Overview

Björn Dick (HLRS)

Installation schedule

2020-02-14	2048 nodes ready for user access
2020-02-25	Final shutdown of Hazel Hen
2020-04-08 till 2020-04-17	(Parts of) Hawk might be unavailable due to benchmarks / etc.
2020-04-20	All nodes ready for user access
~ May 2020	Regular operation

Still subject to change!!!

Production System

- HPE Apollo 9000
- 5632 compute nodes with 128 cores each
  → 720.896 cores

  → hybrid programming models are essential!
- ~27 PFlop/s peak
- 26 PB flash boosted (transparent) Lustre
- Sharing Pre-/Post-nodes with Vulcan
- Access to HPSS
- RHEL 8 (frontends) / CentOS 8 (compute nodes)

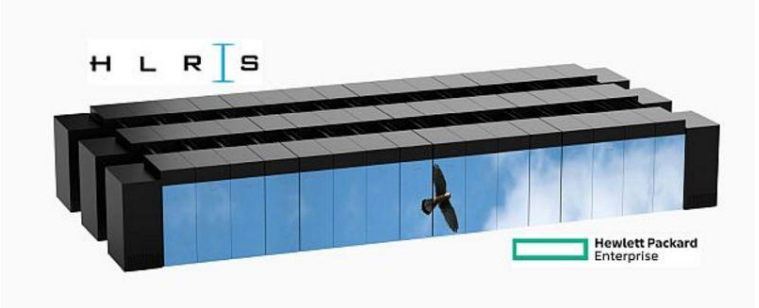

- **`https://kb.hlrs.de/platforms/index.php/HPE_Hawk`**

Test and Development System

- Most of nodes have different micro-architecture (Zen1 aka “Naples”) than production system
  → use those nodes for scaling tests only and treat them as a generic x86_64 black box
- Two nodes have same micro-architecture like production system (Zen2 aka “Rome”)
  → use these nodes for node-level performance optimization

- InfiniBand EDR, one switch

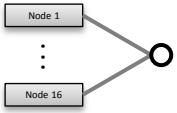

- Software environment is almost the same as will be on production system
  - Currently RHEL/Centos 7, will be updated to v8
  - Further modules / versions might be added
- ws9 mounted

Rome-based nodes

rome1

- AMD EPYC 7702, 64 cores @ 2GHz
- 512 GB DRAM @ 2.666 GT/s
- Settings:
  - 4 NUMA nodes per socket
  - SMT enabled
  - Numa balancing disabled
  - Using transparent huge pages

rome2

probably!

- AMD EPYC 7702, 64 cores @ 2GHz
- 256 GB DRAM @ 2.666 GT/s
- Settings:
  - 1 NUMA node per socket
  - SMT enabled
  - Numa balancing disabled
  - Using transparent huge pages

→ do performance measurements on the same node every time

Support staff

- HPE admins: 3 people
- SiVeGCS / HLST: ~10 people
- Cooperation of HLRS and HPE / AMD:
  - 3 people from HPE
  - Maybe 2 additional people from AMD

- *Work with all of them in bi-annual optimization workshops!*
  https://www.hlrs.de/training

Access

- via usual VPN
- up to now, selected accounts only!

- **`hawk-tds-login1.hww.hlrs.de`**

Storage

- DDN ExaScaler-based High Performance Lustre File system solution
- Using HLRS workspace mechanism
- 26 PB usable capacity
- DDN IME solution works like buffer cache
  - NVMe SSD based storage
  - Bandwidth 250 GB/s
  - Capacity ~ 800 TB
- 48 InfiniBand links into the fabric
- **This is also an expensive resource, pay attention!**

IME®

Fast Data
NVM & SSD

Software

- Performance Tools:
  - Score-P
  - Vampir
  - Scalasca
  - CUBE
  - Extrae / Paraver
  - likwid
  - Darshan
  - AMD uProf (similar to VTune)
  - (VTune Amplifier / Advisor / Inspector)

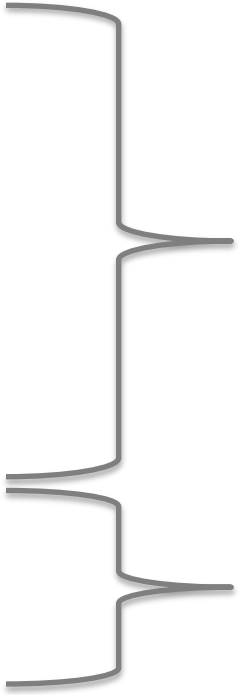

Independent of manufacturer

manufacturer-specific