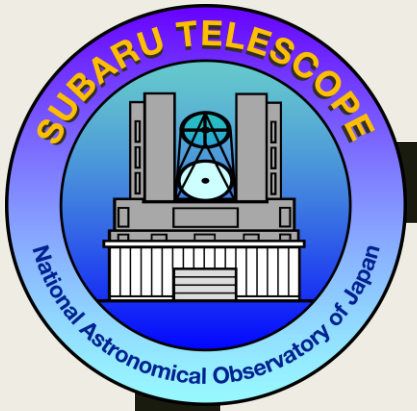

SUBARU TELESCOPE'S COMPUTING SYSTEM ENVIRONMENT (STN6)

Kody Haleamau Rubio & Kiaina Schubert
Subaru Telescope - Hilo & Mauna Kea, Hawaii
Computer and Data Management Division
January 2024

Abstract

1) STN5 monolithic servers and storage environment.
    - The rental contract of Subaru computer system STN5, ended in 2023.
      STN5 included the **Data Archive**, **Observatory Management**, **Network** and **Tape backup**.
    - STN5 implements forty-four physical devices to create a “no downtime environment”

2) Introduction of STN6
    - Implementation of Hyper-Converge Infrastructure (HCI) for a suitable replacement
    - Creating a 95-percentage virtual machine environment with external large storage servers for observation and daily user data
    - CDM to provide a similar environment with a smaller footprint, while meeting budgetary constraints
    - HCI with a separate STARS storage proven to be a reliable and efficient replacement for STN5

3) Network
    - Implementation of a modular infrastructure environment
    - Maximizing network speeds throughout heavy used big data systems

STN5

Hybrid-Physical Virtual Environment (Traditional, Non-Converged)

Breakdown of composition

- 31 Physical Servers
- 25 Virtual Machines
- RedHat Server OS
- Tape backup system LTO7
- 5 Rack Units(4 Base /1 Summit)

Main Specifications

- CPU: 1664 cores (31 physical servers)
- Memory: 3328 GB or 3.3TB
- Storage: 840 TB

Power Consumption

- Base and Summit: 12,000 Watts

- Hardware heavy system

- Typically, each server had about 256GB of memory, which is usually used for high performance servers
- Allocated amount of resources (memory and CPU cores) were over provisioned based on actual performance
- CPU Load averaged less than 1% a month
- Highest load maximized at 20%
- Required manual intervention
- A failure of any type could lead to a catastrophic failure
  - Unexpected downtime
  - Unknown duration
  - Potential data loss

- Prone to failures due to how dependent hardware each server required
- Parts and support scarcity

STN5 Memory, CPU & Storage Usage Summary

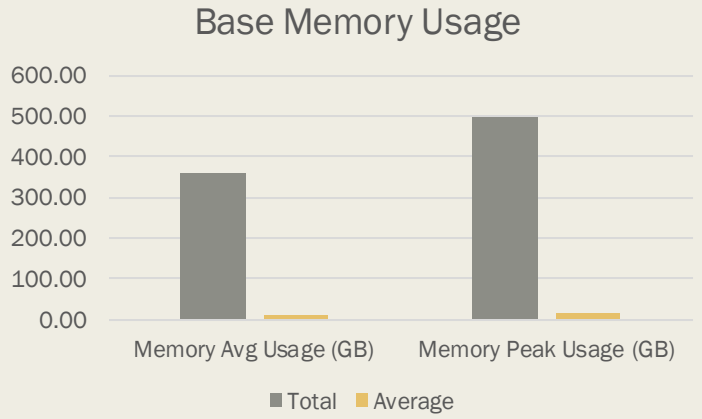

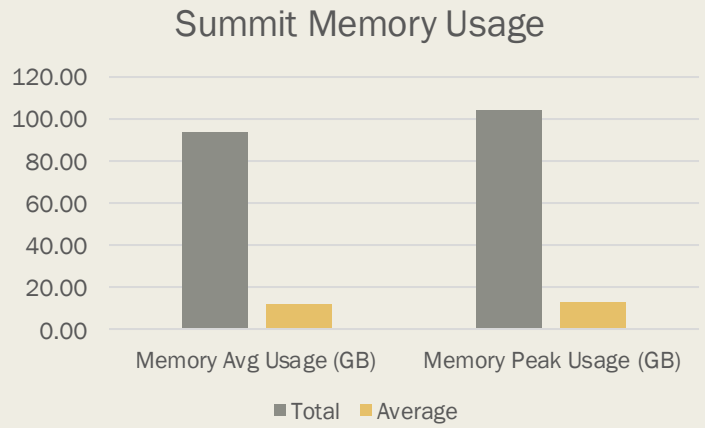

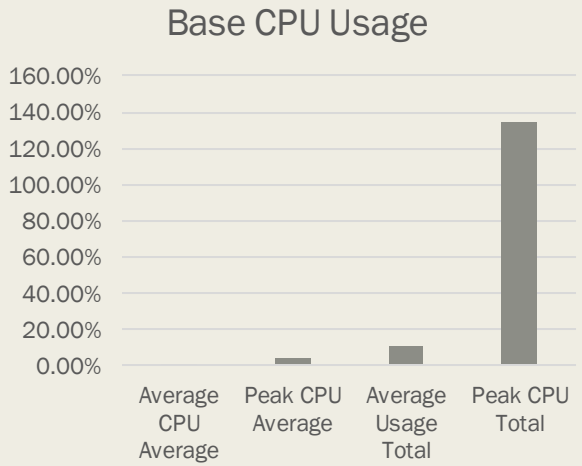

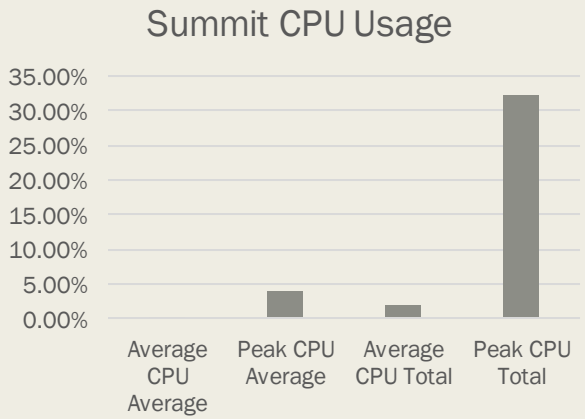

Values and Percentages are based on a single server in STN5 cluster

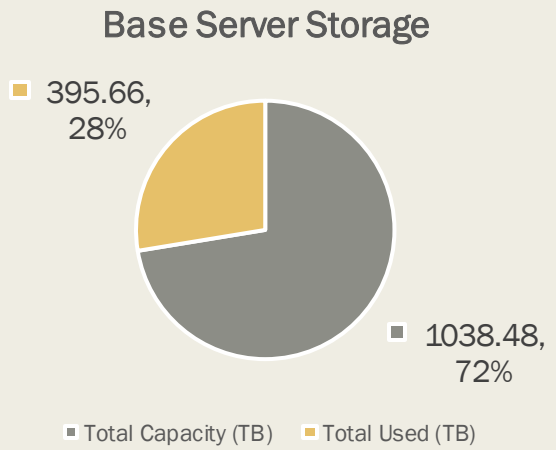

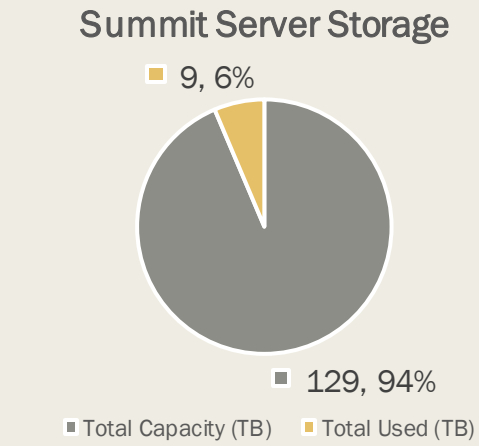

Values and Percentages are based STN5 cluster as a whole

HCI – Hyper Converged Infrastructure

- Cluster of nodes, sharing disk space, memory and processors
- Seamless movement of virtual systems within a cluster
- Administration and maintenance is be reduced
- Load balanced and resources the proper amount of memory and CPU cores into a virtual machine based on the server's actual performance
- Storage in a shared environment across all nodes
- Recovery factor of 2
  - *Provides data redundancy, reduces risk of data loss and increases recovery time.*
  - *All data is located in two different locations*
- Ensures zero downtime in case of hardware or software failure
- Data loss is little to none
- No catastrophic failures in terms of HCI
  - *Term "catastrophic failure" in HCI is minimized to more of a "failure"*
- Eliminates urgent manual intervention
- Automates migration of virtual machines during a "failure"

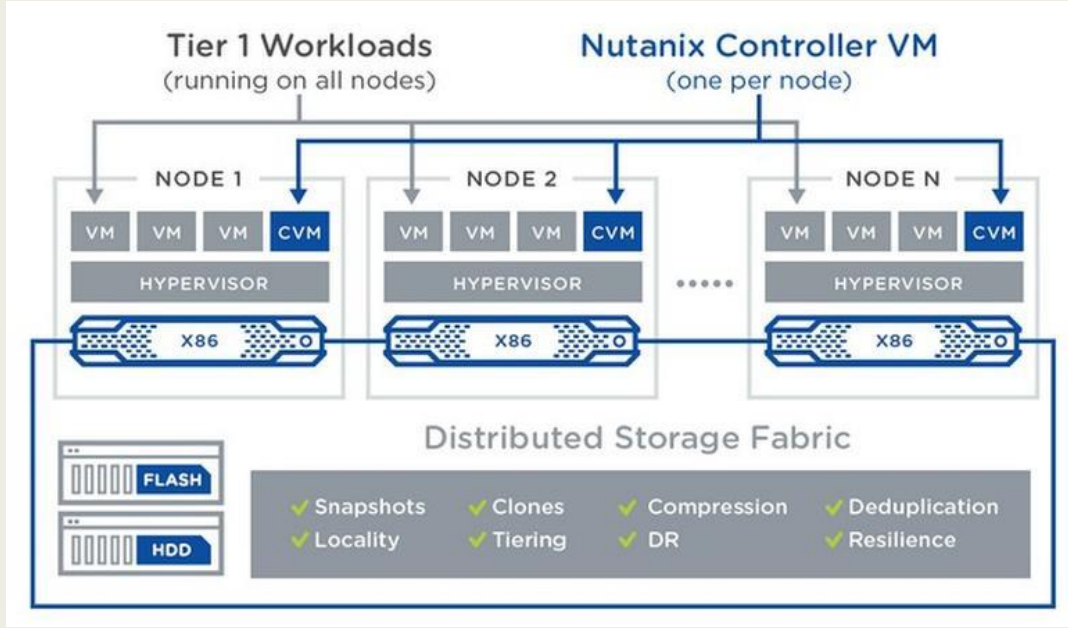

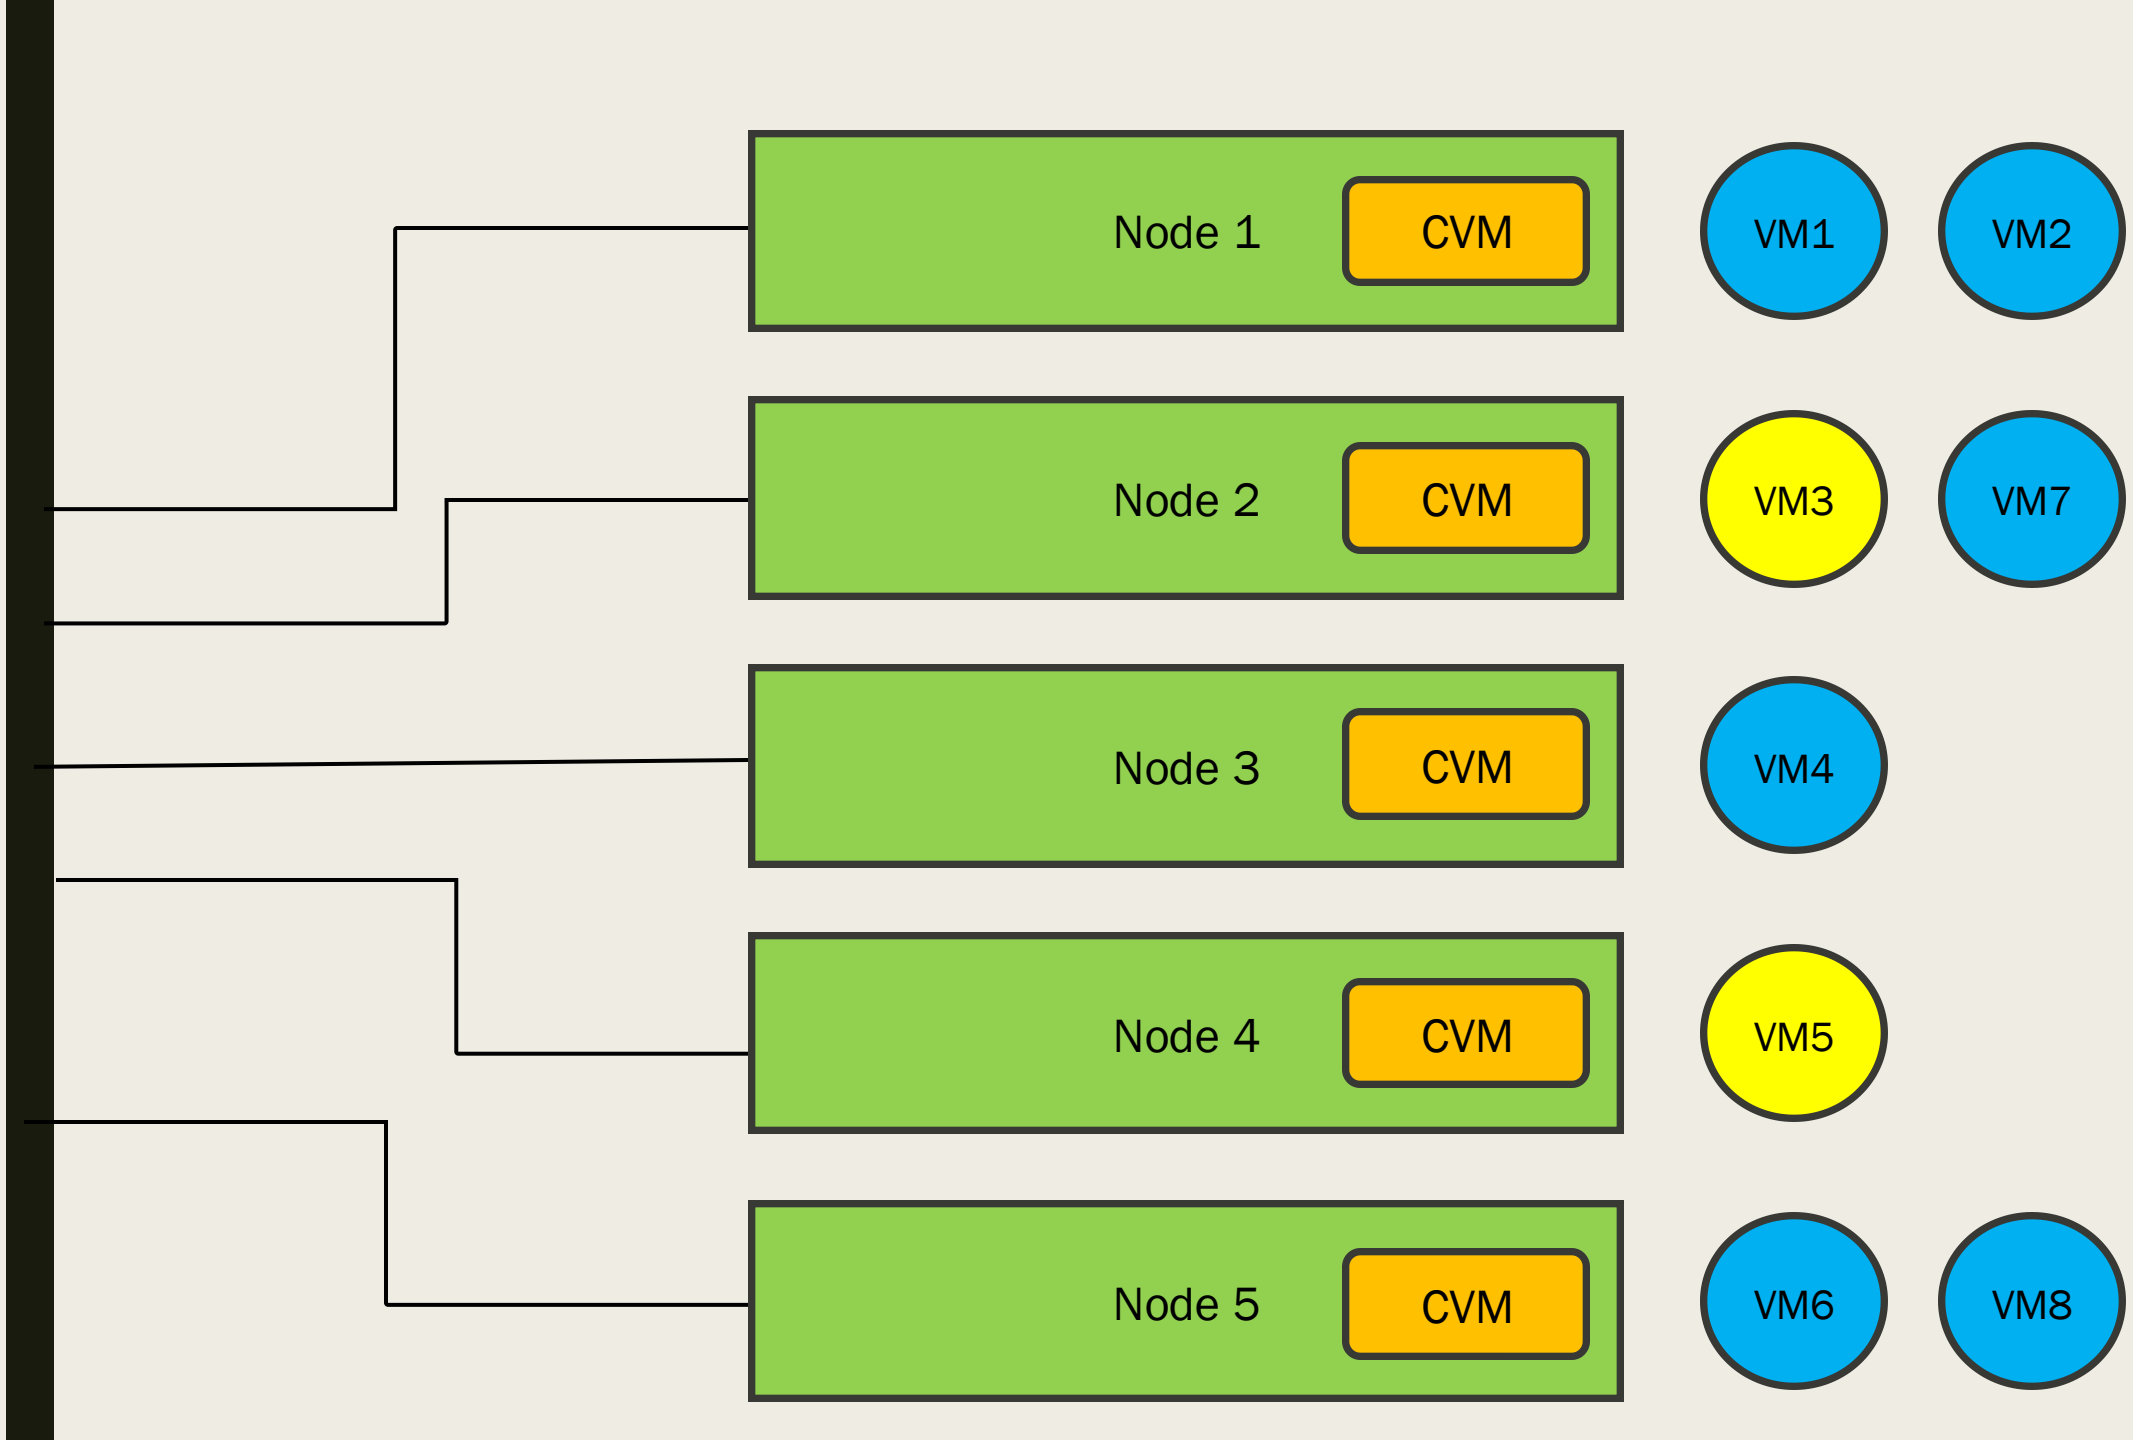
Node 1
CVM
VM1
VM2
Node 2
CVM
VM3
VM7
Node 3
CVM
VM4
Node 4
CVM
VM5
Node 5
CVM
VM6
VM8

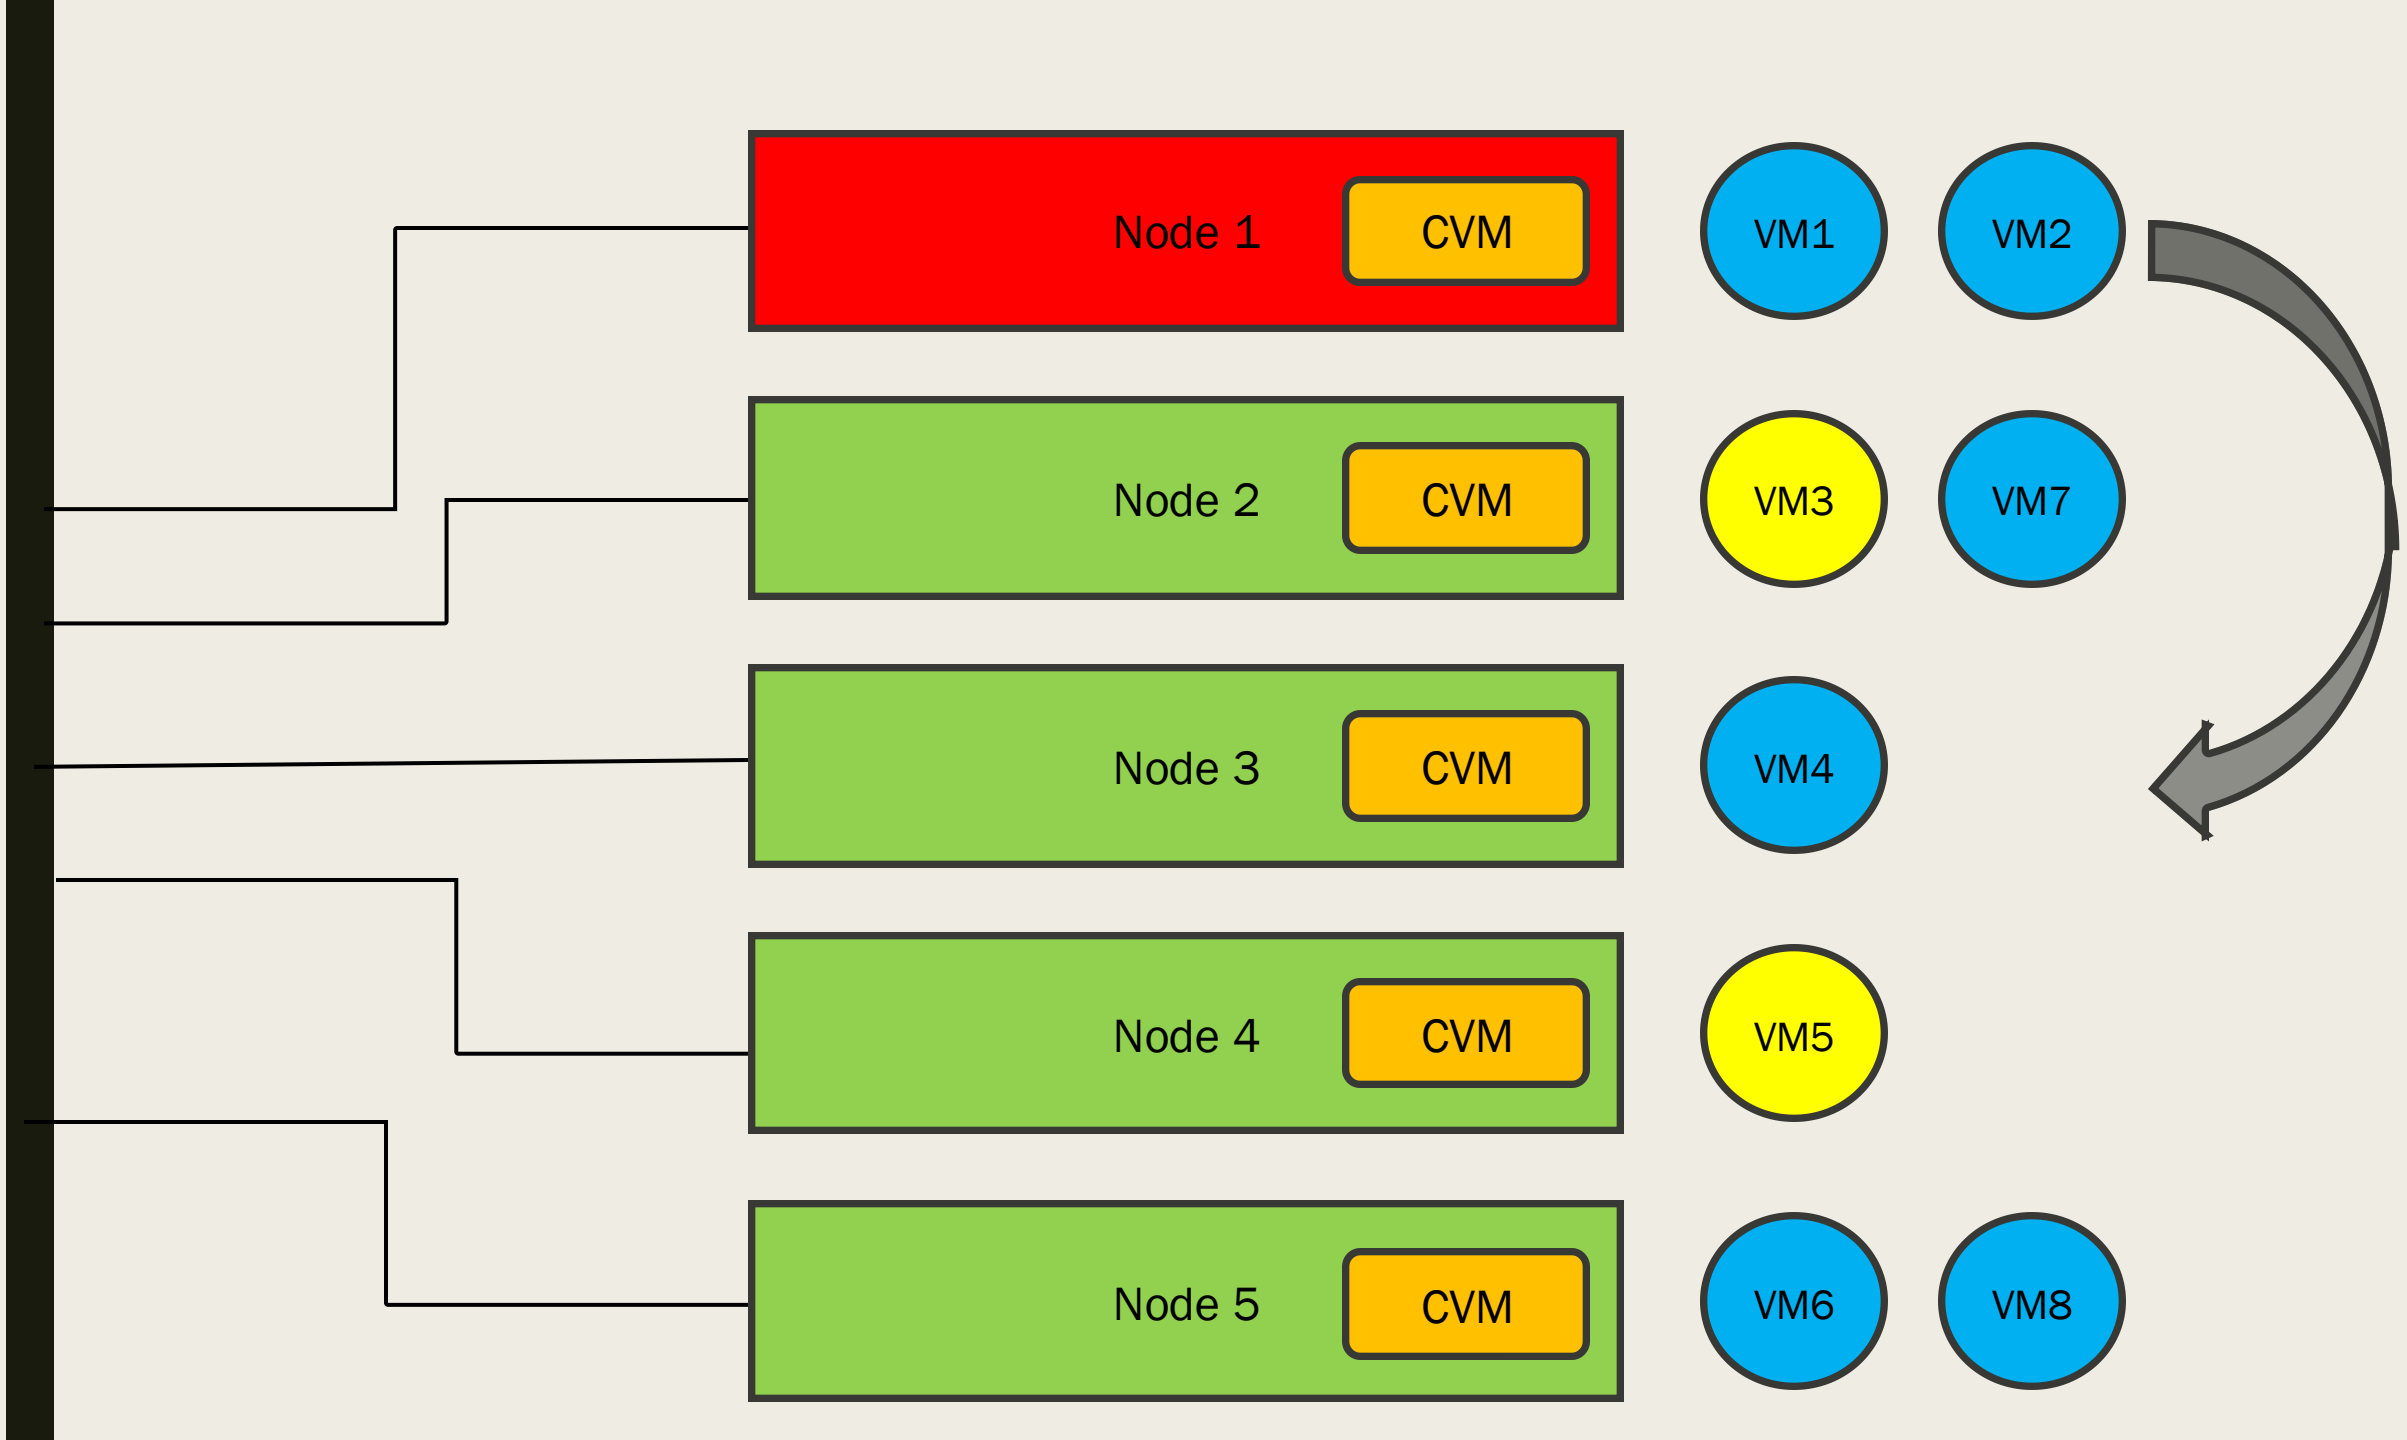
Node 1
CVM
Node 2
CVM
Node 3
CVM
Node 4
CVM
Node 5
CVM
VM1
VM2
VM3
VM7
VM4
VM5
VM6
VM8

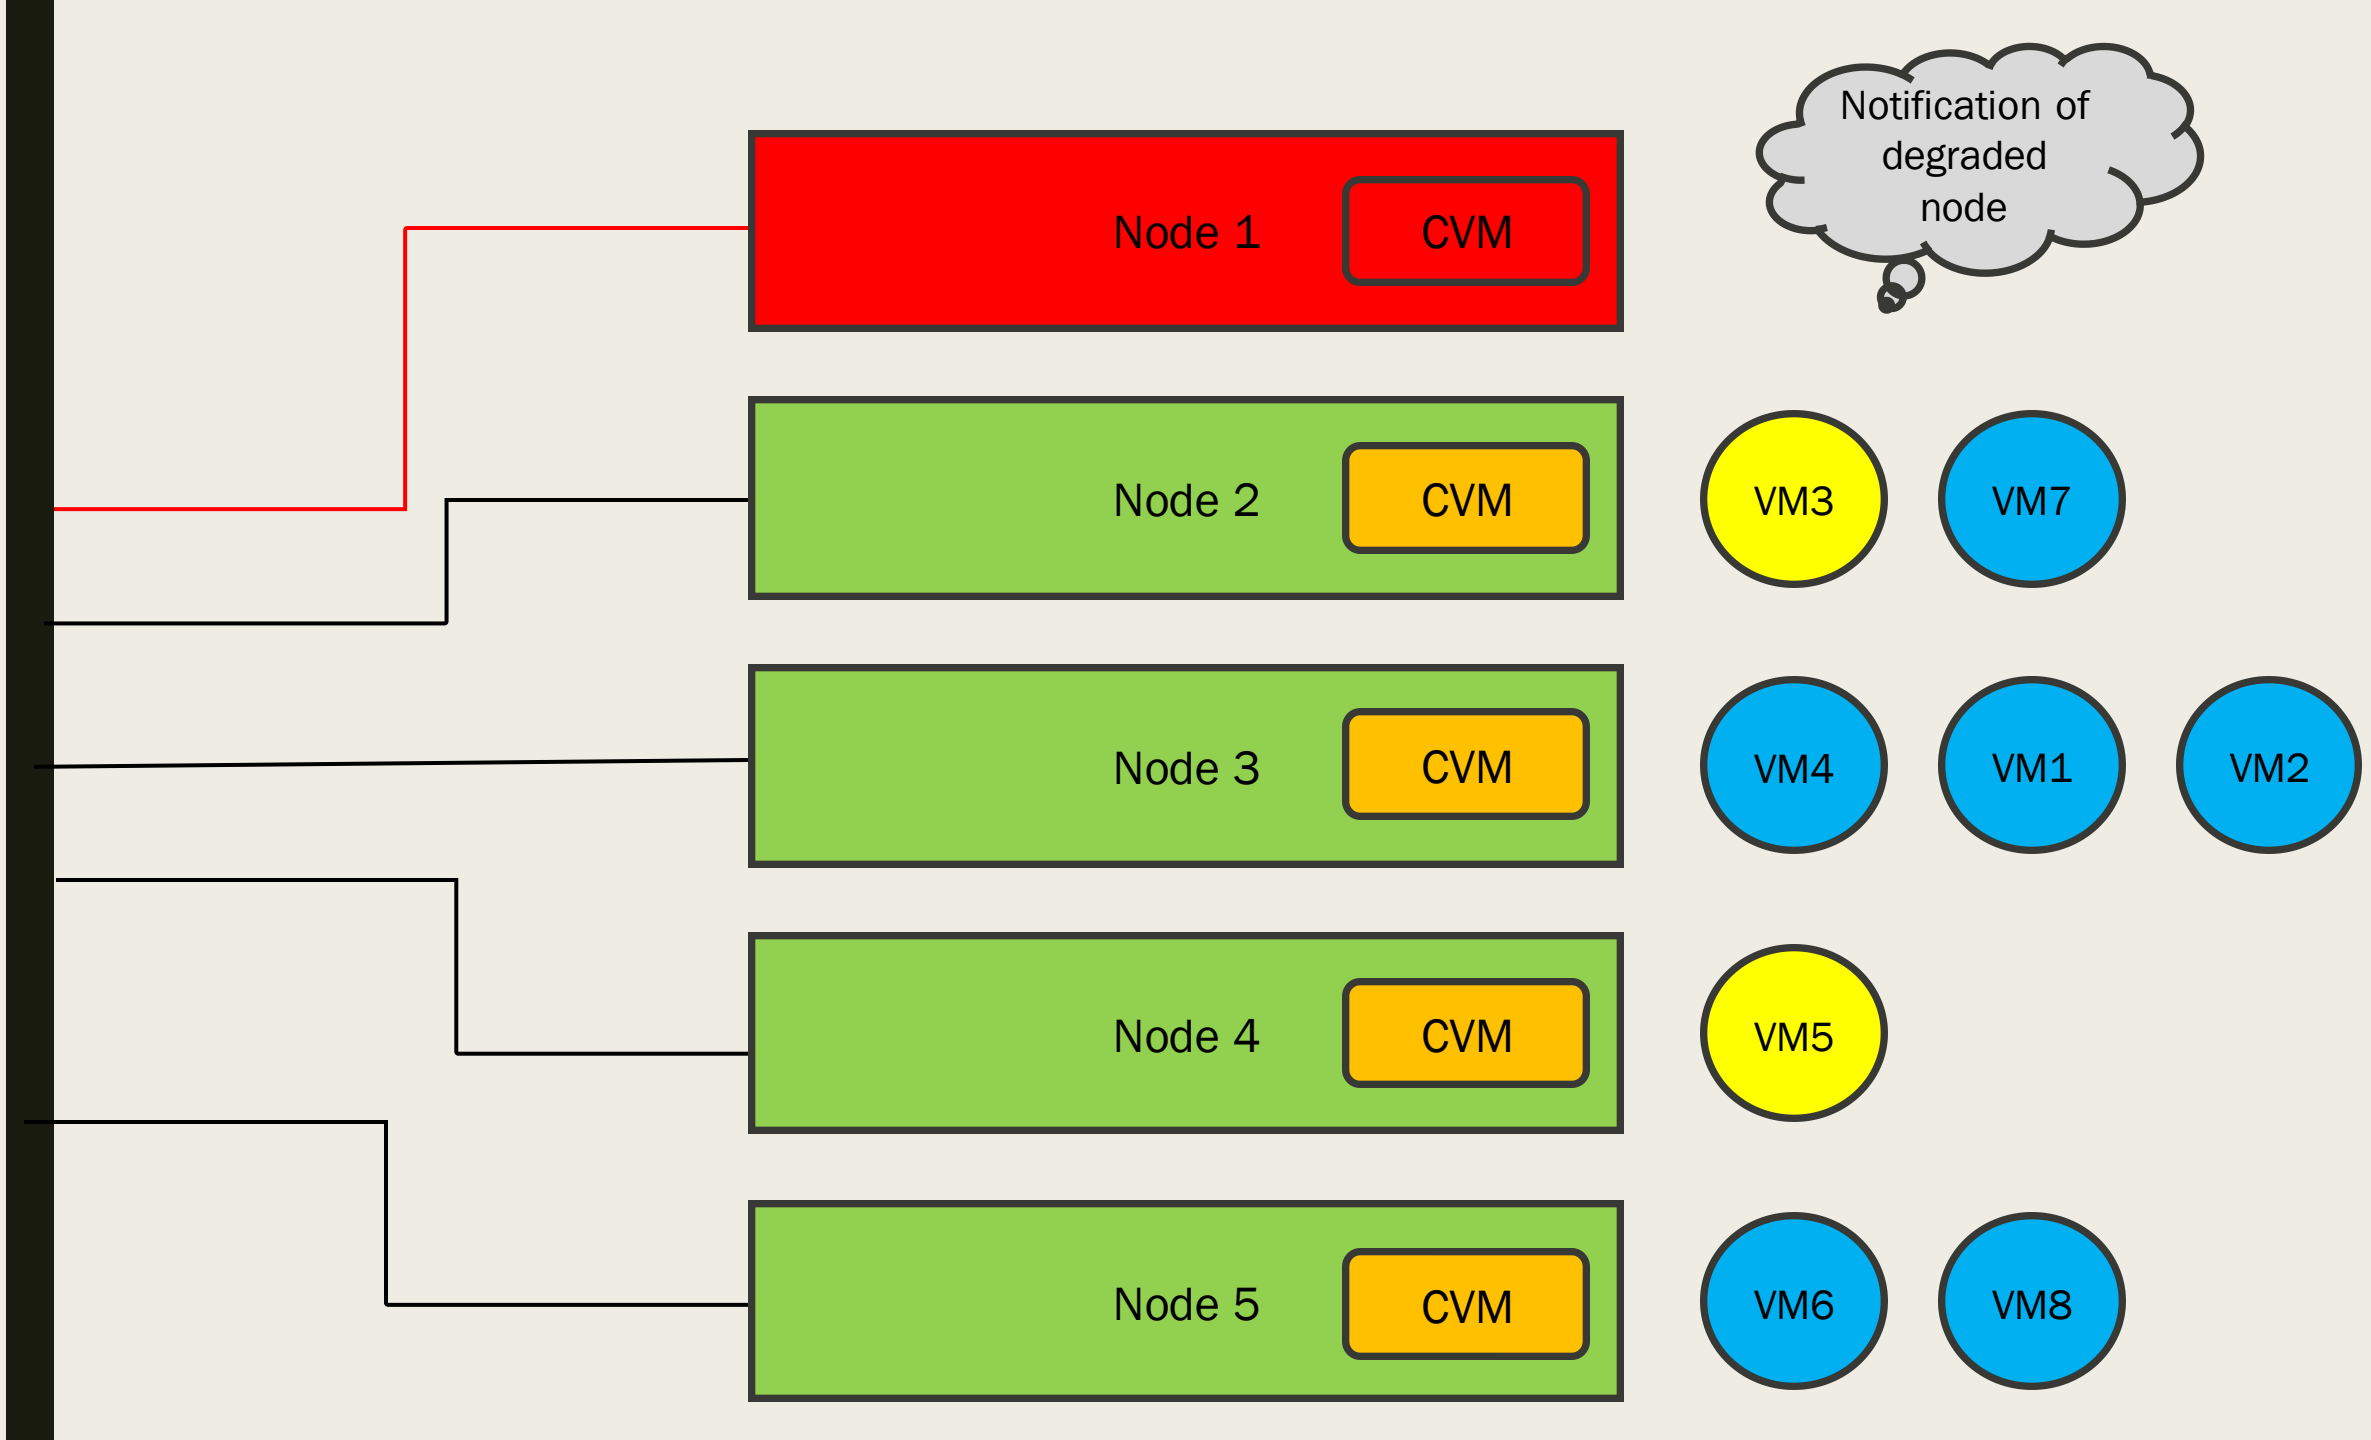
Notification of degraded node
Node 1
CVM
Node 2
CVM
VM3
VM7
Node 3
CVM
VM4
VM1
VM2
Node 4
CVM
VM5
Node 5
CVM
VM6
VM8

STN6

Hyper Converged Infrastructure (HCI)

Breakdown of composition

- 10 Physical Servers
- 8 HCI Servers, Hosting 40+ Virtual Machines
- Hyper Converged System Software
- Tape Backup System LTO9 w/Server (5) Archive Storage w/server
- 3 Rack Units (2 Base / 1 Summit)

Main Specifications

- CPU: 168 cores (8 physical servers)
- Virtual CPU cores: 672
- Memory: 3072GB or 3TB
- Storage: 375 TB
- Archive Storage: 719 TB(1 dedicate server w/ two storage chassis)

Power Consumption

- Base and Summit: 5,050 Watts

- Load balanced and resourced the proper amount of memory and CPU cores into a virtual machine based on the server's actual performance.
- Made proper use of using virtual CPU cores instead of depending on physical CPU cores.
- Creating STN6 systems to make more efficient use with the allocated hardware (memory and CPU).
- Each virtual machine that was priorly a physical server, has become more efficient.
- Higher performing CPUs, allows proper use of virtual cores rather than just the physical cores, as previously used.
- Depreciated about 20+ physical machines and turned them into virtual machines for both the Summit and Base.
- Single managed environment to support virtual environment.
- Maximize efficiency of CPU and memory and a data-redundancy, to minimize data lost and increase service recovery.

STN6 Storage, CPU, & Memory Usage Summary

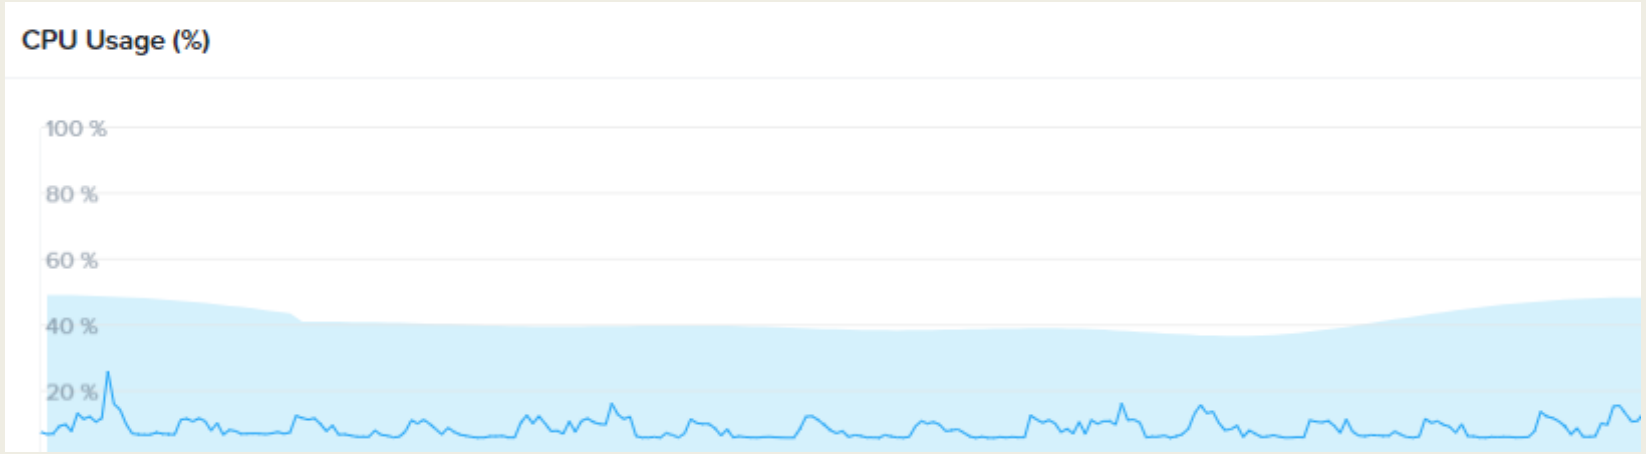

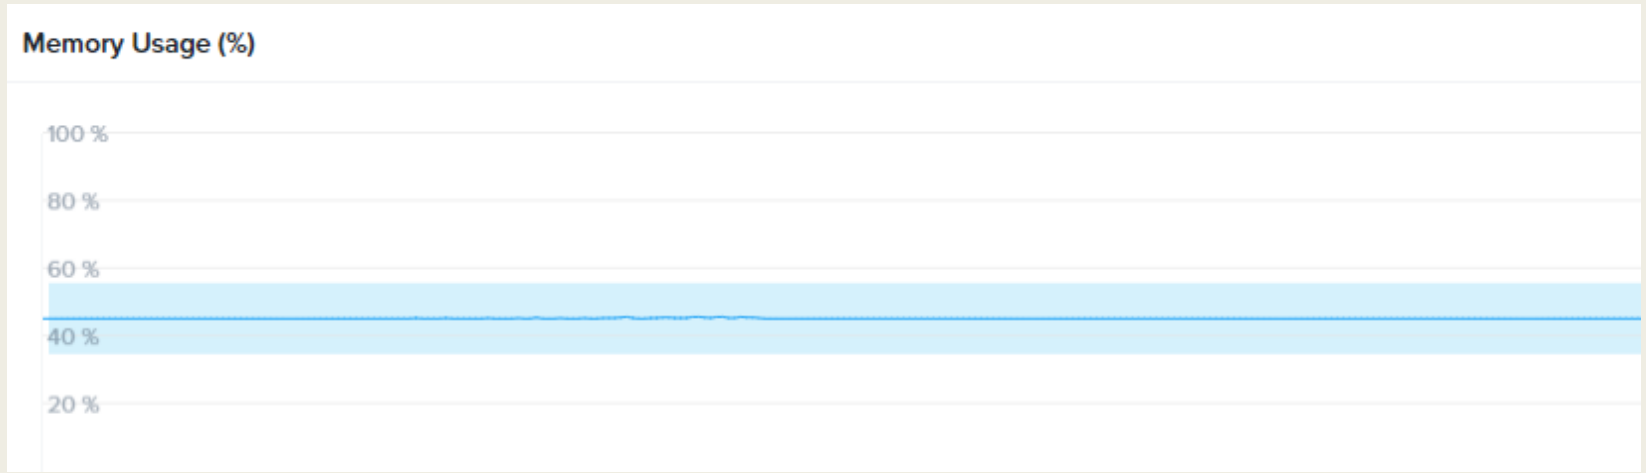

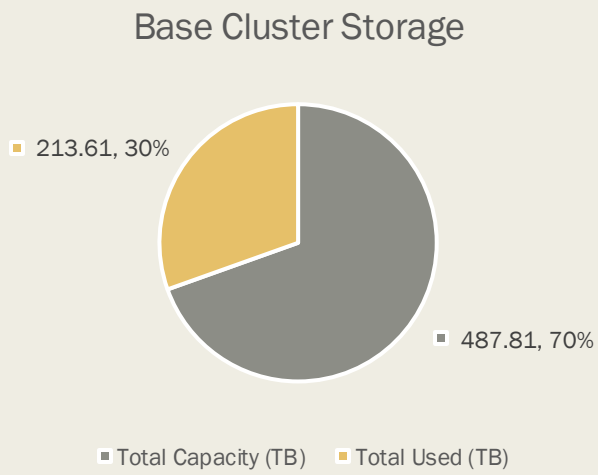

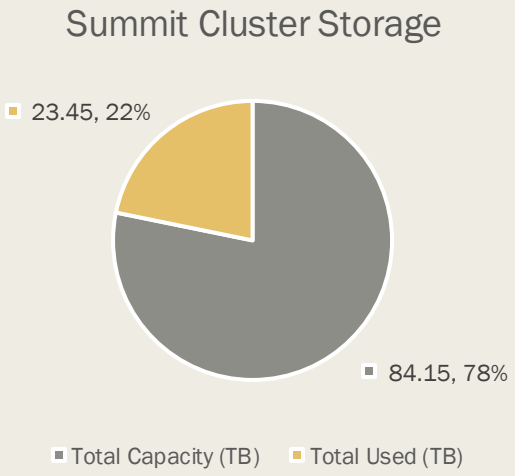

STN5 and STN6 Comparison

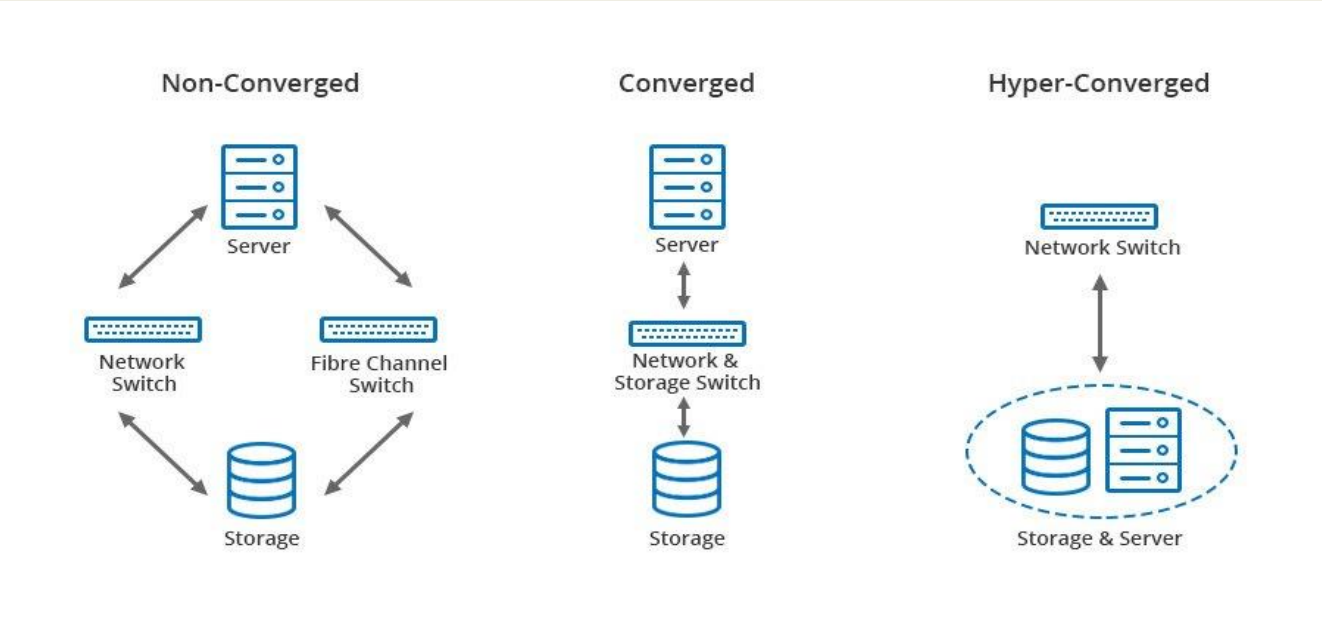

	STN5	STN6	Decreased	Increased
Physical Servers	44	10	77.27%	
Rack Enclosure Usage	5 Units (4 Base, 1 Summit)	3 Units (2 Base, 1 Summit)	40%	
Power Consumption	12,000 Watts	5,050 Watts	57.917%	
Virtual Machines	25	55		120%
Facility UPS Battery Uptime	**60 minutes**	**90-100 minutes**		50 - 66%

Key VMs STN6 Hosts and Supports

- Internal and External webpages
- AllSky Camera
- Dome Safety PLC
- Hilo Base and Mauna Kea Summit Observatory Management
- Mauna Kea Summit Facility Camera System
- Subaru Telemetry System
- Instruments:
  - MOIRCS
  - IRCS
  - HSC
  - PFS
  - Ultimate

Network

- Depreciated large chassis infrastructure at Hilo Base and Mauna Kea Summit
- Implemented a modular solution
- New core switches consists of multiple switches, in a "stacked" environment
- Core switch is managed from one central interface regardless of quantity switches/modules stacked together
- Allows ease of management if a switch/module failure was to occur
- End devices are redundant and split across multiple switches or also known as "blades"
- Implementation and use of High-Speed Networks: 25Gbps, 40Gbps, and 100Gbps physical interfaces