



# Facility Services

Glenn Cooper / Scientific Computing Facilities Department

2023 FCRSG

15–16 February 2023

# CSAID/SCSS/SCF Department

- Shared, often underlying, services:
  - System administration, hardware + OS: install, monitor, fix, retire computing equipment for other services: compute, storage, databases, ...
    - Coordinate with data center services, network services
  - Server & disk storage purchases: budget, specifications, procurement process
  - Software container infrastructure, currently OKD/Kubernetes
  - Underlying tools: monitoring, configuration management, others; used by us and by other service providers
- Other services used by experiments:
  - Interactive computing (“gpvm” and CMS LPC interactive nodes)
  - OKD
  - Jenkins CI/CD service

# Linux future: AlmaLinux

- Fermilab and CERN recommended CentOS Stream in October 2021, but ...
  - 5-year support cycle proved unpopular
  - Some initial issues with broken updates for Stream
  - Rebuilds announced in 2021 (Rocky Linux, AlmaLinux) seem stable, have good support
  - CMS offline (cmssw) chose AlmaLinux
- Revised Fermilab/CERN recommendation for AlmaLinux in December 2022
- Relevant dates:
  - Scientific Linux 7 support ends June 2024
  - Alma 8 support ends May 2029
  - Alma 9 support ends May 2032; we encourage moving directly to 9

# Jenkins CI / CD


- Central dashboard to configure, run, and see status:

S	W	Name	Last Success ↑	Last Failure	Last Duration
✓	☀	<a href="#">mu2e_muse_ci_build</a>	2 hr 22 min <a href="#">#396</a>	N/A	45 min <a href="#">▶</a>
✓	☀	<a href="#">novasoft-git_SRT_slf7_release_build_output</a>	6 hr 37 min <a href="#">#47</a>	N/A	7 min 19 sec <a href="#">▶</a>
✓	☀	<a href="#">novasoft-git_SRT_slf7_release_build</a>	8 hr 15 min <a href="#">#97</a>	N/A	1 hr 38 min <a href="#">▶</a>
✓	☀	<a href="#">decisionengine_pipeline_el</a>	11 hr <a href="#">537#master</a>	2 days 11 hr <a href="#">529#master</a>	3 min 41 sec <a href="#">▶</a>


- Hooks to local repos, GitHub, etc.
- Linux build nodes

# Interactive computing

- **GPCF** (interactive General Physics Computing Facility; “gpvms”)
  - Used for a variety of purposes by most experiments (other than CMS)
- **CMS LPC interactive** nodes: also VMs, using same tools as GPC



GPCF nodes over time

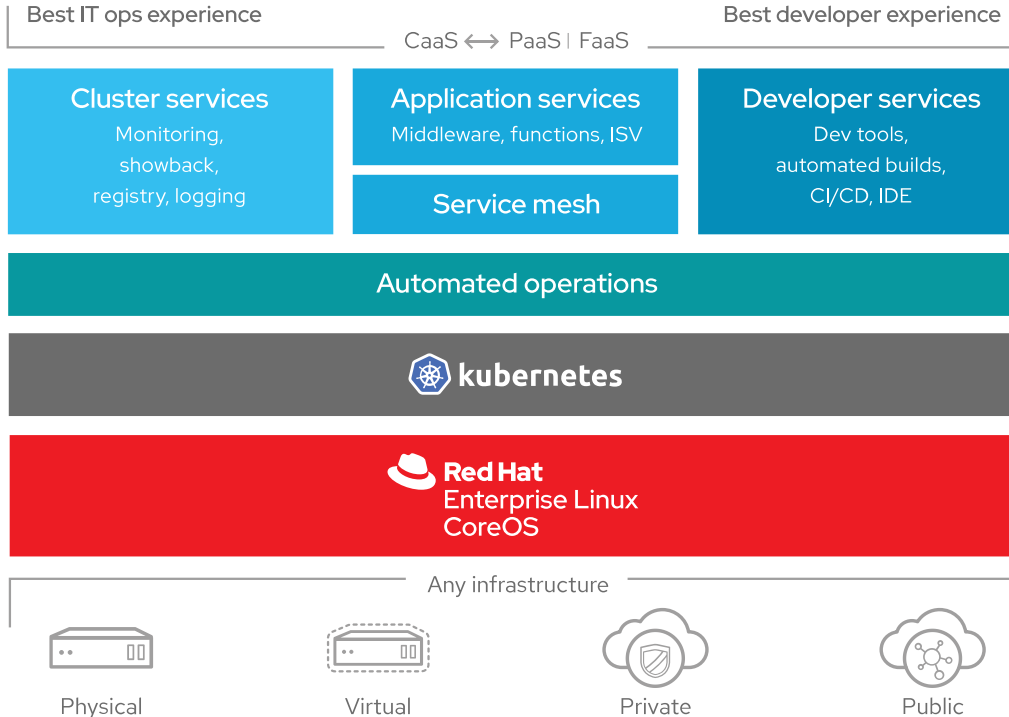


CMS LPC nodes—load average

- Familiar, well established; more supportable than large numbers of physical nodes
- Less flexible than, e.g., elastic analysis facility

# Container infrastructure

- Running on OKD (open-sourced Red Hat OpenShift):  
**multi-tenant container ecosystem**



# Container infrastructure

Used for:

- SCSS-supported **services**:
  - Rucio
  - DUNE hardware database
  - FTS3
  - Others
- Elastic **analysis facility** <https://analytics-hub.fnal.gov/>
  - JupyterHub, BinderHub                      Coffea-Dask
  - Access to GPUs                                Triton Inference Server
  - Burst to batch farms                         Various other applications
  - Now in production (though offerings and scope are still growing)

# Upcoming challenges

- Upgrade ~3000 systems from SL7
- Prepare for HL-LHC, DUNE, Mu2e, others
- Continue progress toward containerization, isolation from base OS
- IPv6 mandate
  - 20% IPv6-only by end of FY23
  - 50% IPv6-only by end of FY24
  - 80% IPv6-only by end of FY25
- Rapidly evolving hardware landscape: GPUs, FPGAs; CPU architectures (ARM, Power, x86 innovations); new SSD types and other storage; etc.
- Longer lead times for purchases: regulatory environment; supply chain (improving but still a factor)