

Scientific Computing Facilities

January 10, 2023

Major Incidents

Date	Service	Scope	Comments

Announcements

- Third-Wednesday maintenance next week, 1/18:
 - Kernel updates/reboots on most interactive nodes

Of Interest

- Increased VM-to-VM and pod-to-pod network performance in OKD by adding more network-handling threads, and going to jumbo frames (settings in several places).
- Normal requests for InCommon certs are still broken. Able to get certs for now using a temporary workaround provided by Jeny.
- Various batches of new hardware have arrived recently, are arriving now, and are due soon. See below for updates.

Off-hours Pages

Upcoming Change Requests

Incident Summary

11 incidents opened in last 21 days: 9 closed, 2 in progress.

Request Summary

5 requests opened in last 21 days: 3 closed, 2 in progress.

Other Activities

- OKD
 - Improved network performance
 - Transfer from physical node to physical node was ~80Gb (~93+ with jumbo frames), as expected for 100GE connections
 - Increased VM to VM performance from ~12Gb to ~65Gb (~75Gb with jumbo frames) by adding more processes to handle network traffic—it had been CPU limited with a single process
 - Transfer through VXLAN:
 - Direct: ~12Gb and ~80Gb with jumbo frames (not sure why the difference here is so huge)
 - Between containers: Was still poor, ~10Gb, even with jumbo frames; but found another parameter to also set MTU for pods, and this increased to 65-70Gb
 - Fixed an issue where some nodes were not able to reach ssiweb
 - It seems that the issue was the firewall
 - Helped Lorena to migrate Fermi FTS app to OKD4
- Miscellaneous
 - Issues with getting InCommon certs remain. We have been told we will be given a special temporary key/cert pair that will allow us to request certs (until Jan 20, 2023). That is not working yet; Jeny is on it.
 - Lots of network issues Tuesday on the CMS network related to planned maintenance—the DNS anycast pool did not behave as they expected. Some more issues Thursday morning some other Network maintenance knocked the F5 offline
 - Removed unneeded condor NFS mounts from HPC nodes
 - wc2 mostly operational as secondary Wilson login node
 - More work making “jobsub_lite” work... Added a ‘testing’ yum repository for the developers to use
 - Work on migration consdb to python3 continues
 - Some issues with Nova online (novadaq-far-farm-48)
 - Submitted request to Central Web for star certs to replace expiring Harbor InCommon certs, in case InCommon renewals are not working soon enough.
- New installations/moves/decommissions:
 - New FermiGrid nodes arriving this week.
 - New public dCache nodes are in racks; OS installation/configuration coming soon.
 - New virtualization (GPCF, OKD) hosts are in racks; no networking yet.
 - 14 head nodes for CMS dCache at Receiving; will get them shipped to FCC.

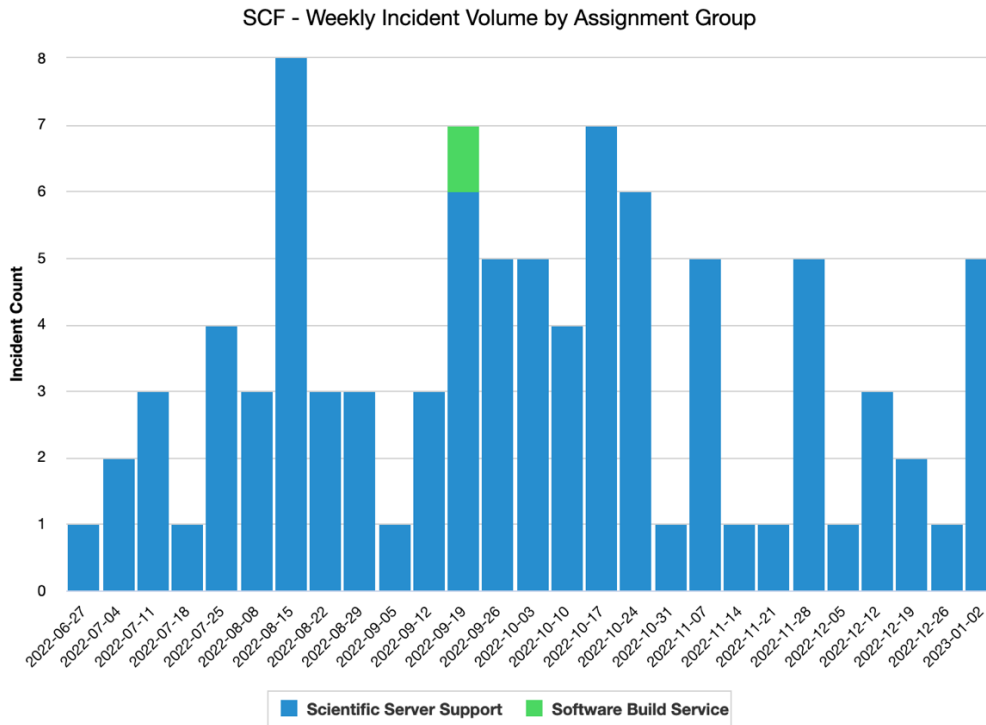
- 7 Ceph servers reinstalled/renamed for SDS department
- Installing AlmaLinux 9 on some FermiCloud hosts

- Normal operations: 161 logged interventions—replacing disks, rebooting hung systems, fixing corrupt RPM databases, retiring systems, fixing xinetd issues, clearing caches, fixing hung NFS mounts, fixing Puppet errors, etc.

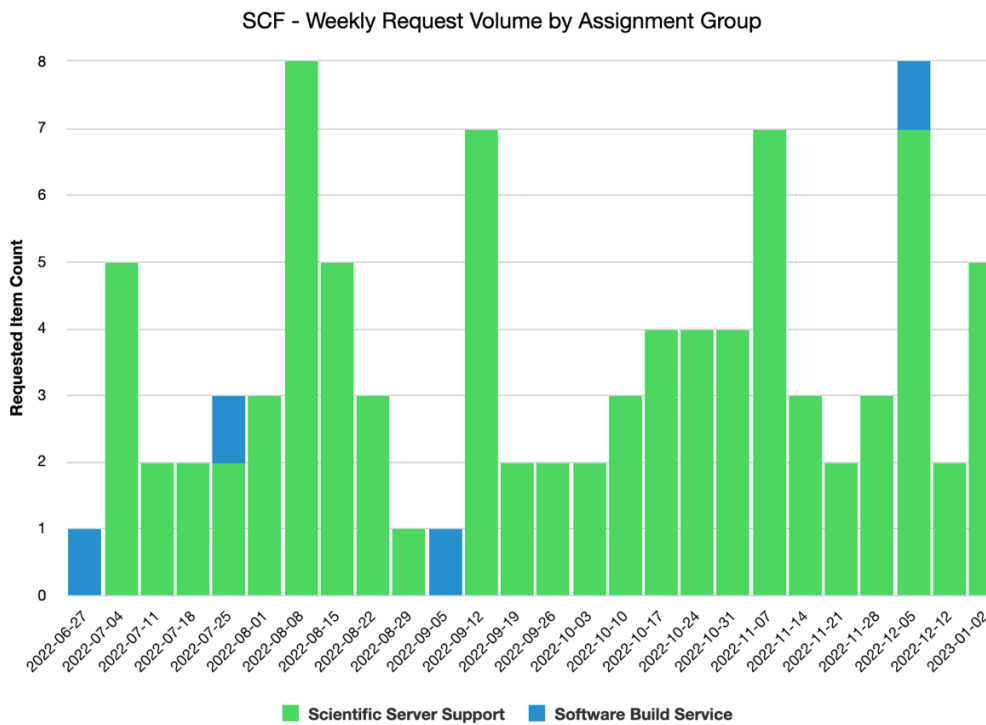
HPC (High Performance & Parallel Computing)

- Bids received for LQ2. Working with Procurement and Compliance to see whether they are acceptable.
- Performance issues on some LQ1 nodes appear to have been solved by rebooting all nodes, on a rolling basis.
- Normal operations

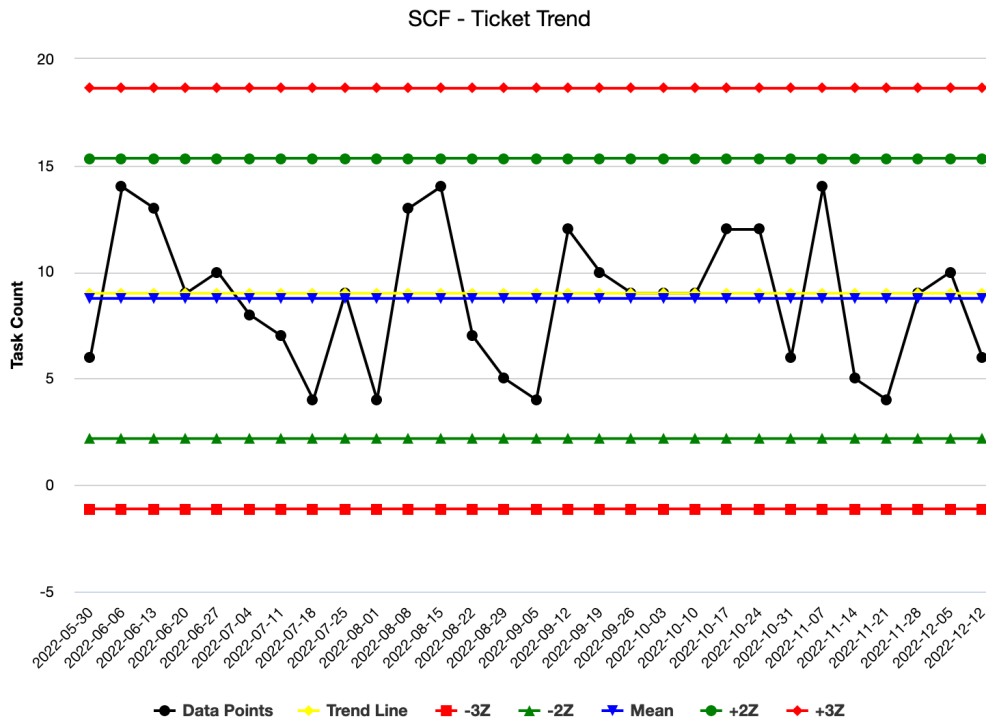
Incidents



Requests



Ticket Trend



GPCF

Totals	
# Hosts	22
# VMs	257
CPU Allocated	50.1%
Storage Allocated	37.3%

Details: <http://ssirhevm1.fnal.gov/report-gpcf/>