

# Experimental Computing Laboratory Meeting January 25, 2024 12:30 PM

Advanced Computing Systems  
Research Section

Jeffrey Vetter, Section Head

Steve Moulton, Systems Engineer

Aaron Young, Software Engineer

ORNL is managed by UT-Battelle LLC  
for the US Department of Energy



U.S. DEPARTMENT OF  
**ENERGY**

# Infrastructure/System Status

- Gradual move to Ubuntu for all research systems and Rocky9 for infrastructure systems, can do RH9 on request.
- Infrastructure VMs moving to less-old hypervisors.
- Infiniband Switch for BlueField2 cards (milan0 & milan1) bricked pending reflash

# System State & Changes (1/2)

- All new heterogeneous-support servers deployed
  - Milan0:
    - 2 \* Nvidia A100 GPUs. Available. MIG available if needed.
    - BlueField 2 card
  - Milan1:
    - Groq AI Accelerators – both deployed on Milan1.
  - Milan2:
    - 8 Nvidia V100 Cards (no nvlink). May prove useful in scaling studies (gift from NCCS). Pending power cabling
  - Milan 3: No accelerators

# System State & Changes (2/2)

- Lewis (workstation)
  - Nvidia GeForce RTX2080 deployed (gift from another group)
  - Existing Nvidia T1000
    - Note both of these are workstation designs, so not in servers.
- Pharoah & Justify moved to Ubuntu 22.04 (from Centos)
  - Nice homogeneous (wrt OS) cluster
  - Reduced software support effort.

# Software installations of note 1/2

- ROCM updated on all AMD GPU systems
  - Multiple versions available via module
  - If you don't pick one, you will get 6.0.0, or nothing, depending
  - aomp installed on all AMD GPU systems
  - system list:
    - explorer (2 \* Instinct MI50)
    - cousteau (2 \* MI100)
    - radeon (Radeon VII)
    - zenith (Radeon RX 6800)
- Module avail to see which ROCMs are deployed
- Multiple ROCM SDK version maintenance now simple (but still not documented by AMD 😊)

# Software installations of note 2/2

- Vanilla OpenMPI 5.0.1 on all Ubuntu 22.04 systems
- OneAPI (base and hpc) deployed on equinox and milan0. Can be deployed elsewhere on request.
- nvhpc updated on all nodes.

# Accounts and Other Management Activity

- All requested accounts in place, except those involving PAS requests. Those should complete next week
  - *Ticket queue will then be empty! Yeah!*
- Siemens project in place, awaiting account requests
  - Siemens licenses received today; will deploy to license server

# Certificate Furball

- As originally deployed,
  - Excl.ornl.gov pointed to AWS machine
  - \*.excl.ornl.gov pointed to on-prem machines.
- Unfortunately this causes problems with new certificate vendor (Sectigo).
- Actions
  - Excl.ornl.gov to be renamed [www.excl.ornl.gov](http://www.excl.ornl.gov)
  - CNAME record (excl.ornl.gov → [www.excl.ornl.gov](http://www.excl.ornl.gov))
  - Certificates updates have been tested and now work.
- User Impact
  - At this point, none. The intrusive parts are done.



# FY 2023 purchases

- Unusual approval delays (not CSMD), some product orders not completed
- In house, or in transit
  - RFSOC 4x2.  
*The RFSOC 4x2 board is a complete, ready-to-use system built around AMD's ZYNQ Ultrascale+ RFSOC ZU48DR device. Featuring four 5 GSPS ADCs with 6 GHz RF input ...*
  - Zenith V2  
*Heterogeneous development system based on consumer-grade GPUs, FPGAs*
  - DVXplorer  
*Dynamic Vision Sensor for neuromorphic projects*
  - Connectx-6 Network adapters and cabling for SDR (2 \* 100Gb interfaces)
  - Multimeter (ooooooh! aaaaah!)  
Useful lab tool

# TRC (Translational Research Center) is coming

- But so is Christmas (but maybe not this year)
- Appears that construction activity safety issues continue
- No additional information since last meeting
- Will house design projects
  - Will have at least 10Gb ExCL network access

# Plans for next three months

- Deploy all received systems
- Continued documentation improvements
- Looking into acquiring Hopper GPUs and other wish-list items.
- Move oswald02 & 03 to Ubuntu 22.04
- 5100/227 (lab space) cleanup and redeploy
- Stabilize non-ORNL-managed UID & GID assignments ExCL-wide to simplify service deployment
  - Also has implications for docker groups

# Application monitoring

- Currently host and application monitoring is via metrics revealed by the CheckMK agent and custom modifications. This is used primarily for system reliability and utilization metrics.
  - Agent is somewhat heavy weight, limiting data granularity. In most cases metrics are collected every 60+ seconds.
    - I.e., agent query is instatiated 60 seconds after previous query completes.
- Not well suited (but can be done) for instrumenting applications

# Current monitoring

- Metric displays are via the CheckMK interface, which is tailored with system engineers in mind.
- This interface also has rule-based metric thresholds (**ok**, **warn**, **crit**) for automatic notification of threshold crossings
- Cool to look at, but not much fun for software engineers and researchers (except maybe OS and monitoring researchers).
- CheckMK can export to Graphite, which can then export to Grafana

login.ftpn.ornl.gov						
State	Service	Icons	Summary	Age	Checked	Perf-O-Meter
OK	Check_MK	☰	[agent] Success, execution time 5.9 sec	2023-12-29 07:01:25 - 27 d	2024-01-25 11:36:30 - 42.8 s	5.86 s
WARN	Check_MK Agent	☰	Version: 2.1.0p21, OS: linux, TLS is not activated on monitored host (see details) <b>WARN</b> . Update error: HTTPConnectionPool(host='checkmk.ftpn.ornl.gov', port=443): Max retries exceeded with url: /checkmk/check_mk/deploy_agent.py (Caused by SSLError(SSLError('bad handshake: Error([('SSL routines', 'tls_process_server_certificate', 'certificate verify failed')])))) <b>WARN</b> . Time since last update check: 345 days 0 hours (warn/crit at 2 days 0 hours/never) <b>WARN</b> . Last update: Feb 14 2023 11:14:33, Agent plugins: 8, Local checks: 3	2023-01-23 18:28:54 - 367 d	2024-01-25 10:41:05 - 56 m	
OK	Check_MK Discovery	☰	no unmonitored services found, no vanished services found, no new host labels	2023-08-15 11:31:44 - 163 d	2024-01-25 09:55:50 - 101 m	
OK	Check_MK HW/SW Inventory	☰	Found 8036 inventory entries	2020-03-03 15:14:29 - 3.9 y	2024-01-24 16:15:14 - 19 h	
OK	certManagerCertificate	☰	certificate /etc/pki/lls/certs/login.excl.ornl.gov.cer notAfter=Mar 2 23:59:59 2024 GMT	2023-02-14 15:11:24 - 345 d	2024-01-25 11:36:30 - 42.8 s	
OK	CPU load	☰	15 min load: 0.22, 15 min load per core: 0.06 (4 cores)	2023-08-04 09:59:11 - 174 d	2024-01-25 11:36:30 - 42.8 s	0.31
OK	CPU utilization	☰	Total CPU: 4.70%	2018-08-14 16:26:48 - 5.5 y	2024-01-25 11:36:30 - 42.8 s	4.69%
OK	Disk IO SUMMARY	☰	Read: 0.00 B/s, Write: 599 kB/s, Latency: 1 millisecond	2018-08-14 16:26:48 - 5.5 y	2024-01-25 11:36:30 - 42.8 s	0.00 B/s / 584.89 kB/s
OK	Filesystem /	☰	73.88% used (12.55 of 16.99 GB), trend: +1.76 MB / 24 hours	2023-04-07 05:02:12 - 293 d	2024-01-25 11:36:30 - 42.8 s	73.88%
OK	Filesystem /boot	☰	18.39% used (186.49 of 1014.00 MB), trend: 0.00 B / 24 hours	2019-04-03 16:43:03 - 4.8 y	2024-01-25 11:36:30 - 42.8 s	18.39%
OK	Filesystem /tmp	☰	12.95% used (2.02 of 15.62 GB), trend: +9.44 MB / 24 hours	2023-02-10 09:32:26 - 349 d	2024-01-25 11:36:30 - 42.8 s	12.95%
OK	HTTP login.excl.ornl.gov	☰	OK - Certificate 'login.excl.ornl.gov' will expire on Sat 02 Mar 2024 11:59:59 PM GMT +0000.	2023-12-29 07:01:27 - 27 d	2024-01-25 11:36:35 - 37.8 s	
OK	Interface eth0	☰	[2], (up), MAC: 52:54:00:6B:08:83, Speed: unknown, In: 5.63 kB/s, Out: 7.73 kB/s	2019-10-08 12:51:54 - 4.3 y	2024-01-25 11:36:30 - 42.9 s	45.0 kbit/s / 61.9 kbit/s
OK	Interface eth1	☰	[3], (up), MAC: 52:54:00:1C:6F:E9, Speed: unknown, In: 3.53 kB/s, Out: 4.22 kB/s	2019-10-08 12:51:54 - 4.3 y	2024-01-25 11:36:30 - 42.9 s	28.2 kbit/s / 33.8 kbit/s
OK	Kernel Performance	☰	Process Creations: 12.17/s, Context Switches: 1492.16/s, Major Page Faults: 0.00/s, Page Swap in: 0.00/s, Page Swap Out: 0.00/s	2021-08-02 15:43:24 - 906 d	2024-01-25 11:36:30 - 42.9 s	0/s
OK	Log /var/log/messages	☰	No error messages	2018-08-14 16:26:48 - 5.5 y	2024-01-25 11:36:30 - 42.9 s	
OK	Logins	☰	On system: 10	2023-08-18 16:13:42 - 160 d	2024-01-25 11:36:30 - 42.9 s	
OK	Memory	☰	Total virtual memory: 4.65% - 1.93 GB of 41.45 GB, 10 additional details available	2022-05-03 06:46:03 - 632 d	2024-01-25 11:36:30 - 42.9 s	9.87%
OK	Mount options of /	☰	Mount options exactly as expected	2018-08-14 16:26:48 - 5.5 y	2024-01-25 11:36:30 - 42.9 s	
OK	Mount options of /boot	☰	Mount options exactly as expected	2019-04-03 16:43:03 - 4.8 y	2024-01-25 11:36:30 - 42.9 s	
OK	Mount options of /tmp	☰	Mount options exactly as expected	2019-08-27 10:27:00 - 4.4 y	2024-01-25 11:36:30 - 42.9 s	
OK	NTP Time	☰	Offset: 0.0276 ms, Stratum: 3, Time since last sync: 7 minutes 55 seconds	2024-01-25 02:53:52 - 8 h	2024-01-25 11:35:34 - 98 s	27.6 μs
OK	Number of threads	☰	518, Usage: 0.33%	2022-12-14 13:21:50 - 407 d	2024-01-25 11:36:30 - 42.9 s	518

Open menu

datasource

graphite.ornl.gov


Instance

excl

host


login\_ftpn\_ornl\_gov

### CPU Usage



wait Min: 0.00790% Max: 3.09% Avg: 0.0568%  
 steal Min: 0.00394% Max: 0.0321% Avg: 0.0121%  
 system Min: 1.96% Max: 15.8% Avg: 2.67%

### Load Average & Thread Count




load1 Min: 0 Max: 1.61 Current: 0.100  
 load15 Min: 0.0900 Max: 0.880 Current: 0.190  
 load5 Min: 0.0350 Max: 1.15 Current: 0.190

### Memory In Use / Committed



mem\_used Min: 1.86 GiB Max: 2.43 GiB Last \*: 1.92 GiB  
 0 Min: 0 B Max: 0 B Last \*: 0 B

### Memory Distribution



buffers Min: 78.1 MiB Max: 93.6 MiB Current: 90.8 MiB  
 cached Min: 13.7 GiB Max: 15.8 GiB Current: 15.1 GiB  
 mem\_used Min: 1.86 GiB Max: 2.43 GiB Current: 1.92 GiB

### Memory Committe



mem\_total Last \*: 19.4 G  
 mem\_inx\_committed\_as

### Swap utilization



### Disk Throughput




Disk\_IO\_SUMMARY.disk\_read\_throughput Min: 0.0 B Max: 195 KiB Current: 195 KiB  
 Disk\_IO\_SUMMARY.disk\_write\_throughput Min: 61.1 KiB Max: 391 KiB Current: 391 KiB

### Overall Disk Utilization



### One Minute Load Average



### CPU Utilization (Total)



### Network Usage Bytes



Interface\_eth0.in Min: 944 B Max: 2.04 MiB Current: 3.74 KiB  
 Interface\_eth0.in Min: 800 B Max: 6.90 MiB Current: 1.54 KiB  
 Interface\_eth0.out Min: 1.40 B Max: 337 KiB Current: 2.75 KiB

### Network Usage -- Packets



Incoming Min: 24.6 B Max: 5.30 KiB Current: 43.9 B  
 Outgoing Min: 2.37 B Max: 5.23 KiB Current: 20.6 B

# Prometheus – a different approach

- Polling metric collection, much like CheckMK
- Agents (Exporters) are generally lighter weight and much more focused. Due to lighter weight, more frequent data collection feasible
- Instrumentation can be added to applications to expose desired metrics via http.
  - Client libraries are available in Go, Python3, Java, Rust & Ruby
- Viewing Prometheus data via Grafana is directly supported. Building queries in Grafana is straightforward.
- **Prometheus is focused on infrastructure performance & application performance monitoring.**



# Instrumentation

- Counters: number or size of events
- Gauge: current state snapshot
- Summary: return multiple counters
- Histogram: set of counters in quantiles
- Buckets: Arbitrary set of counters, scaled by some value (linear or exponential)

# Protocol and Operations

- Runs as an embedded HTTP (apparently not HTTPS) service
- Content is named/tagged/value sets
- Easily human readable
- Prometheus server runs standalone, allows querying and viewing of metrics
  - Not deployed yet; should not be problematic. This is modern yet mature software.
  - Configuration via static YAML files

# Prometheus state in ExCL

- Speculative. Not deployed.
  - Adequate hardware available
- Discussion was stimulated by research inquiry.
- Happy to discuss this one-one or separate meeting

# Questions/Discussion?

- Documentation how tos.
- Rfsoc – where/how to deploy?
- Potential ExCL 1-day workshop
  - AM training session and project discussion
  - PM lightning talks
  - Highlight current work and previous successes