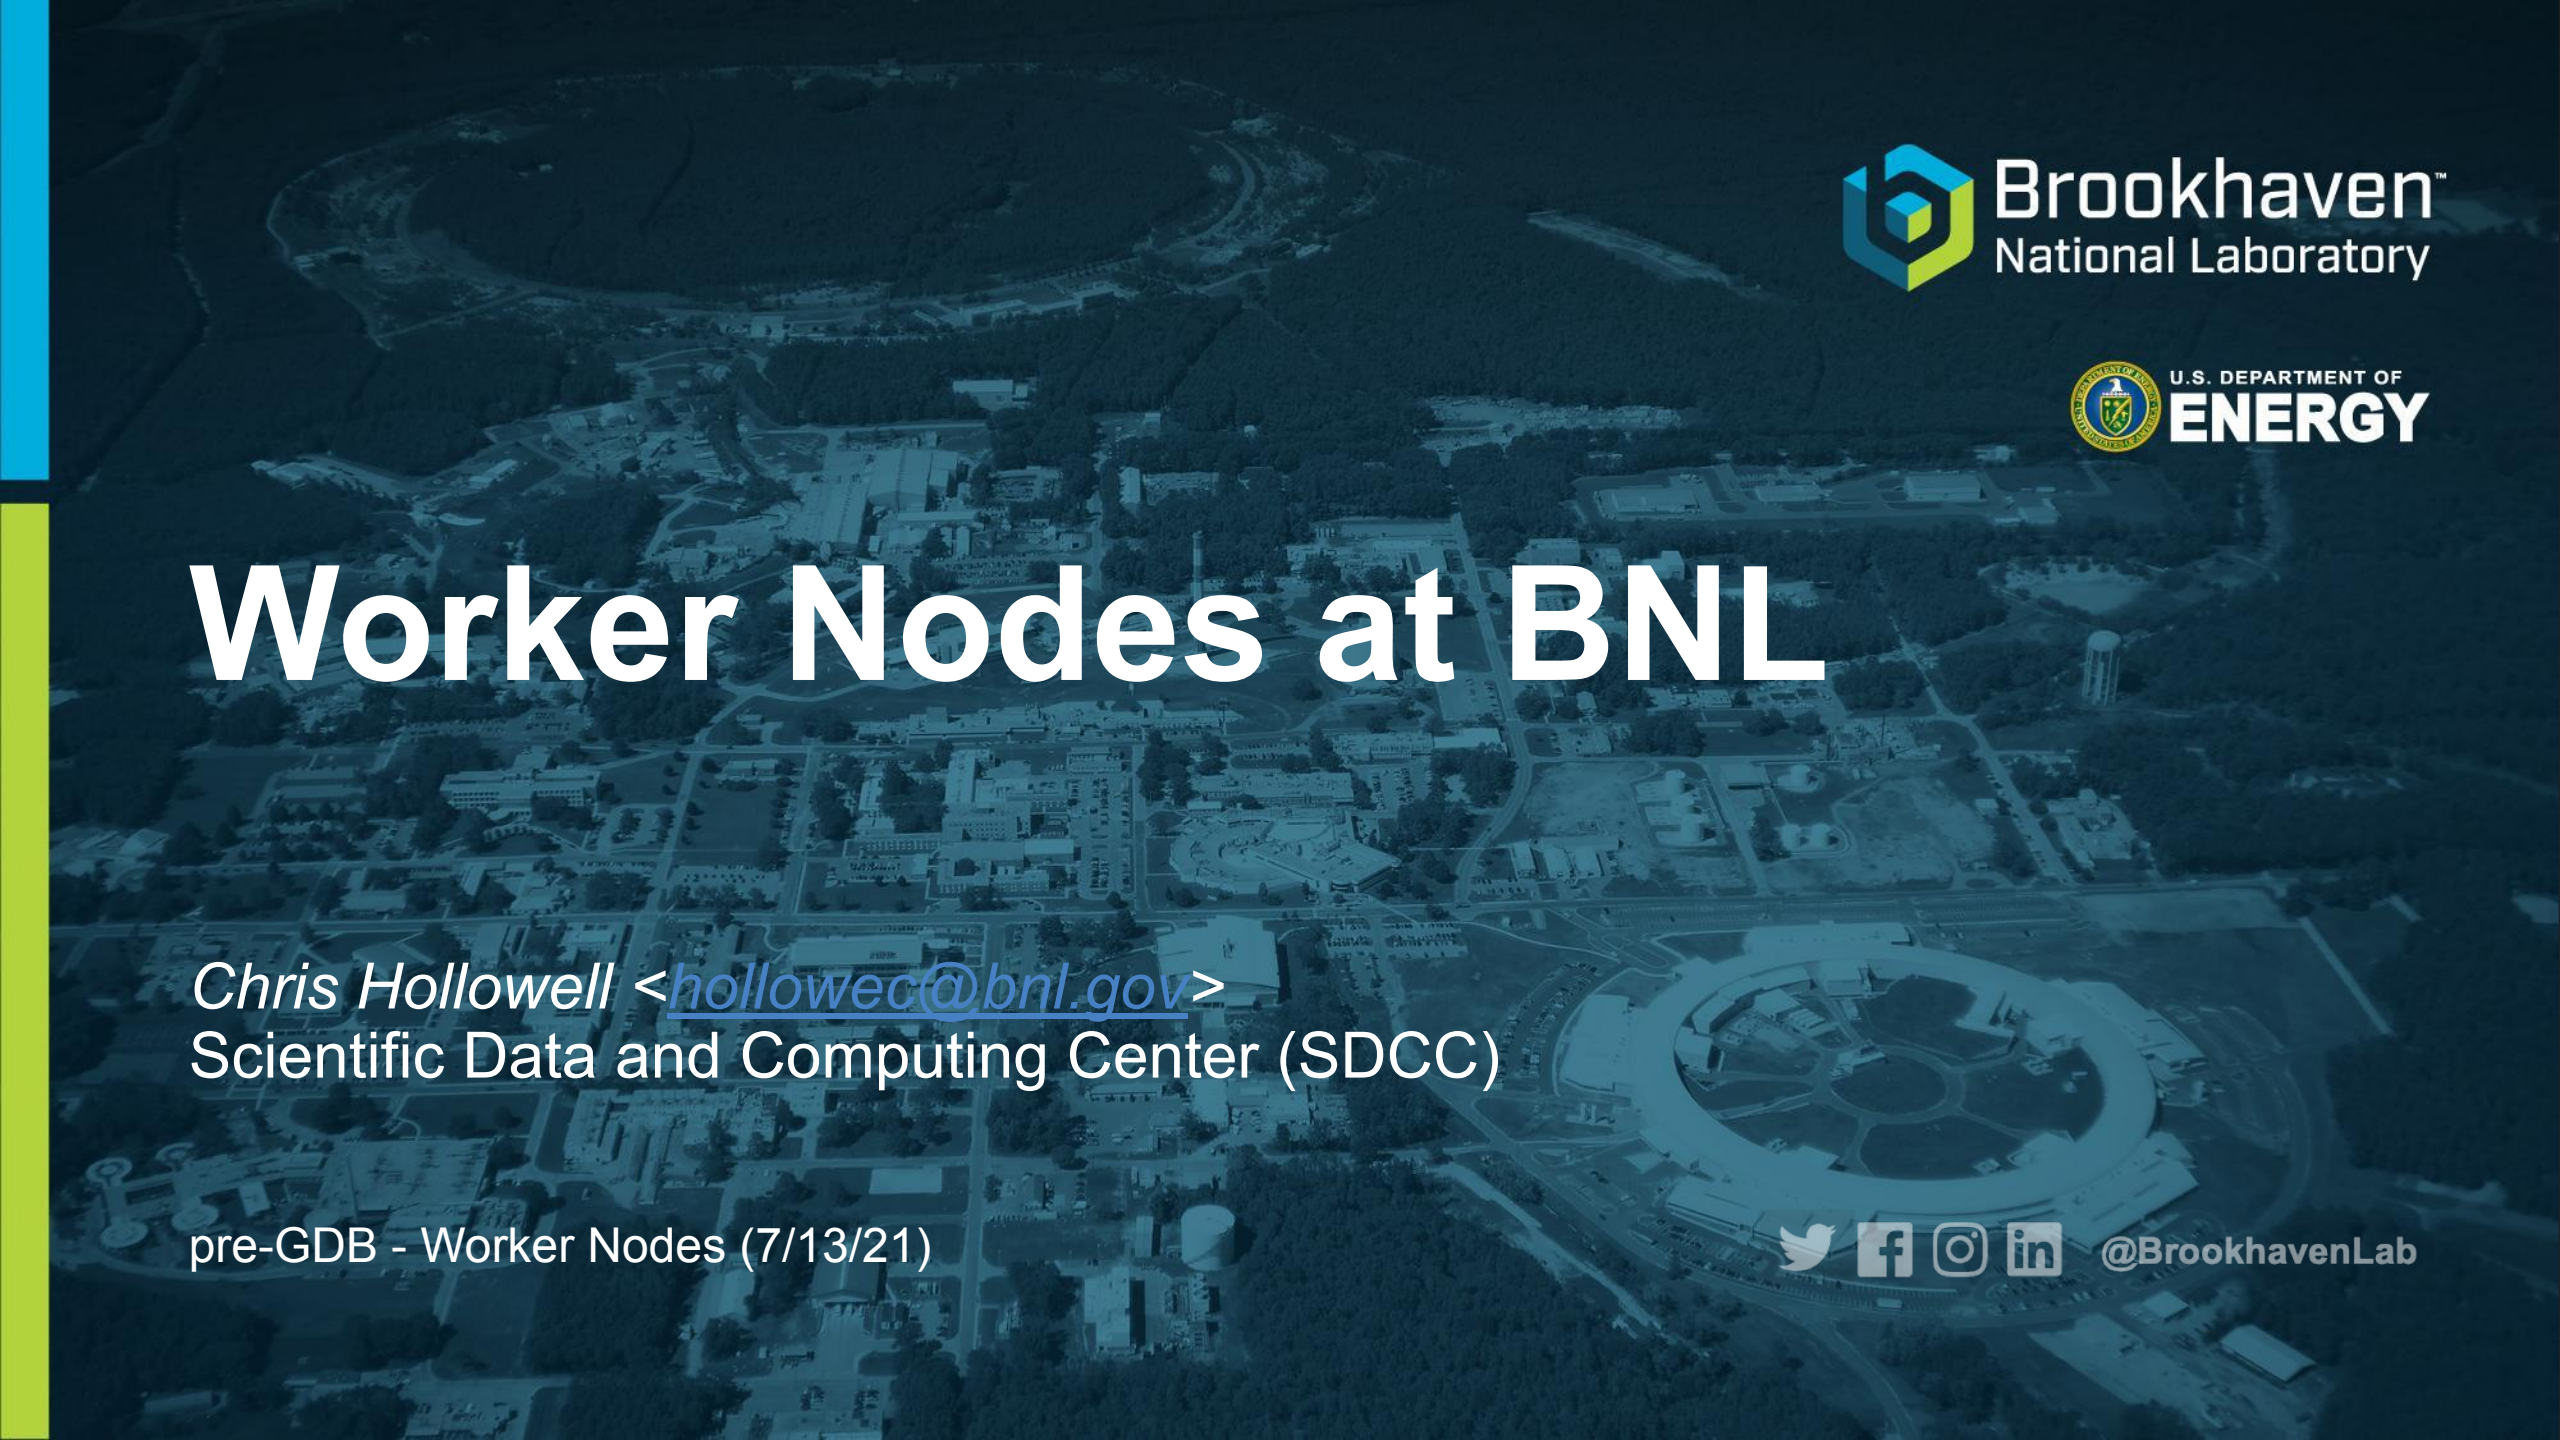

Worker Nodes at BNL

Chris Hollowell <hollowec@bnl.gov>
Scientific Data and Computing Center (SDCC)

pre-GDB - Worker Nodes (7/13/21)

@BrookhavenLab

BNL HTC Overview

- **Support High Throughput Compute (HTC) for a number of HEP/NP experiments**
  - ATLAS (US T1)
  - RHIC T0: STAR, PHENIX, sPHENIX, EIC
  - Belle2

- **~2,000 HTC worker nodes total, ~80k log. cores, ~885 kHS06**
  - A mix of various generations of equipment
    - Primarily Intel-based, but also a number of AMD-based hosts
  - HTCondor utilized as batch system
    - Split into separate “shared” and ATLAS T1 pools
    - Partitionable slots allows for execution of multi-core and high memory jobs
  - All hosts currently running Scientific Linux (SL) 7
  - PXE installations managed by custom PXEUtils software, which utilizes Kickstart
    - Puppet used for node configuration management

BNL HTC Overview (Cont.)

- **Most recent HTC node purchases:**
  - 111 Supermicro SYS-6019TR4 servers (~1140 HS06/node)
    - Dual Intel Xeon Cascade Lake 6252 CPUs @ 2.4 GHz
      - 96 log. cores total
    - 12x 16 GB (192 GB total) DDR4-2933 MHz RAM
    - 4x 1.8 TB SSDs-1U form factor
    - 2x 1 Gbps LACP connectivity
  - Can only fit ~19 nodes in a rack in current datacenter due to thermal/power requirements
    - New datacenter coming online in August has rear door heat exchangers (RDHx), and can increase density to 36 of these nodes per rack
    - 10 Gbps will also be utilized to simplify networking at that density

Rack of Supermicro SYS-6019TR4 Servers

BNL HPC Overview

- Also provide a number of HPC clusters for our user community, available via SLURM batch system
  - **Institutional Cluster (IC)**
    - 216 HP XL190r Gen9 nodes with EDR IB
    - 2x Intel Broadwell Xeon E5-2695v4 CPUs (36 cores total)
    - 256 GB RAM (DDR-2400)
    - 2x K80 or P100 GPUs
  - ATLAS has 6 of these nodes dedicated for their use
    - Available on the grid via ARC-CE
  - Cluster is ~5 years old: in the process of evaluating and purchasing a replacement

  - **ML Cluster**
    - 5 HP Proliant XL270d Gen10 nodes with EDR IB
    - 2x Intel Cascade Lake Xeon Gold 6248 CPUs (40 cores total)
    - 768 GB RAM (DDR4-2933)
    - 8x V100 GPUs

  - **KNL Cluster**
    - 142 KOI S7200AP nodes with dual-rail Omnipath
    - 1x Intel Xeon Phi 7230 CPU (256 log. Cores total)
    - 192 GB RAM (DDR4-1200)

Institutional Cluster

BNL HPC Overview (Cont.)

- **Skylake Cluster**
  - 64 Dell PowerEdge R640 nodes with EDR IB
  - 2x Intel Skylake Xeon Gold 6150 CPUs (36 cores total)
  - 192 GB RAM (DDR4-2666)
- **New HPC cluster for NSLS2**
  - 30 1U Supermicro nodes with EDR IB
  - 2x Intel Cascade Lake Xeon Gold 6252 (48 cores total)
  - 768 GB RAM (DDR4-2933)
  - 12 of the hosts with 2x V100 GPUs

NSLS2 HPC Cluster

BNL Worker Node Purchasing Process

- Typically, worker nodes are spec'd and purchased as follows:
  - **Attempt to consolidate compute purchases for our different experiments as much as possible to take advantage of economies scale and reduce procurement-related effort**
  - **Obtain remote access to a number of vendor labs to test/benchmark new CPU models**
    - HEPSPEC06
    - HEPscore
  - **Based on remote test results, attempt to get some test hardware onsite as well**
    - Vendors often have "try-and-buy" programs
    - Best to get the exact config you want (for all components), physically handle the systems, and run tests with our own node OS image
      - Know exactly what you are getting from a physical management perspective
      - Ensures node OS kernel compatibility/performance

BNL Worker Node Purchasing Process (Cont.)

- **Based on test results, retail pricing, and our own requirements, select one Intel-based and one AMD-based node configuration and ask vendors to quote each**
  - HS06 scores are different, so different quantities for each config are specified to meet pledge
  - Helps select the most cost effective solution between Intel and AMD
- **DOE requires competitive bid with three vendors for large purchases without extensive justification, or purchasing at pre-negotiated rates**
  - Have found that competitive bid tends to yield best cost
  - But introduces some delays in the procurement process
- **After delivery, all systems must pass QA tests via custom RQA software before acceptance**
  - Verifies system configuration (CPU model, memory size, etc.)
  - Performs various CPU/memory/disk stress tests

Vendor Requirements to Minimize SDCC Staff Effort

- **Require full-inside delivery of completely integrated racks**
  - Vendor completes integration offsite
  - All nodes in place, and cabled up appropriately
  - BIOS settings pre-set to our request
    - SMT/HT enabled (yields ~20-30% performance increase), other performance settings
    - Boot mode/order
    - etc.

- **Vendors supply all hostnames, serial numbers and MAC addresses in CSV spreadsheet before delivery**

- **Require HTML5-based full KVM console access via out-of-band management interface on all nodes**
  - Typically just a low-cost add-on license to the management controller
    - Many systems come with this functionality by default now
  - IPMI SOL can be difficult to deal with
    - Doesn't get you access to multiple virtual terminals needed to debug certain issues, particularly with installation failures
    - Need to setup OS getty and BIOS properly in order to utilize
  - Was particularly useful this past year when most staff were working remotely

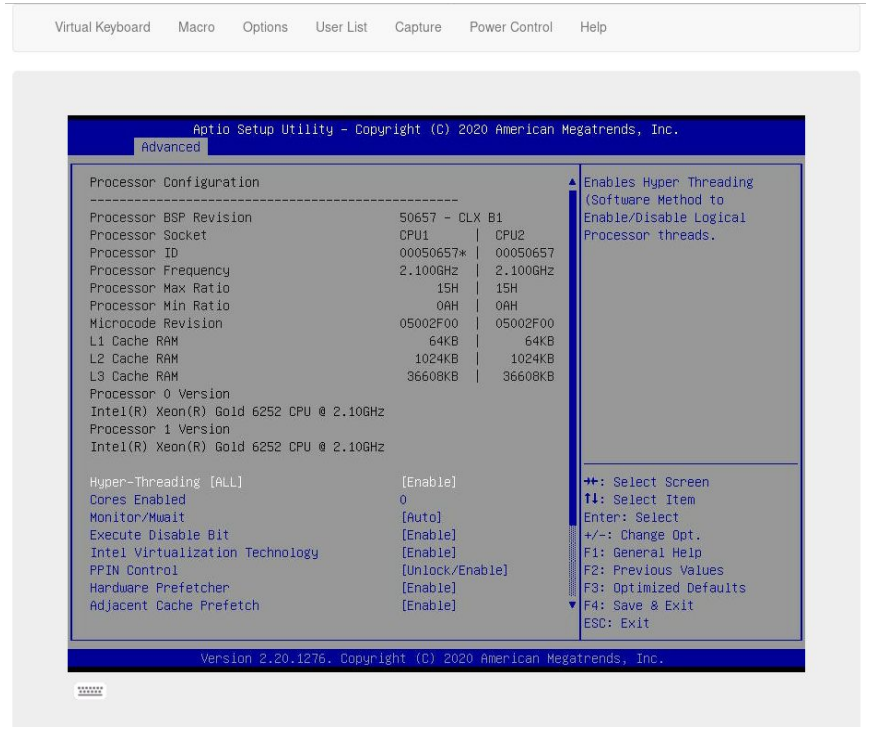

Supermicro iKVM Interface

Considerations When Spec'ing Nodes

- We support multiple experiments, with different node requirements
  - Local I/O characteristics for jobs
  - Required local scratch space per job
  - Job network bandwidth requirements
  - Memory requirements
  - Can be a complicated process to create a configuration that works for everyone

- While it may seem simplest to just purchase as many CPU cores as one can per host to optimize HS06 rating, local disk and network requirements per job are limiting factors
  - More cores = more HTC jobs = more network and disk requirements
  - Retail cost of high-core-density CPUs contribute to the decision to scale-back core count per node as well

- Matching memory size requirements to CPU memory channels can also be complex
  - Want to ensure all available memory channels are populated, and populated equally, to ensure optimal performance

- The density of SSDs has increased substantially over the years, while pricing has also come down considerably
  - All new nodes at BNL purchased with SSDs instead of spinning drives
  - The leading component failure in our nodes has been spinning hard drives
    - Use of SSDs has greatly reduced our hardware maintenance overhead
  - Allows for significantly more local IOPs, so can increase core-density

Potential Issues For Upcoming Purchases

- **Pandemic-related hardware supply issues, but our sources indicate that delays are not expected to be too lengthy (perhaps weeks, not months)**
  - Situation seems to be worse on the storage side
- **CentOS 8 Early EOL**
  - Adoption of CentOS/RHEL 8 for worker nodes has been delayed in the community while a decision on the path forward is made
  - Rocky Linux 8.4 was released in late June
    - BNL is currently testing, but this is likely the solution we will adopt for compute
  - AMD Milan CPUs not officially supported in RHEL/CentOS/SL 7
    - https://access.redhat.com/articles/5899941
    - Sites continuing to run RHEL/CentOS/SL 7 on bare metal wanting an officially supported configuration must continue to purchase AMD Rome, or Intel (Ice Lake or Cascade Lake) CPUs

Conclusions

- BNL HTC compute farm consisting of ~2,000 hosts, and ~80k logical CPU cores
  - Also providing our user community with significant HPC resources in various clusters
- Our purchasing process is typically based on competitive bid, where we allow both Intel and AMD solutions for optimal pricing
- We require that vendors fully integrate our equipment into racks offsite before shipment, that they preset BIOS settings according to our specifications, and that they provide a CSV containing all system serial numbers and MAC addresses
  - Reduces staff manual effort, and time to bring equipment online
- Spec'ing node configurations for multiple tenants can be a complex process
- Pandemic-related delays, and the uncertainty brought on by the CentOS 8 early EOL have complicated the wokernode node purchasing process this year