


Optimizing Architectures for Multi Mission Archives

Gretchen Greene


Space Telescope Science Institute




STScI Archive and Data Center



- ⊕ STScI hosts operational archives and data processing for multiple missions
 - ⊞ The big active missions: HST, Kepler, GALEX
- ⊕ MAST: Multi-mission Archive at Space Telescope
 - ⊞ Established in 1997 as NASA's Optical/UV archive
 - ⊞ Supports both active and legacy missions (IUE, FUSE, EUVE, ...)
 - ⊞ Images (DSS), spectra, catalogs (GSC2), time series
 - ⊞ Community Contributed Products - High Level Science Products (HLSP)
 - ⊞ Currently integrating the Hubble Legacy Archive (HLA)
- ⊕ JWST Science Operations Center in design and development
 - ⊞ Early production phase for the Science instrument Integration and testing Data (SID) Archive – ground test data archive utilizing MAST & HST data archive and distribution system
 - ⊞ Primary mission science data processing and archive operations



Balancing Architecture Goals




⊕ Support Mission Requirements


- ⊞ HRAS (High Reliability, Availability, Serviceability)
- ⊞ Provide capability to tweak mission dms operations priorities
- ⊞ Shared costs for multiple mission architecture support
- ⊞ Plan optimal hardware life cycles through life of mission
- ⊞ Performance management, continuous and benchmarked

⊕ Organizational archive planning for community

- ⊞ Optimize the usefulness of the FULL archive content for the next generation of archival science research – data mining
- ⊞ Design systems with flexibility for innovation
- ⊞ Integration with external archive resources



MM Architecture Planning Team



✦ HST Architecture Transition Planning Team:

✦ John Scott, Tom Walker, Steve Slowinski, Patrick Taylor, Rich Kidwell, Mike Swam, Faith Abney, Prem Mishra, Eugene Mindel, Tony Nueslein

✦ Archive Infrastructure Planning Team (AIPT):

✦ Brian Mclean, Karen Levay, Ron Russell, Bernie Shiao, Mark Kyprianou, Vera Gibbs, Sandy McCarthy

✦ JWST SOC and Systems Engineering WG (SEWG)

✦ Joe Pollizzi, Daryl Swade + WG

✦ Archive Management Team

✦ Carl Johnson, Rick White, + leads of dms work areas

✦ IVOA and VAO


Innovation

Unified Customizable User Interface – Web 2.0/Visualization

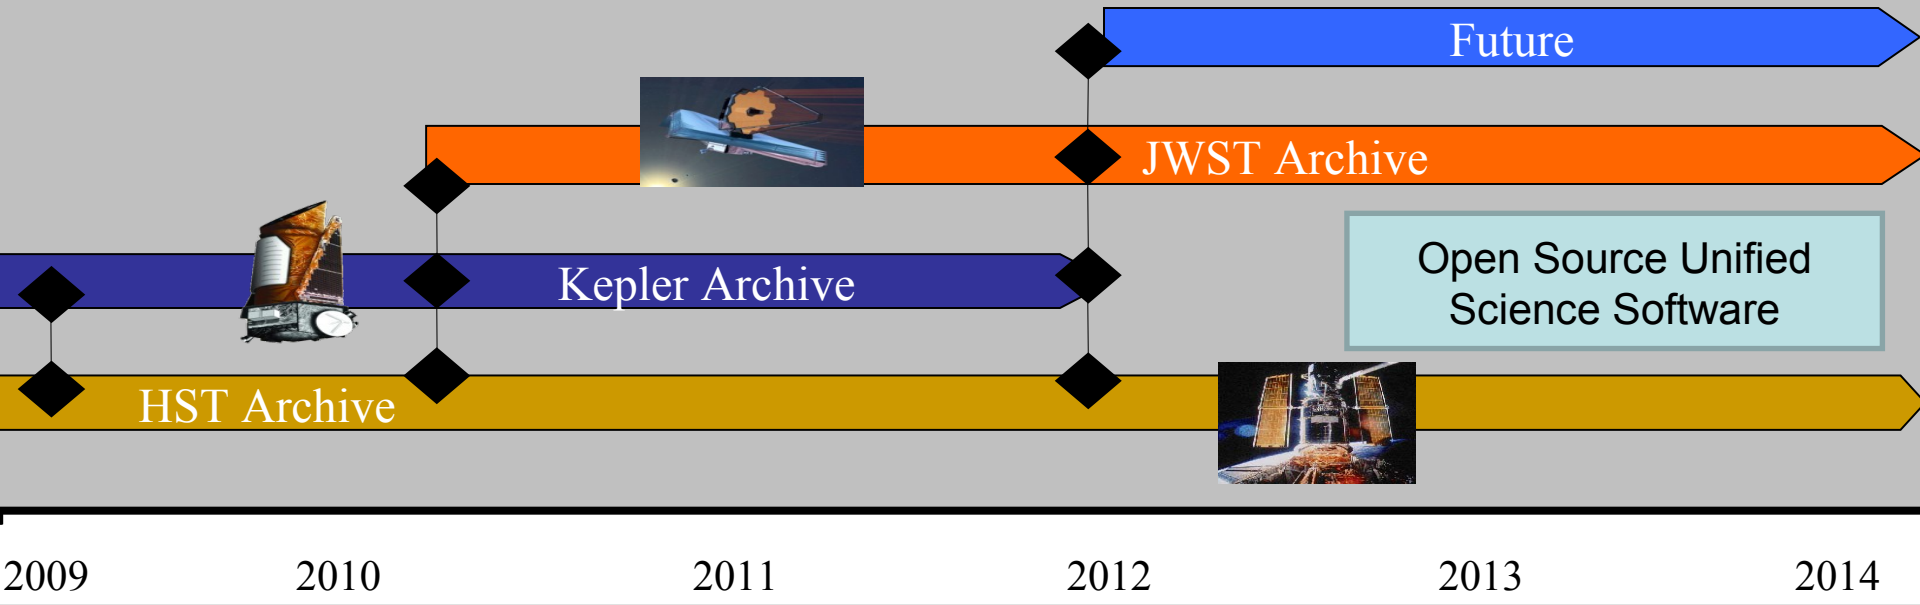
Multi Mission OPS Architecture

Virtualized Architecture

Workflow



Multi Mission Science Archive Operations



Public Science Archive




SAN Upgrade to 8Gb

Network Bandwidth

Service Oriented Architecture


Infrastructure



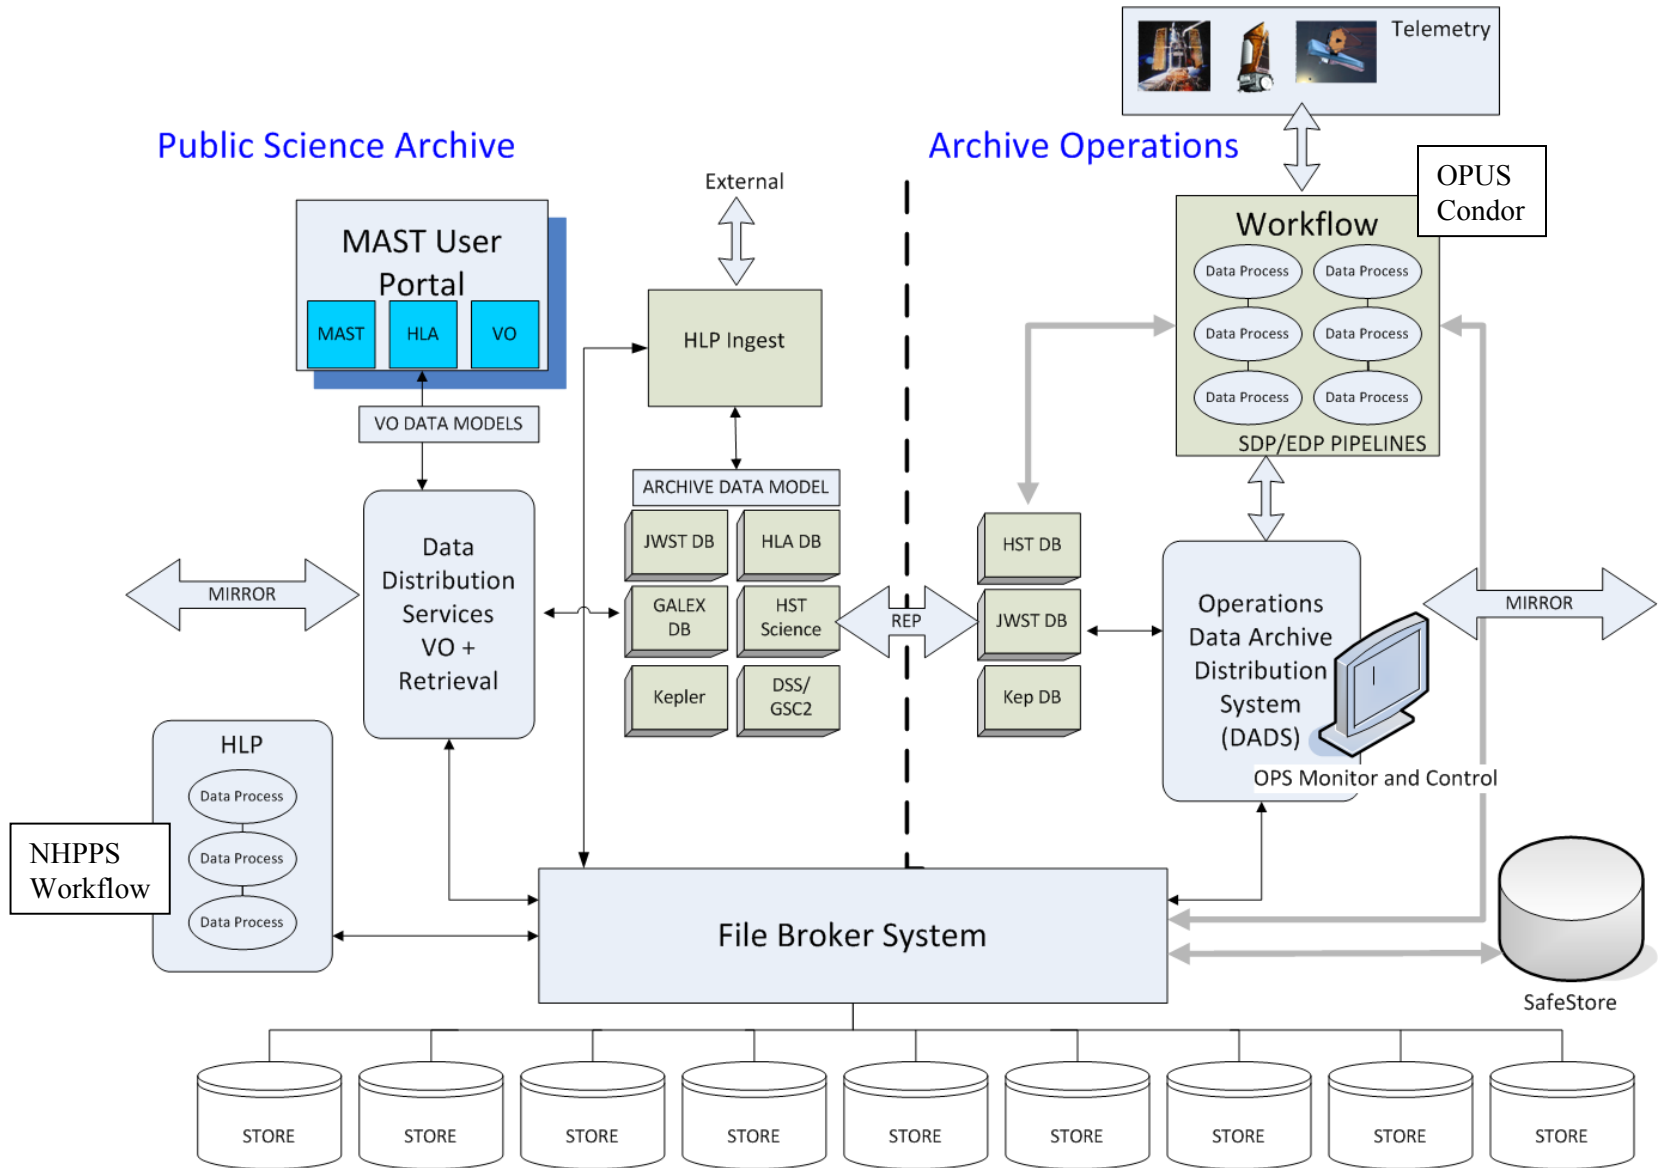

Architecture supported work areas in STScI data management systems (DMS)




- ⊕ Data receipt and processing of science and engineering telemetry data
 - ⊞ Workflow, instrument pipelines, reprocessing for calibration enhancements and higher level science products generation
- ⊕ Archive Systems for storage and distribution of raw and processed science and engineering data
 - ⊞ Operational Monitoring, file brokering, safestore, and database systems
 - ⊞ Science user interface , data retrieval from core archive, help desk
- ⊕ Science Calibration Software – Instrument team expertise
 - ⊞ Instrument pipeline steps which perform corrections , combine , and characterizes data
- ⊕ Data Analysis Tools – distributed software
 - ⊞ Instrument scientists use to characterize instrument performance
 - ⊞ End user astronomers generate and visualize results of calibrated data
- ⊕ Interface and service support for other mission subsystems
 - ⊞ Planning and Scheduling of Observations
 - ⊞ Flight Operations Systems , NASA centers




Functional Multi Mission Architecture





Architectural Tiers




❖ Interleaved system models and views


- ❑ Applications – update/modernize functionality of operational systems
- ❑ Network Infrastructure and connectivity to the community – increased bandwidth and IOPS
- ❑ Compute, database, and application servers – new technology and scaling for increased performance demand, fault tolerance
- ❑ Storage Tiers – Internal Private Cloud + External solutions
- ❑ External Mirrors – Replication of Holdings
- ❑ External Archive Interoperability – Virtual Observatory standards

❖ Big Issues to Address



- ❑ Migration of Legacy Systems – full or partial system replacements
- ❑ How mission schedules interface with technology life-cycles
- ❑ Scaling and Flexibility for Growth - choosing the right path forward
- ❑ Who is the customer?



Architecture Development Strategies




- ⊕ Optimize, innovate, consolidate and balance requirements
- ⊕ Build working teams across organizational groups that combine the science, software engineering and IT expertise to bridge communications
- ⊕ Hardware systems share where possible base-lined configurations for consistent deployment and maintenance
- ⊕ Build dedicated systems to address the requirements for development, test, and operations environments
- ⊕ Combine architecture solutions....where possible physical and system configurations for common functional systems
- ⊕ Scalable infrastructure with end-to-end level monitoring
- ⊕ Engage resource sharing through collaborative partnerships

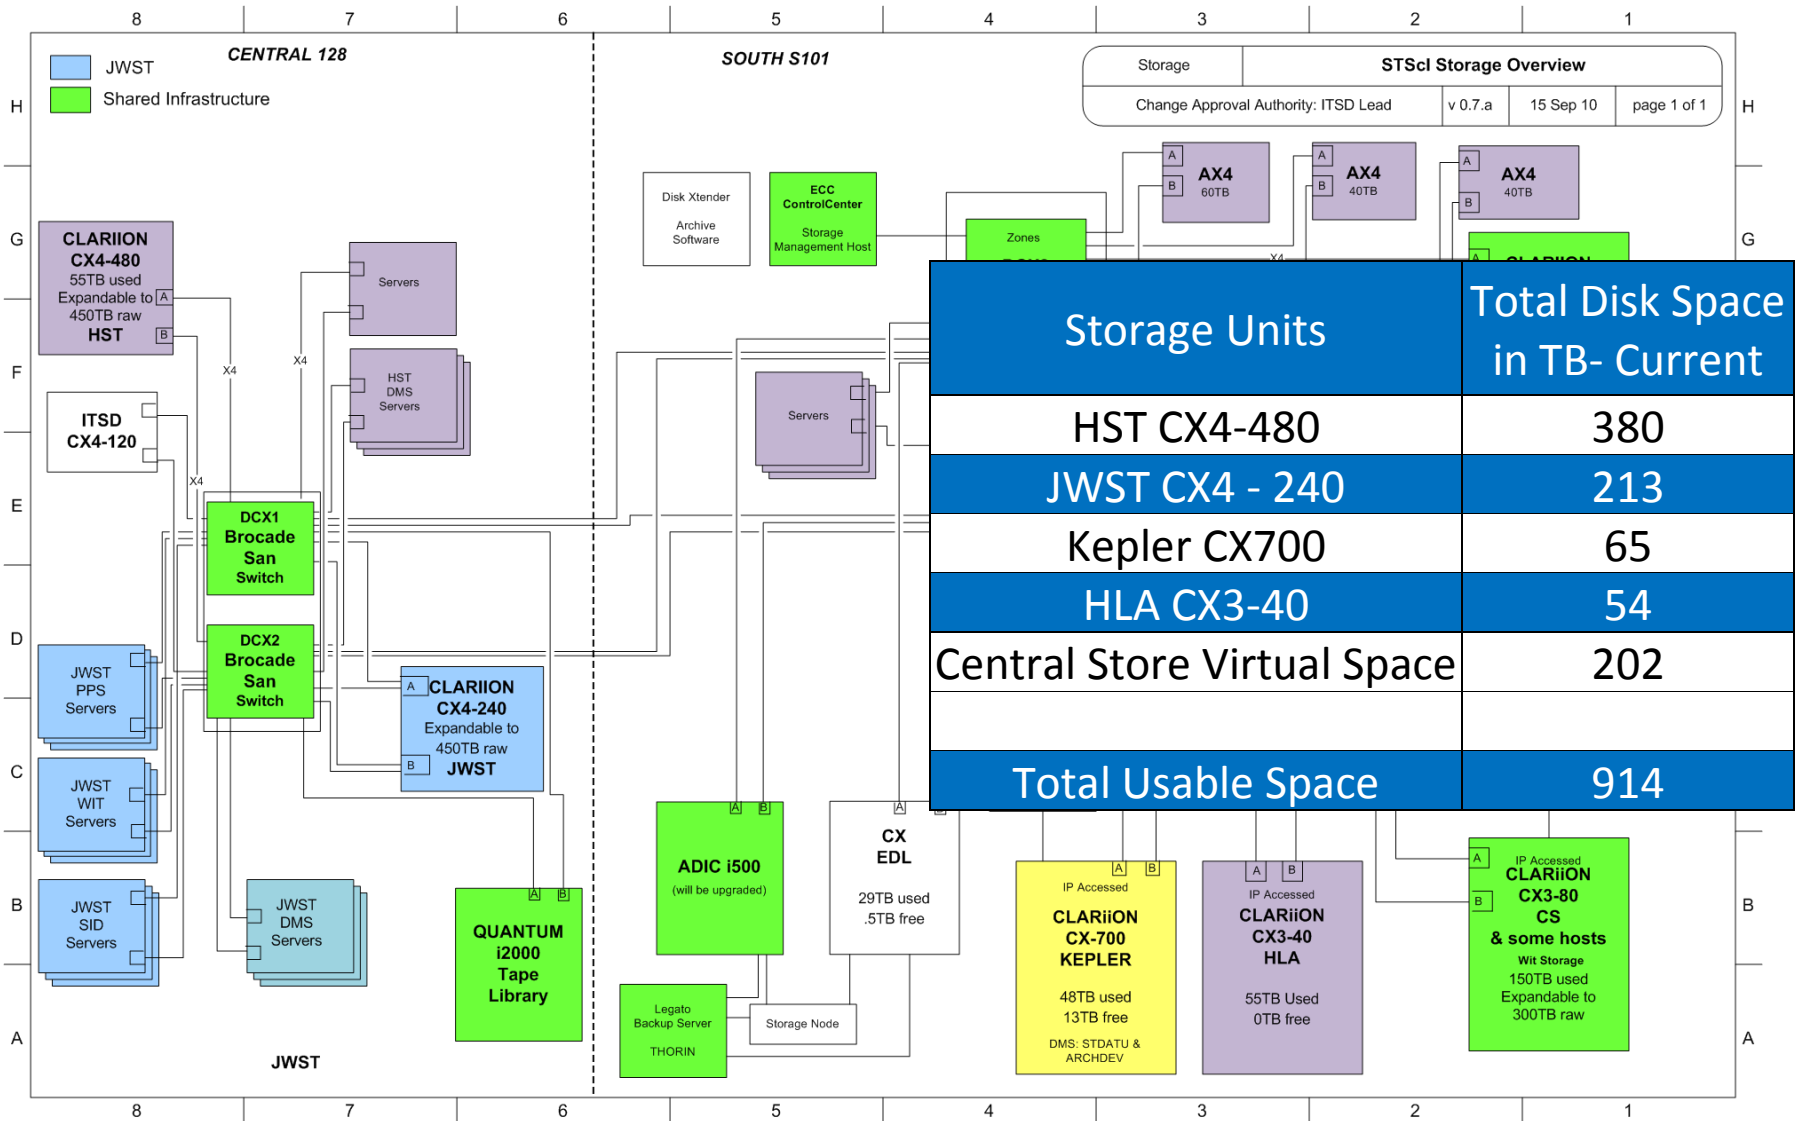




Storage Architecture - Private Cloud Storage Area Network(SAN)

- ❖ SAN shared infrastructure and EMC storage frameworks - 8Gb fiber channel network – commodity integration support cross platform
- ❖ Connectivity of host servers through Host Bus Adapters (HBA) to dual enterprise class switches with multiple fabrics
- ❖ Tiered storage: Flash, SATA (files), and Fiber channel (DBs) drives
- ❖ Storage File Brokering Systems will be adapted to manage both batch and direct file access
- ❖ Developing Total Cost for Ownership (TCO) model understanding




SAN Architecture







Applications architecture considerations



- ⊕ Design applications which function across distributed servers
 - ⊞ Load balance core mission functions: ingest, request processing, and reprocessing
- ⊕ Modernize web technology for the archive user portal services and operations control systems
- ⊕ Increase development of database services to reduce management of meta files – use data models and mapping
- ⊕ Shared software and collaborative exchange throughout systems (internal and external)
 - ⊞ Public Science Archive and Operations Archive
 - ⊞ Produce and use *open source* software – AURA science software effort
- ⊕ Key interfaces between archive and workflow systems requires shared file system architectures
 - ⊞ Application design and file system solutions (GFS-2, NFS, Others)



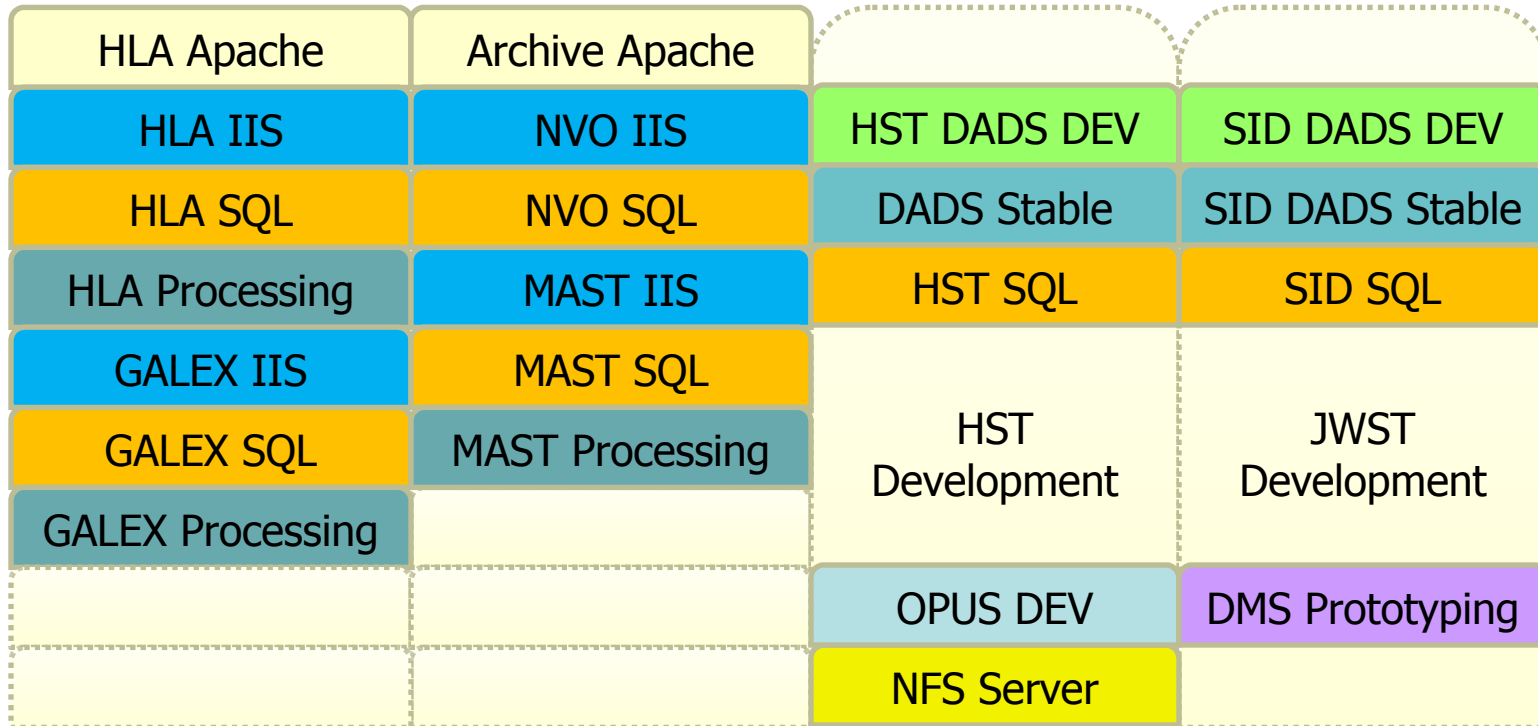
Virtualized Development Infrastructure




Public Science Archive




Mission Operations





VMware ESX cluster – Expand as needed



Realized Virtualization Pros





- ⊕ Ability to map multiple server os environments to single physical node, multi-functional purposed
- ⊕ Migration and sharing of VM server deployments can be 'motioned' across physical nodes for load management
- ⊕ Storage can be presented across ESX servers to facilitate migrations of servers without data copy
- ⊕ Tools work well for automated physical to virtual system configuration migrations and vice versa
- ⊕ Fully configured VM server distribution capability for external locations – OS, DB, APP
- ⊕ Ideal regression testing system for multi-platform



Realized Virtualization Cons

- ⊕ VMWare cpu allocation limits – slow tech change
 - ⊞ Testing demonstrated VM does not scale for HT cpus, leads to underutilized processors > 4 core
- ⊕ VMWare ESX os limits storage lun allocation to 2TB, for large scale configuration increased complexity for abstraction layers
- ⊕ Raw device mapped luns may appear on servers not hosting VMS – potential for data corruption
- ⊕ VMs have no awareness of SAN HBAs and therefore cannot use SAN backup
 - ⊞ Potential bridge of gap with VEEAM still in testing, additional licensing required
- ⊕ Licensing and costs for implementation/core



Public Science Archive


Note 1 : Test/Processing hardware identical to Public systems
(a) Test environment same as Public (b) act as backup

Note 2 : Clusters will start as failover and transition to load balancing

HLA Footprints, GALEX,
Science Catalogs, NVO




Windows cluster



LINUX cluster
LOAD Balancing


HLA/ Archive




LINUX cluster

PUBLIC

HLA Footprints, GALEX,
Catalogs, NVO
Test & Processing

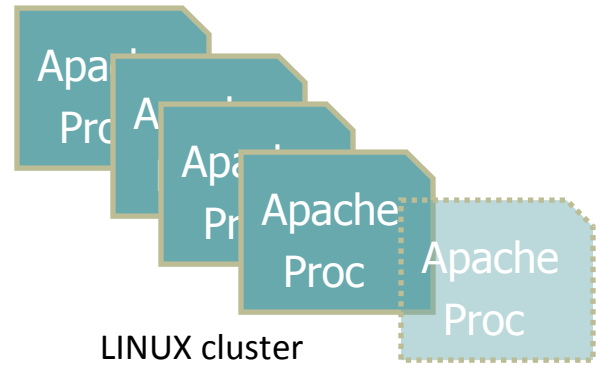


Windows cluster
Add Servers AS NEEDED




LINUX cluster
LOAD Balancing

HLA/Archive
Test & Processing

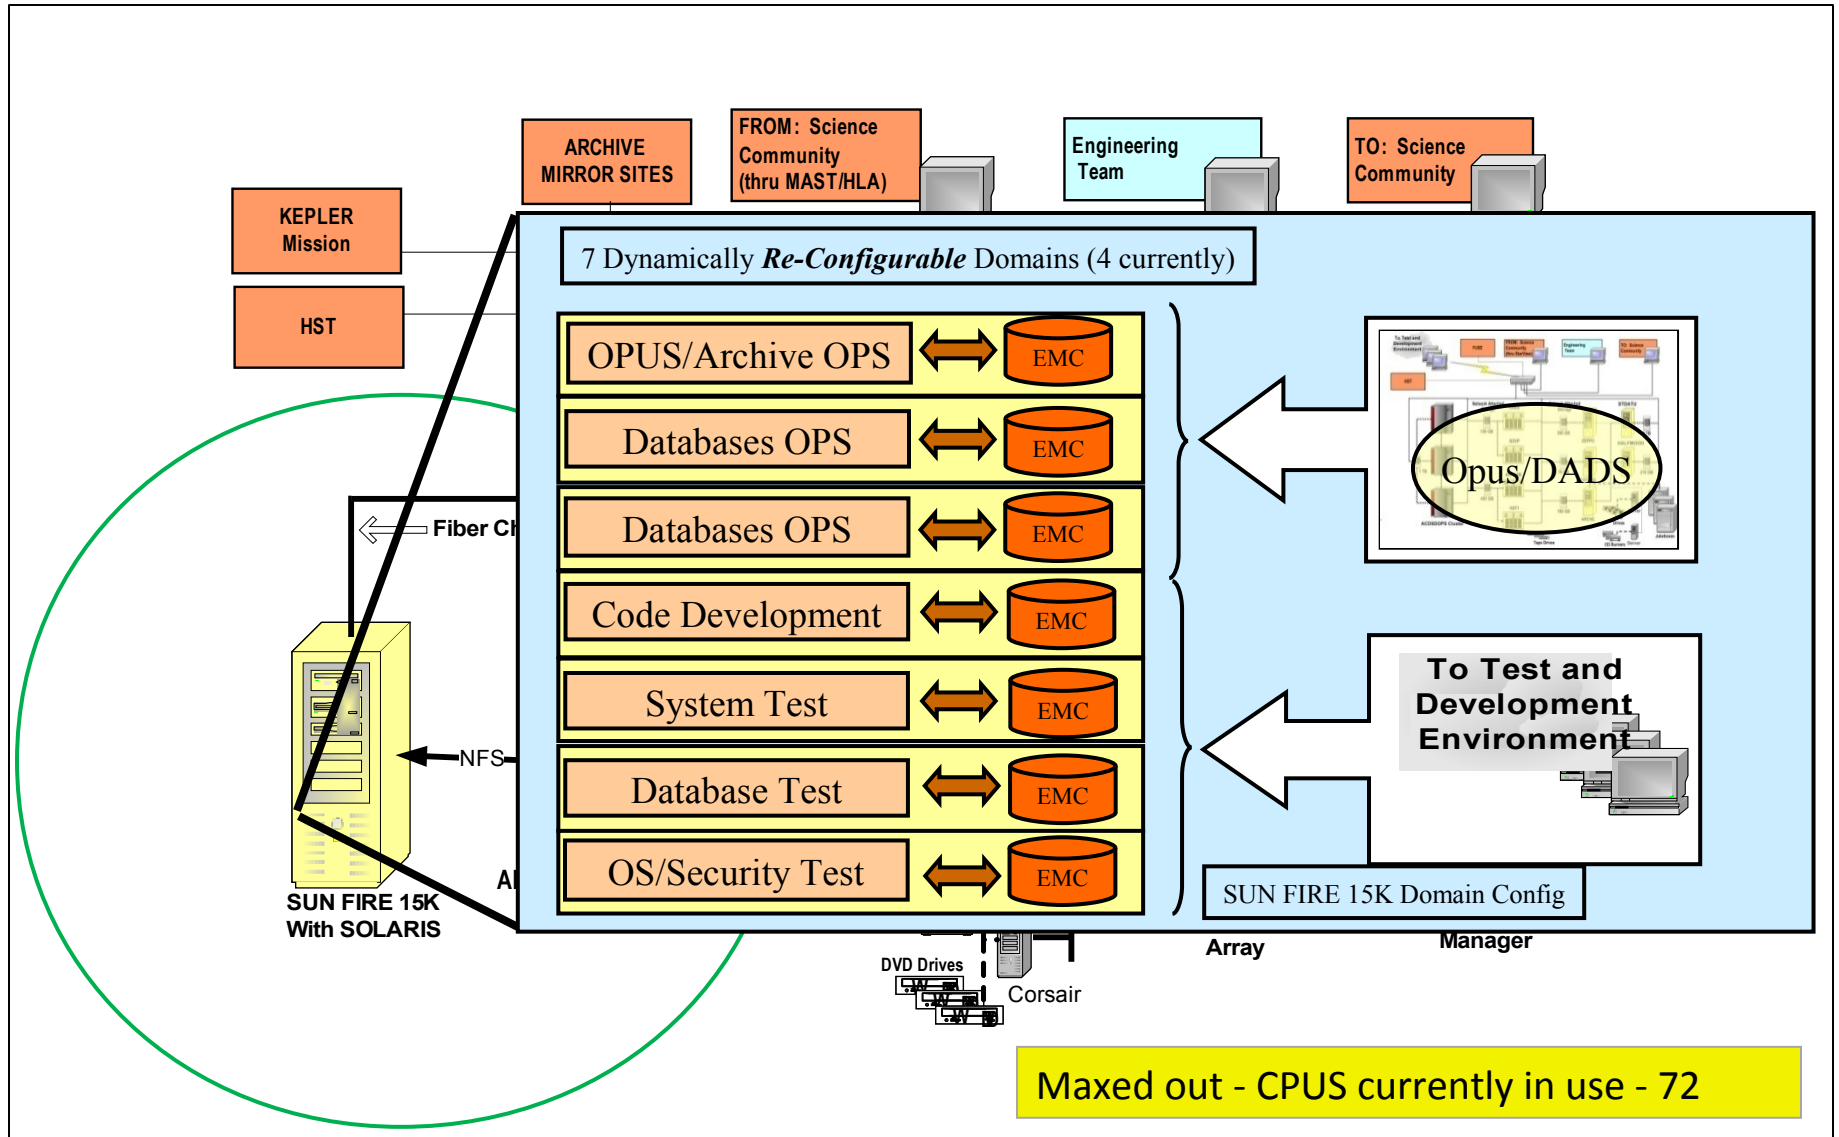




LINUX cluster
Add Servers AS NEEDED

TEST

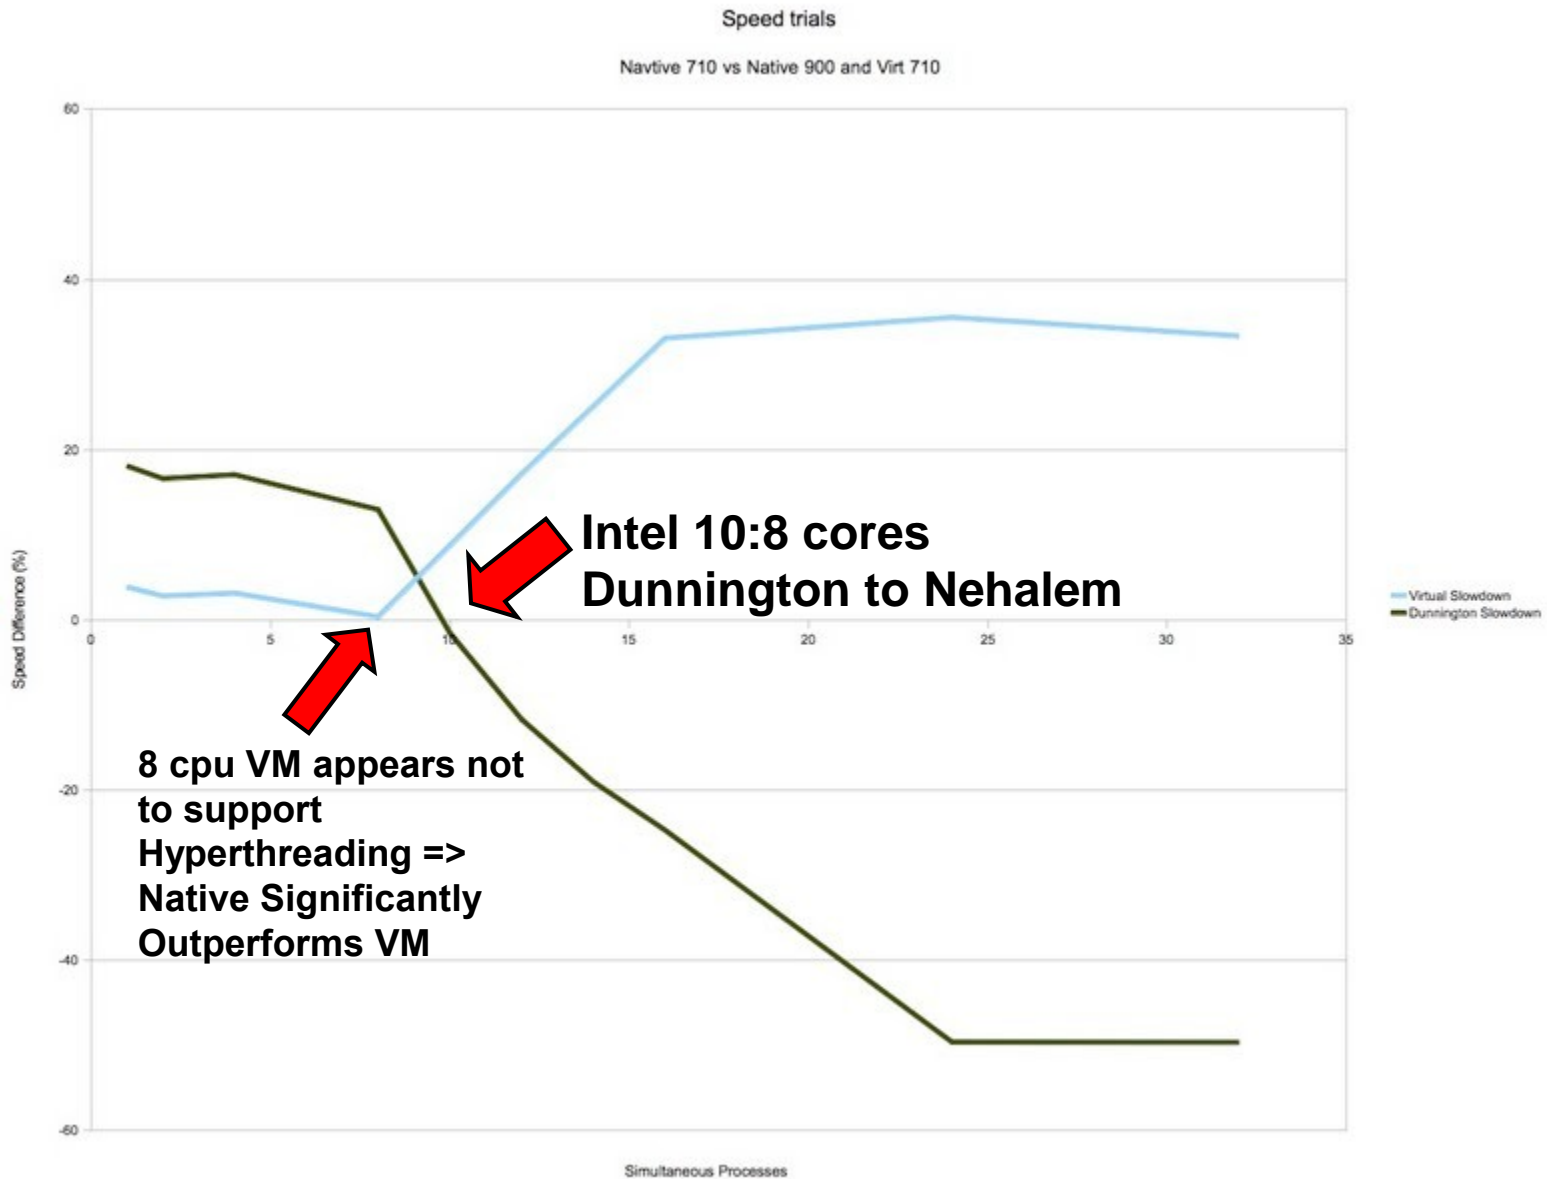




Operation Archive Architecture Today with Modernized Safestore






Performance-processor testing



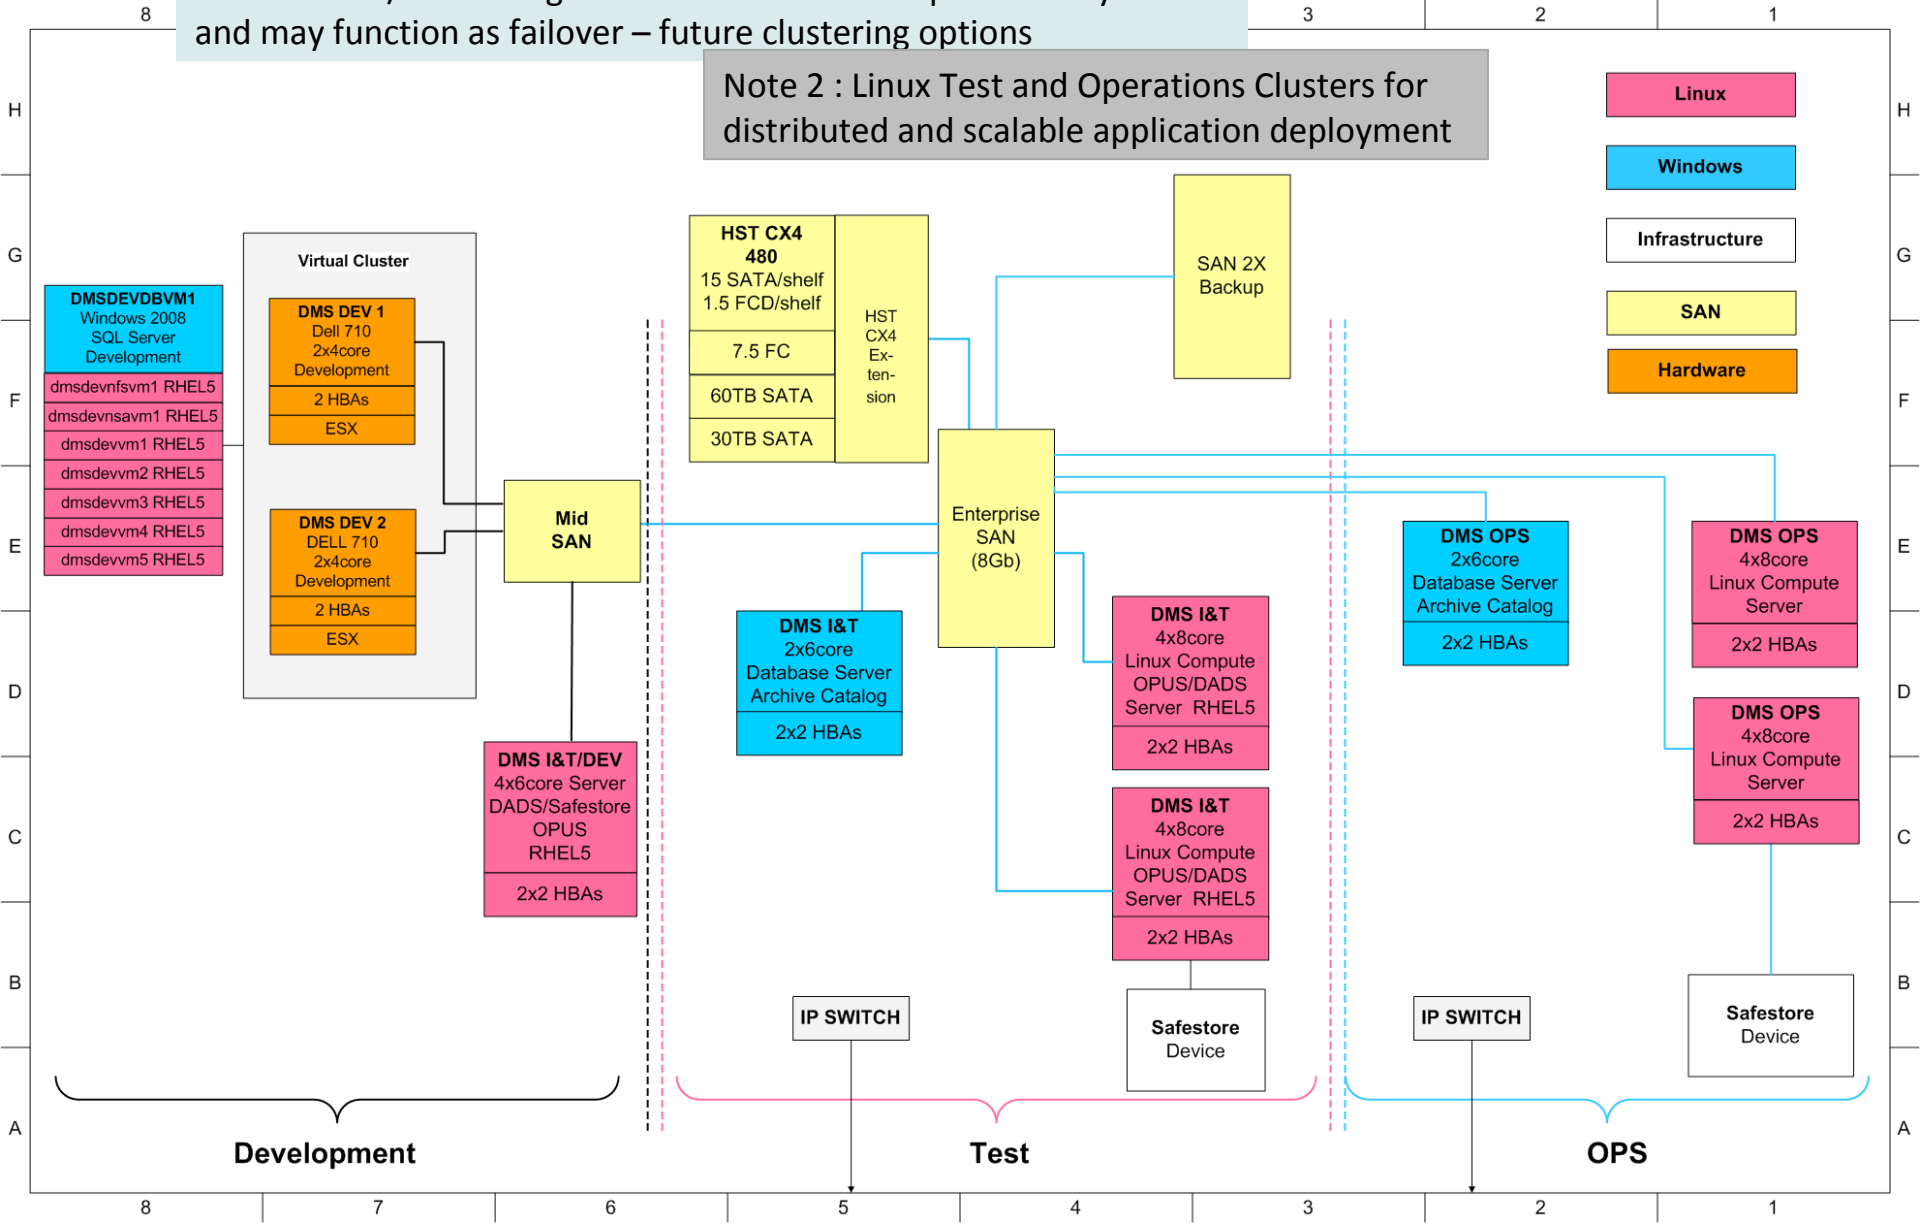




New Operations Server Architecture



Note 1 : Test/Processing hardware identical to Operations systems and may function as failover – future clustering options


Note 2 : Linux Test and Operations Clusters for distributed and scalable application deployment



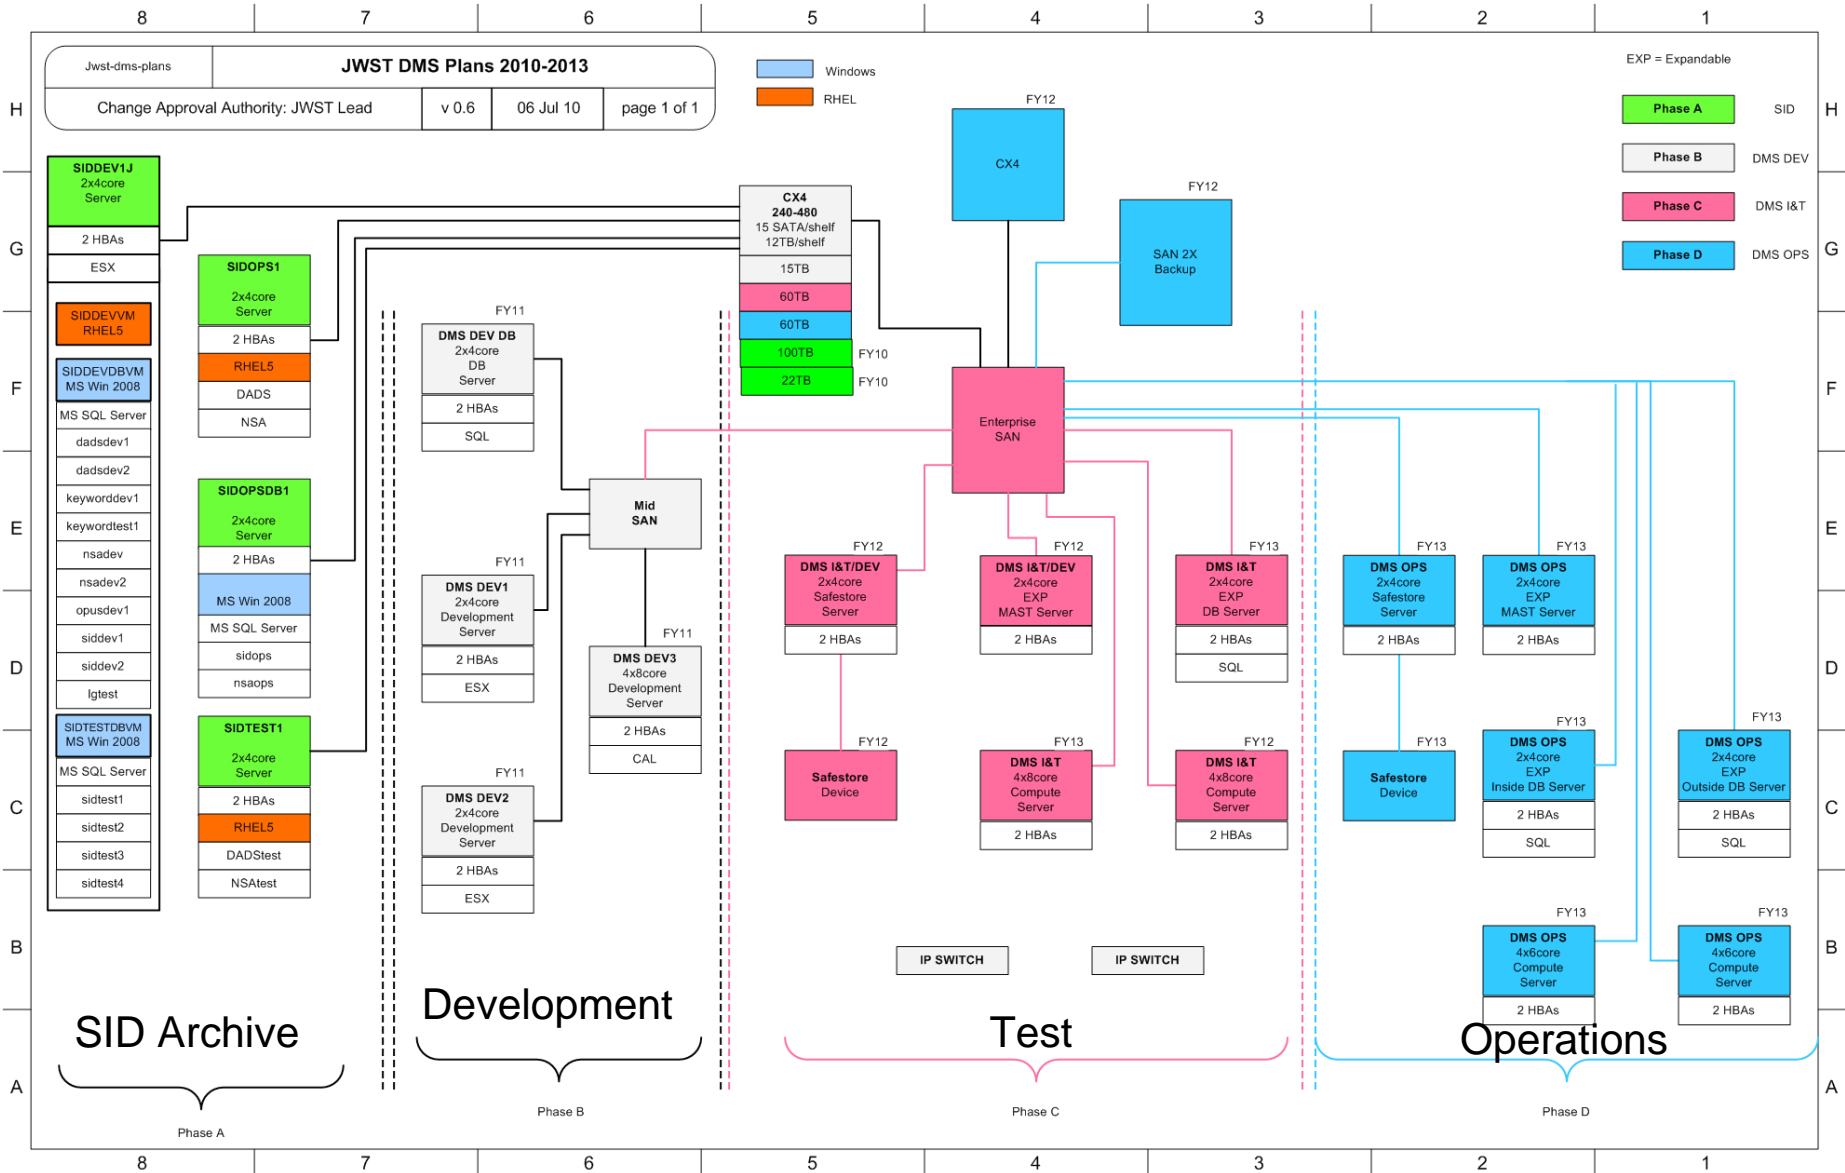




JWST DMS architecture ramp up

- ✦ Science instrument Integration and Test Data Archive (SID Archive)
 - ❑ Ground Test Data for JWST detector testing, instruments, and formal observatory test facility
 - ❑ Virtualized system architecture
 - Initially we planned for dev, test, and ops, yet discovered storage limitations for VMWare allocations to SAN pools
 - ❑ Production system online - scalable to 200TB archive
- ✦ DMS design phase beginning
 - ❑ Data Processing, Archive Systems, Architecture and Calibration Software
 - ❑ Capacity of Holdings estimated 1-2PB for base calibration




JWST DMS architecture






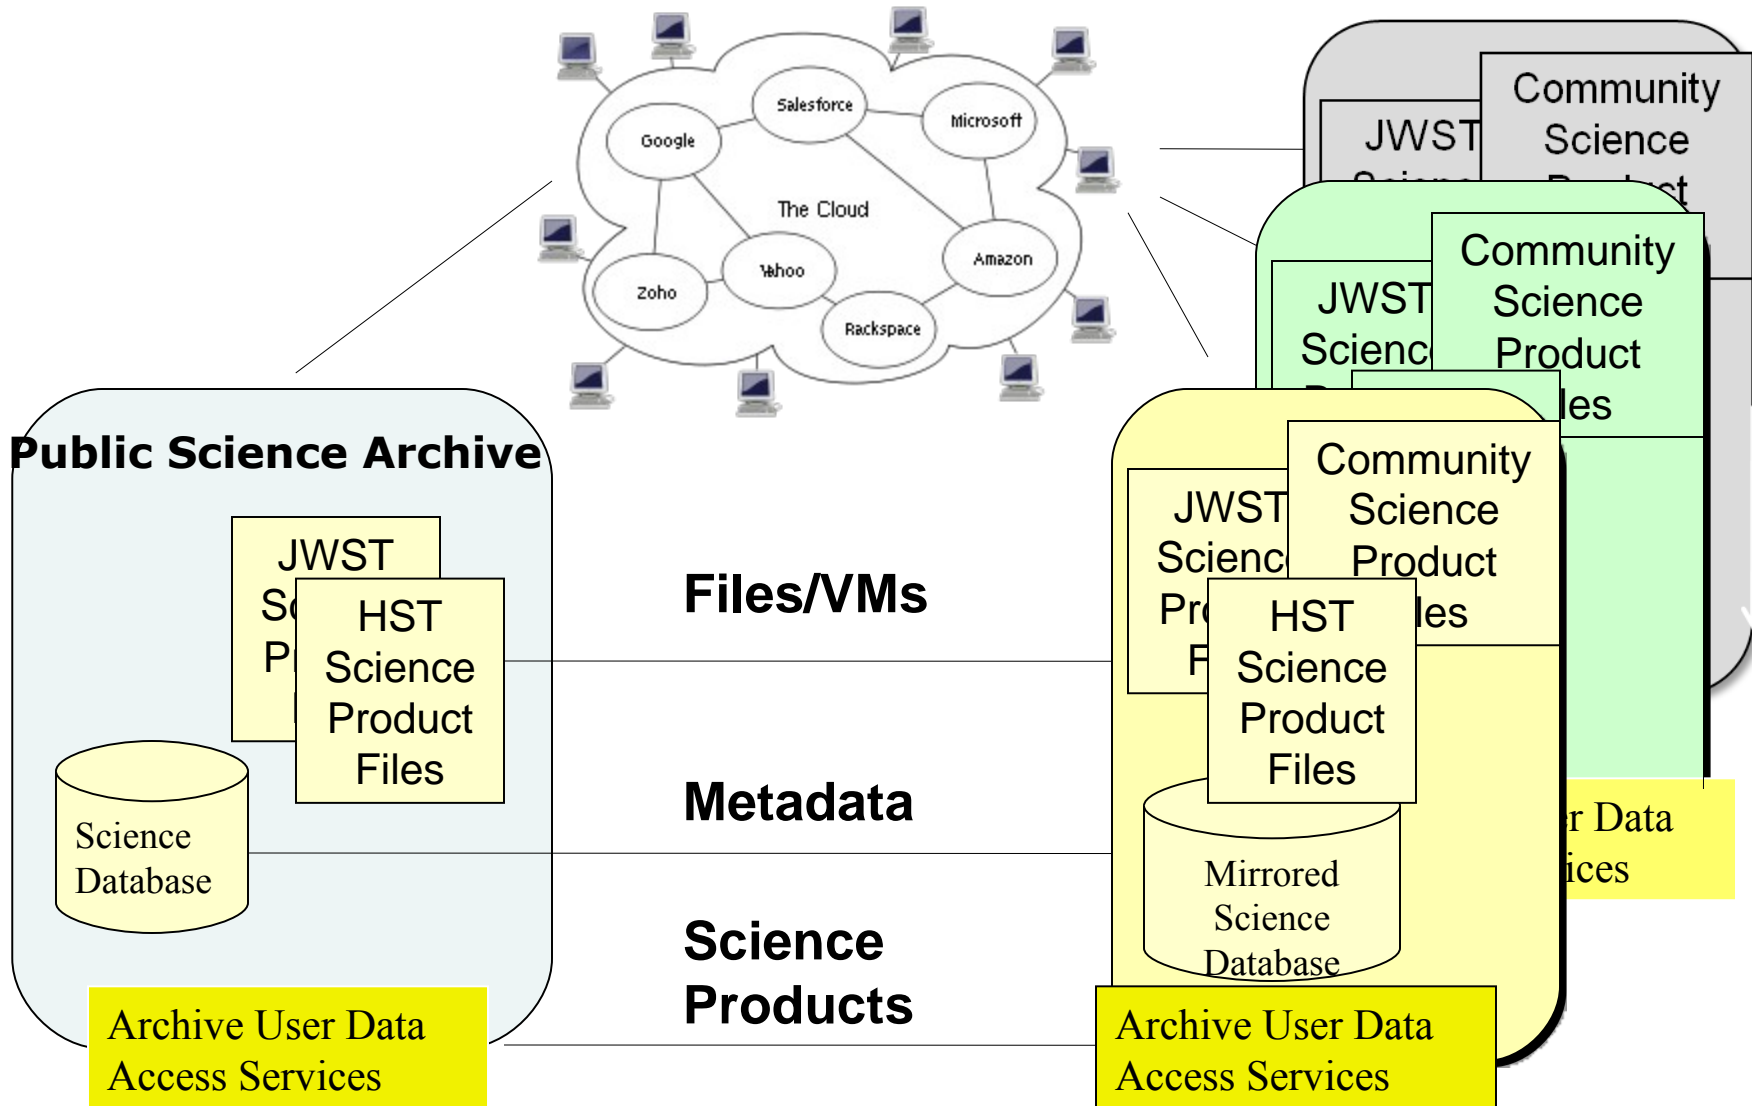

External Architecture Collaboration




- ⊕ Mirroring STScI archive holdings
 - ⊞ Evolving models lean toward Public Science Archive cached calibrated data
 - ⊞ Metadata exchange may be simplified by adopting VO data models
 - ⊞ Considering External Cloud solutions for bandwidth
- ⊕ Catalog cross-matching projects
- ⊕ Synergy with other mission archives
- ⊕ JWST archive collaborations are still TBD
- ⊕ Archive Interoperability, innovation, and technology for enabling scientific data mining across distributed networks




Future External Archive Shared Architectures





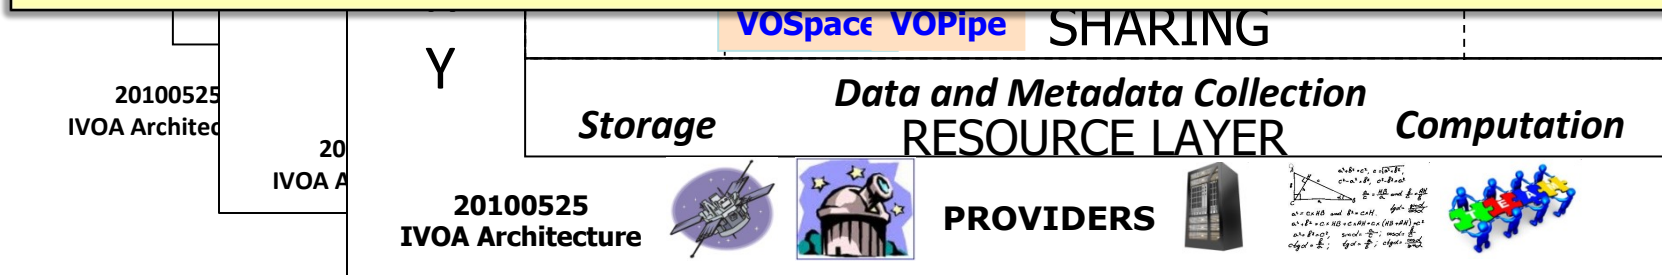
Mapping to IVOA Architecture - Technical Coordination Group (TCG)




- STScI Public Science Archive VO Components:
 - VAO Fully Searchable Registry
 - Data Provider Services (TAP, ObsTAP, Cone, SIAP, SSAP)
- In development
 - Footprint Service Specification
 - Data Models
 - VAO Work Areas: SEDs, Search Portal, Testing, Operations
 - VOTable Validation Service
 - VOEvent for supernova in Multi-Cycle Treasury ongoing
- JWST formal requirements for DMS to implement IVOA standard data provider Services


REC
Progress

D
P
A
T
A
P
R
O
T
O
C
O
L
L
S





Conclusions



- ❖ Multi mission architectures require shared solutions across the strata of archive systems
- ❖ Virtualized solutions work well for dev and distribution
- ❖ Common yet flexible infrastructure is key for scalability, performance and maintenance
- ❖ Synergy both with internal and external solutions provides opportunity for advancing science return
- ❖ Long term planning for technology life cycle peaks ~five years for refresh of architectural components – build systems that can adapt in units rather than full replacement