Vera C. Rubin Observatory
Rubin Observatory Project Office

Summit Computing Cluster

Cristian Silva

ITTN-061

Latest Revision: 2022-07-12

D R A F T

Abstract

Description of the deployment of the summit computing cluster.

Change Record

Version	Date	Description	Owner name
1	YYYY-MM-DD	Unreleased.	Cristian Silva

Document source location: `https://github.com/lsst-it/ittn-061`

Contents

Summit Computing Cluster

1   Introduction

The workload of the commissioning cluster (antu) will be absorbed by the Summit computing cluster (yagan); some of the benefits of this topology are:

- Reduced administration overhead

- Limit security constraints to a single location

- Simplifies configuration deployment

- Promotes summit independence

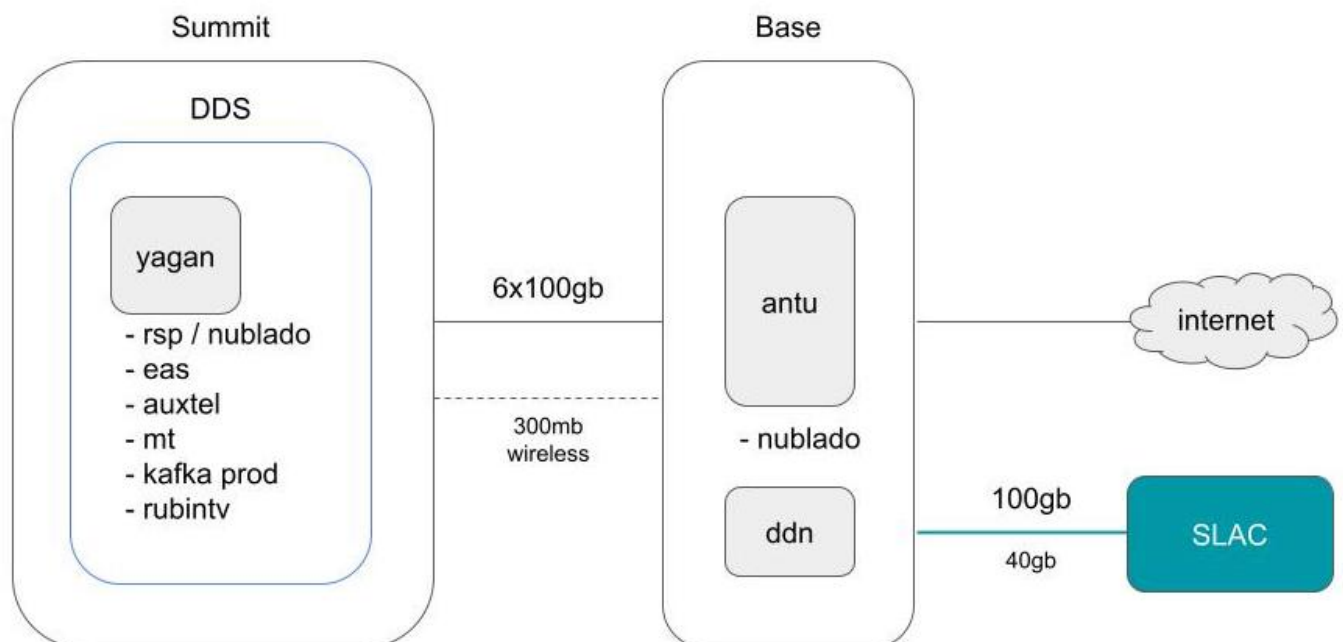The big picture of the current topology is the following:

FIGURE 1: topology

2 Current State

2.1 Hardware

The commissioning cluster runs in the antu nodes, with a capacity of 832 cores and 3.2TB of Ram. Details of its hardware configuration can be found in ITTN-014 Computing Infrastructure

FIGURE 2: antu cluster

The servers composing the antu cluster are a mix of different Dell models, previously deployed as forwarders and DTNs by NCSA. There's also a DDN unit of about 700TB of storage, and it is currently providing NFS mounts for Nublado

The computing cluster at the summit is called yagan. It has a capacity of 576 cores and 2.2TB of Ram.

FIGURE 3: clusters at the summit

The servers in the yagan cluster have the same model and specs. This hardware model is shared with several other services and clusters at the summit, so they all work as a big pool of spares in case a node fails.

2.2   Software

All clusters are provisioned using Rubin Devops Stack (puppet, ipa, etc)

The antu cluster runs:

- Nublado

The yagan cluster runs:

- Rubin Science Platform

- EFD (in migration process)

- EAS CSCs

- Auxtel CSCs

- Kafka Producers

- MT CSCs

- RubinTV

Regardless of the several components running in yagan, its cpu usage stays below 10% and the RAM usage is never above 5%

3   Planned State

Each rack at the summit has a maximum of 48RU, 8RU are used by network equipment such as switches and fiber headers. Therefore each rack could potentially host 40 servers of 1RU each.

The kubernetes clusters at the summit are all of 1RU with 64 cores; hence each rack could deliver up to 2500 cores. However, given the current commissioning and summit cluster size, the maximum size of the new computing cluster would be 1400 cores.

A   References

B   Acronyms

Acronym	Description
DDN	Data Delivery Network
EAS	Environmental Awareness System
EFD	Engineering and Facility Database
MT	Main Telescope
NCSA	National Center for Supercomputing Applications
NFS	Network File System
PMO	Project Management Office
RAM	Random Access Memory