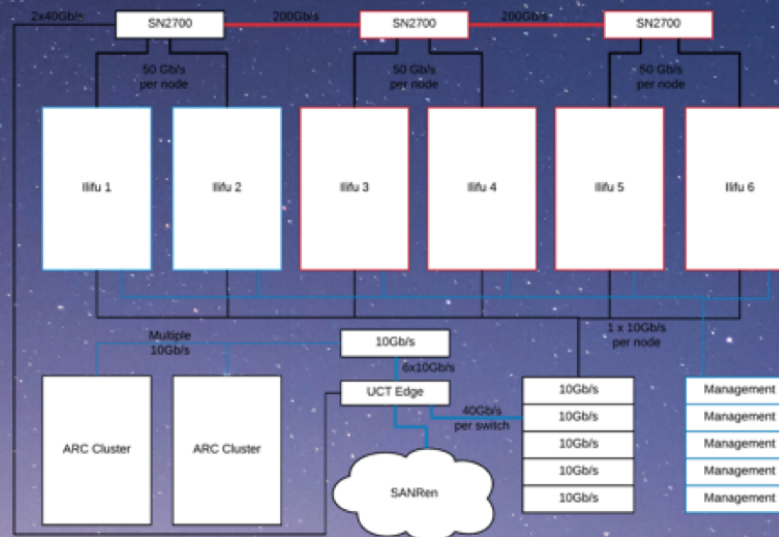
- Systems at UCT and NWU
  - UCT – Compute and mixed storage
  - NWU – Object Storage
- Redundant core management system
- 8 compute nodes
  - 2.6GHz Xeon Processors
  - 24 cores, 385 GB RAM / node
  - 192 cores, 3TB RAM total
- UCT CEPH storage
  - ~300TB usable from 3 targets
- NWU Swift storage
  - ~400TB usable from 8 targets
- 10Gb/s connection to SANReN
  - Not through institution firewalls





# Ilifu architecture

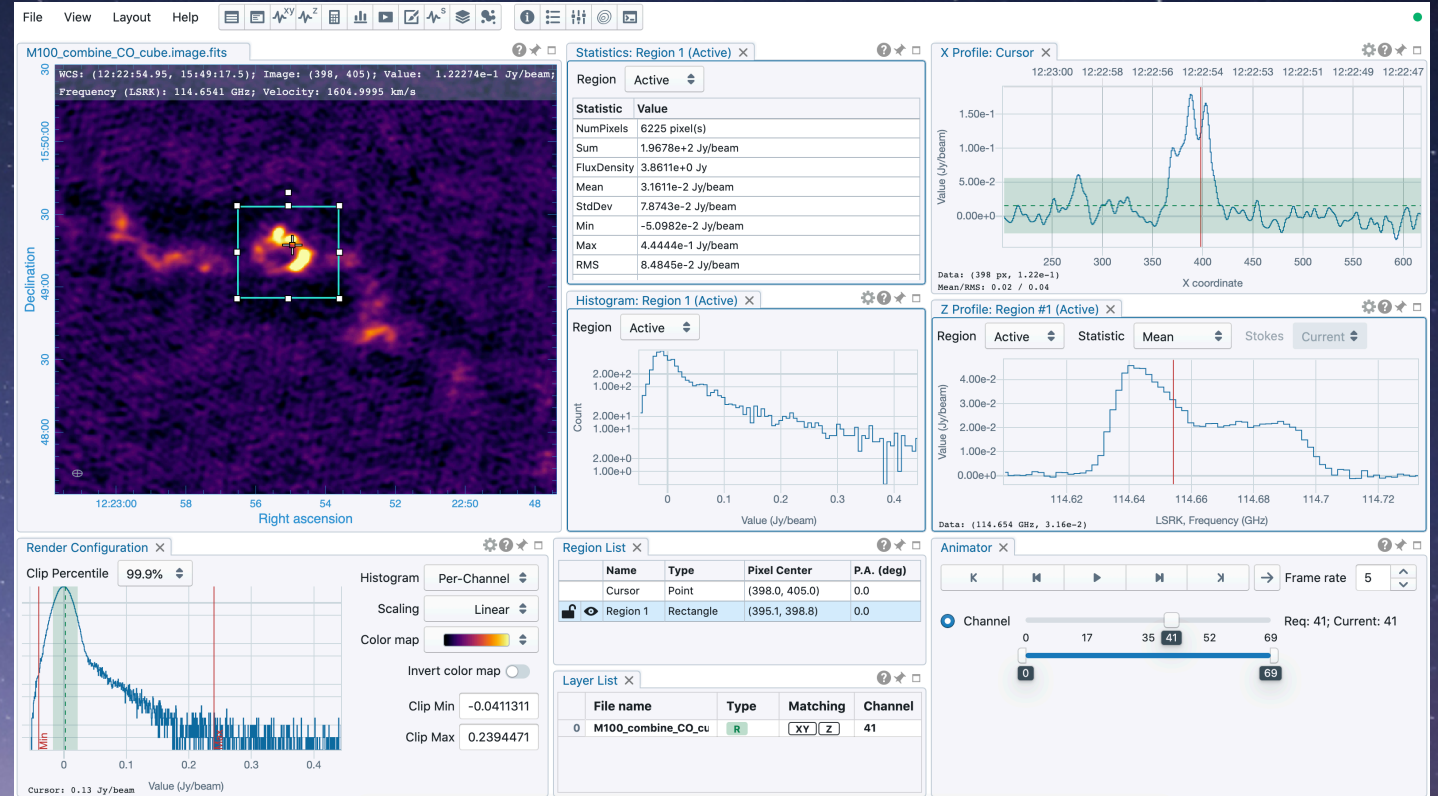
- Main Ilifu cluster now in 7 racks (diagram old)
- ARC nodes now act as development/training area
- Uses BeeGFS for scratch and Ceph for other storage (over 8PB in total)
- 10Gb/s WAN access moving to 100Gb/s Q2
- Moving to hierarchy in MLX backplane to provide 400Gb/s between switches





## Services

- Web services
  - OpenStack
  - Jupyter
  - CARTA
  - Galaxy
- Terminal services
  - SLURM clusters
  - “Fat” nodes
- Data management / movement
  - FTS, Rucio, Globus







Inter-University Institute  
for Data Intensive Astronomy

## AAI

- LDAP used for Unix namespace mappings
- Moving to KeyCloak for AAI integration and SSO
  - In production for OpenStack.
  - In dev environment for other services but will move to production soon
- Using EGI Checkin for:
  - SAFIRE AuthN – SA university / government authentication federation
  - Edugain AuthN
  - Google Enterprise AuthN
  - CoManage role based AuthZ





Inter-University Institute  
for Data Intensive Astronomy

## Ilifu future

- USA funding sought for Bioinformatics
- SARAO providing funding to support MeerKat processing
- Discussions on how to incorporate Ilifu in SKA regional centre





Inter-University Institute  
for Data Intensive Astronomy

## Summary

- Ilifu in full production
- EGI AAI active on OpenStack, moving to production on other services very soon
- Major focus on SA based BioInformatics and support for MeerKat processing
- IDIA keen to provide software to EGI marketplace



