SAND2019-3153PE

INTEL NUC – (p)HPC SECURE STANDALONE

PRESENTED BY

Victor Kuhns 9328 – Systems Architect

March 6th, 2019

Vision - Bringing HPC resources to Mobile / Desktop

Initial Motivation

- Teach HPC to new generations [STEM]

Personal Desktop HPC

- Functions identically (User/Administrative) to capacity clusters
- Minimal training
- Low power
- Low budget

Full Control and Scheduling of small core workloads on the Desktop

- Solidifies R&D work schedule from inconsistent submission times on Capacity clusters
- Affords relief from Capacity cluster resource contention by moving small core workloads to the Desktop

Secure, Scalable Standalone Computing

- No requirement for corporate network
- Does not require Systems Administrator
- Functional in multiple environments
- Compartmentalized
- Run anywhere

History - We are now in a time to consider desktop HPC

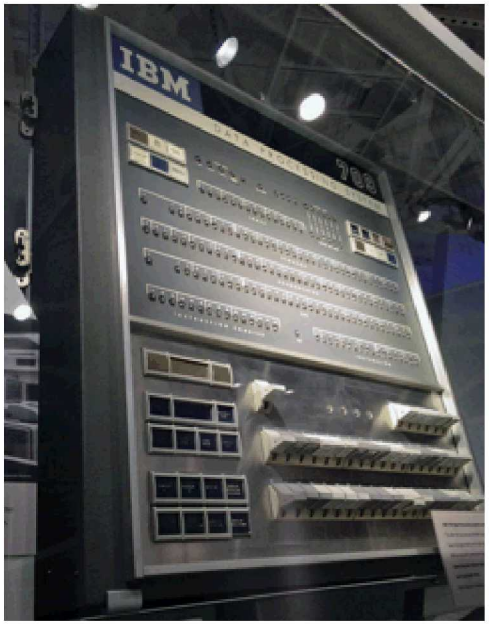

Mainframe

- 1970 – IBM 709
  - 42,000 Addition/Subtraction OPS
  - $2.6 Million
- 100-250 kW for operation and almost as much for cooling

PC – Microprocessor

- Intel 4004
  - 92,600 OPS
  - $5
- Affordable
- No special power and cooling requirements

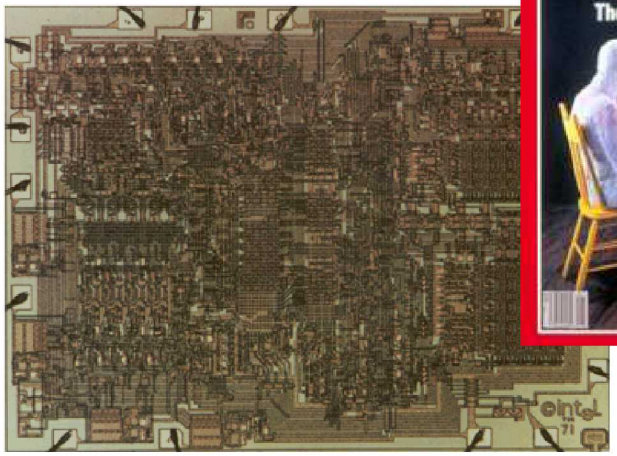

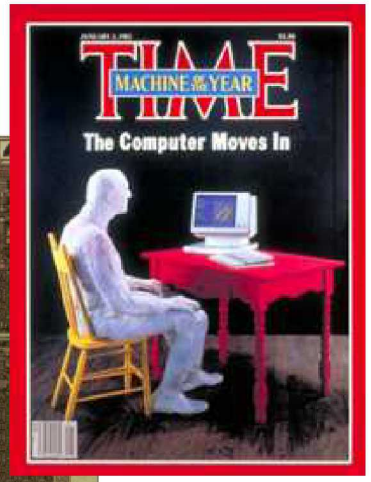

History - ASCI-RED to NUC Desktop in 20 years

ASCI-RED

- $46 Million
- 1.037 TFLOPS
- 1600 SqFT
- 850,000 Watts
- 1816 Compute Blades

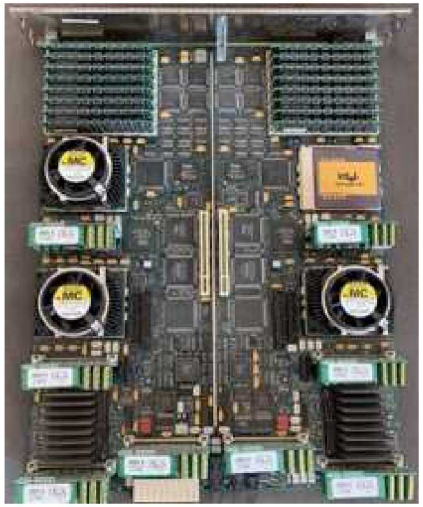

NUC Desktop

- $5100
- 1.2 TFLOPS
- .026 SqFT
- ~75 Watts
- 6 Compute Blades

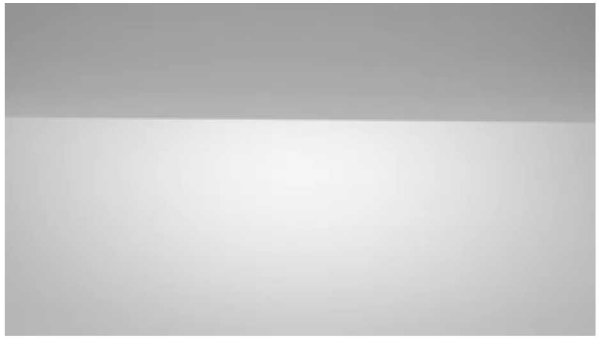

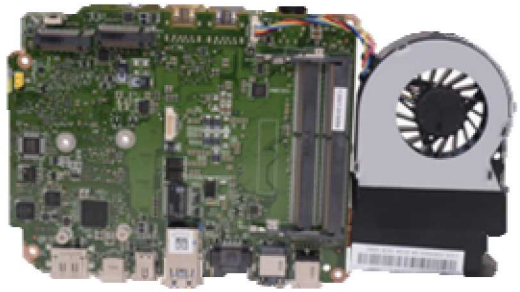

Today - Operational Scalable 1.2 TFLOP Standalone Desktop HPC

Successful Prototype

- Multiple CTH runtimes have completed using input decks straight from Capacity Clusters
- Demonstrated ability to easily boot additional desktop hardware to increase performance
- Primary head node capable of booting production HPC hardware in standalone capability
- Functions in standalone – No external network
- Programming environment identical to current Sandia standards

Current Prototype - Fully functional prototype to demonstrate desktop HPC using SNL production codes

INTEL NUC – HARDWARE

- NUC-RPS – Head node serving TFTP,DHCP,NFS, OS Images
- NUC-LSM[1-6] - Diskless compute (8 Logical Cores, 32GB RAM)
  - 48 Logical Cores
  - 468 Watts
  - ~1.2 TFLOPS. New Hardware has better performance
- Completely Standalone – No network connectivity required
  - Pre and Post processing capability
  - Simply attach monitor and keyboard to head node
- 1 GB or 10 GB(USB-C) Ethernet
- Serviceable Parts (FRU)

SOFTWARE

- RHEL 7.4
  - Munge, dhcp, nfs, iptables, rsync, gcc, etc
- SLURM
  - SLURM 18
  - PAM
- X86_64 Programming Environment
  - GCC, Intel, papi, vtune, make, openmpi, etc

How it works - Operates identical to production HPC Cluster

Hardware

Compute Nodes
Diskless, RHEL 7.4,
PXE BOOT

Compute 1

Compute 2

Compute 3

...

Compute N

Management Node
Diskfull NUC, RHEL 7.4,
PXE, TFTP, NFS, DHCP

Head NUC

Software

SLURM – Queue/Allocate Personal

```
-bash-4.2$ sinfo
PARTITION AVAIL  TIMELIMIT  NODES  STATE NODELIST
ldms-alph    up 2-00:00:00     22  down* ovis-demo-[01-22]
ldms-bet*    up 2-00:00:00      6  alloc nuc-lsm[1-6]
-bash-4.2$ squeue
  JOBID PARTITION     NAME     USER  ST       TIME  NODES NODELIST(REASON)
1000005 ldms-beta vgkuhns_  jpkorbi   R 1-00:49:13      6 nuc-lsm[1-6]
-bash-4.2$
```

OCT1 – CTH Runtime Output

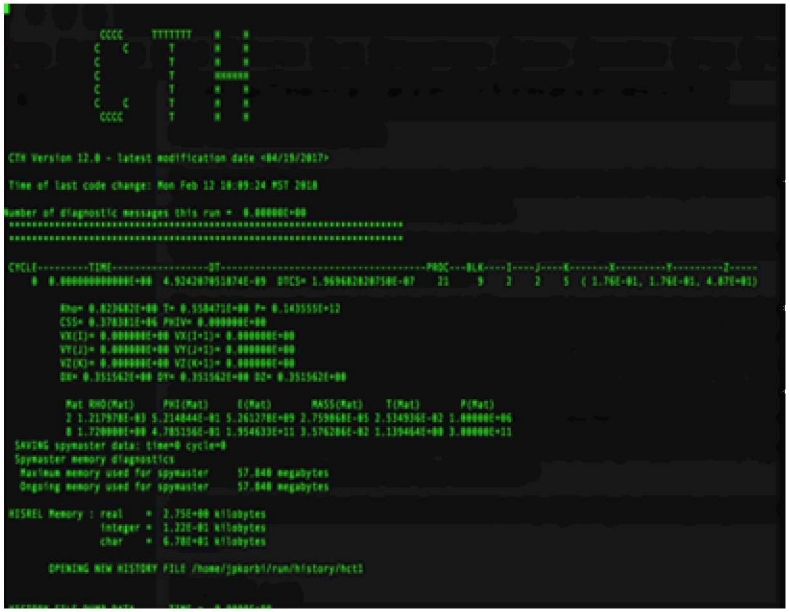

How it works - Scalable from Desktop to Production HPC

Boot either NUC Nodes or Production HPC cabinet nodes

- System has demonstrated scalability from desktop miniPC to 1536 cores on standalone HPC cluster
- All data isolated to management node
- Can seamlessly transfer to larger HPC environment
- Compatible with production submission scripts
- Compartmentalized (Separate management nodes)

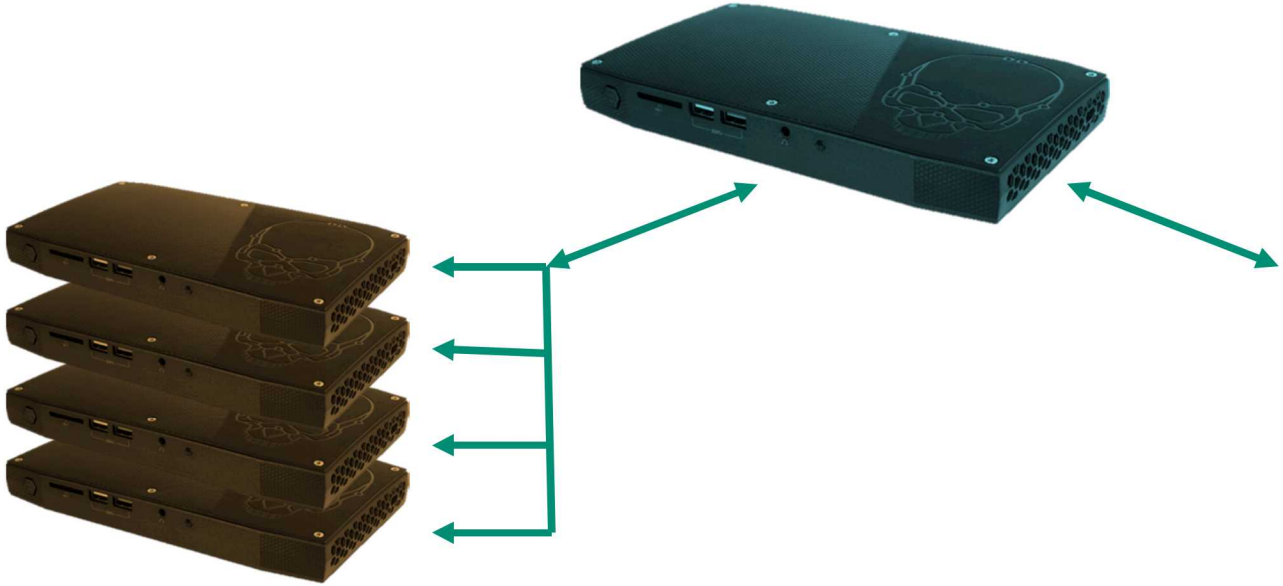

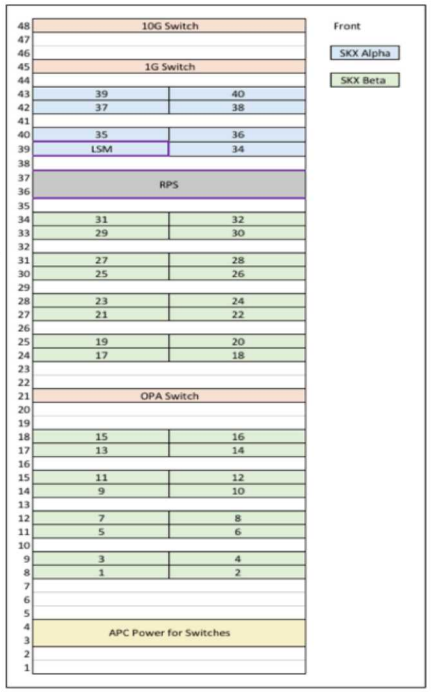

Intel Skylake x86-64,
AVX512 CPU Architecture

How it works - Portable HPC Cluster

Boot NUC Cluster and take HPC with you

- Functional in multiple environments
  - Workstation
  - Remote Facilities (Ie: Thunder Range)
  - Vehicle
  - Airborne
- Allows for HPC at data collection point
  - No need for special cooling or power facilities

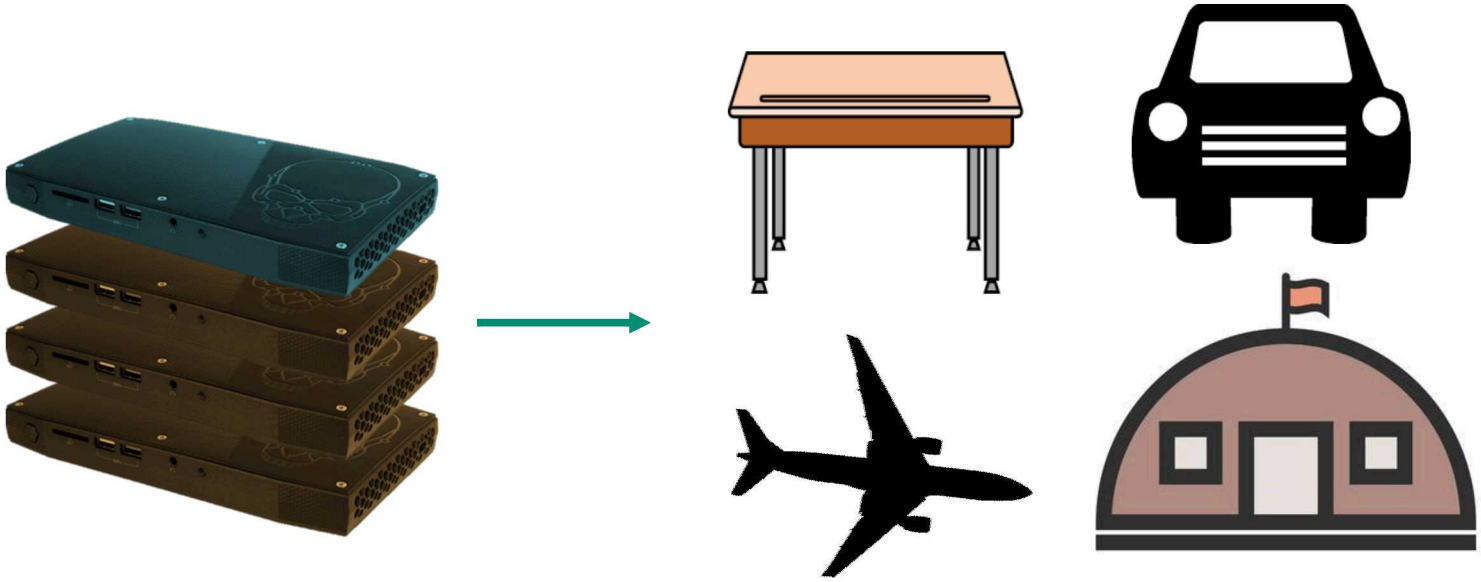

Secure Concept – Compartmentalized Standalone Computing

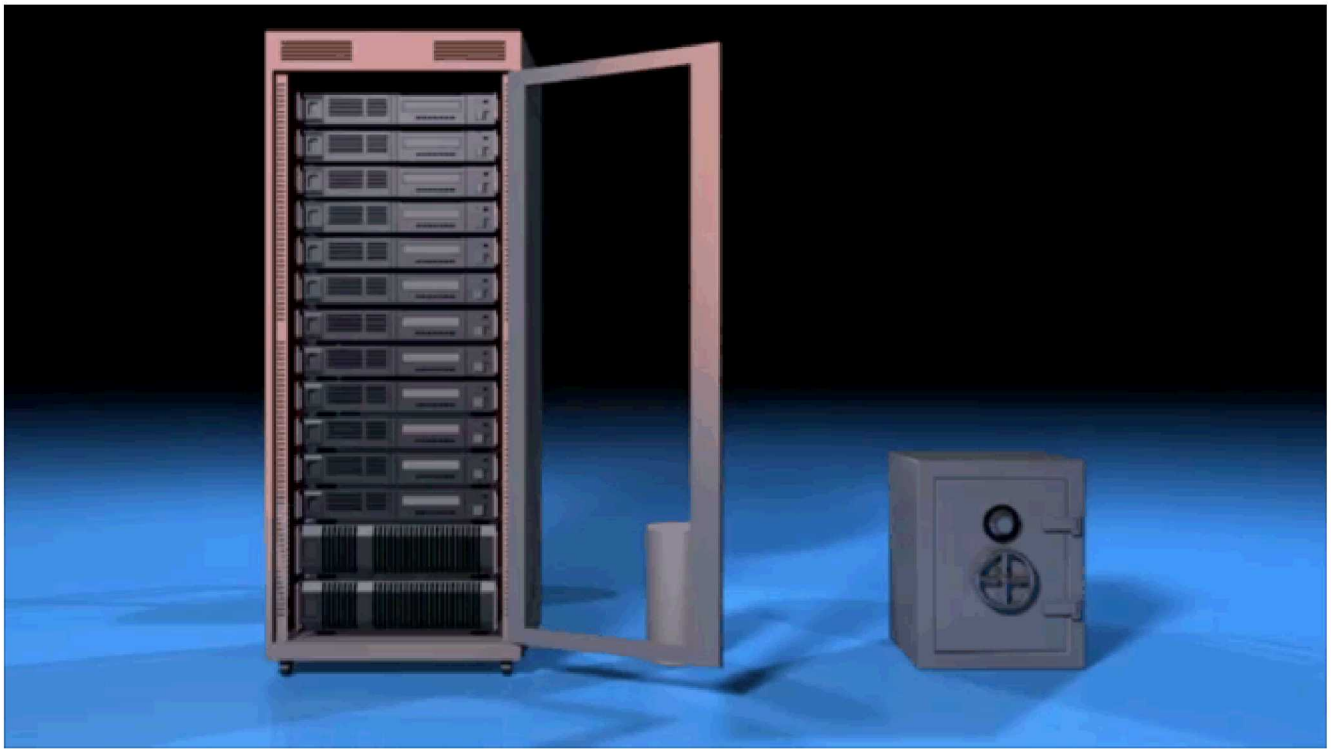

Typical HPC Model

- HPC Cabinet connects and processes data via corporate network
- 010101 – Left Side in beginning

Our Model

- Remove corporate network
- Boot HPC Cabinet in standalone using NUC Head Node
- Secure NUC in safe once work is complete
- Power cycle rack, prepare for next compartment, next NUC.

John Korbin (1555) booting HPC resource from NUC, analyzing runtime data on NUC Console.

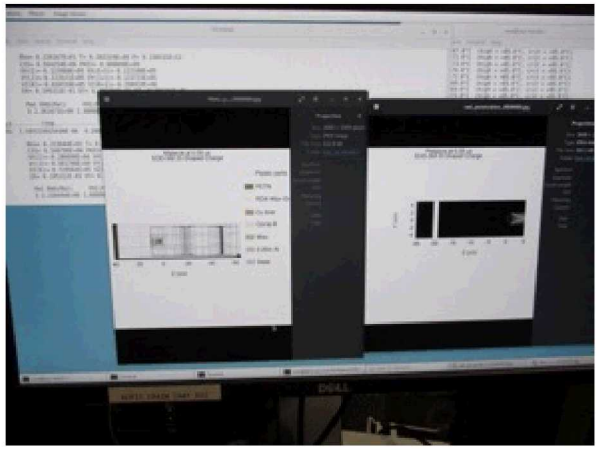

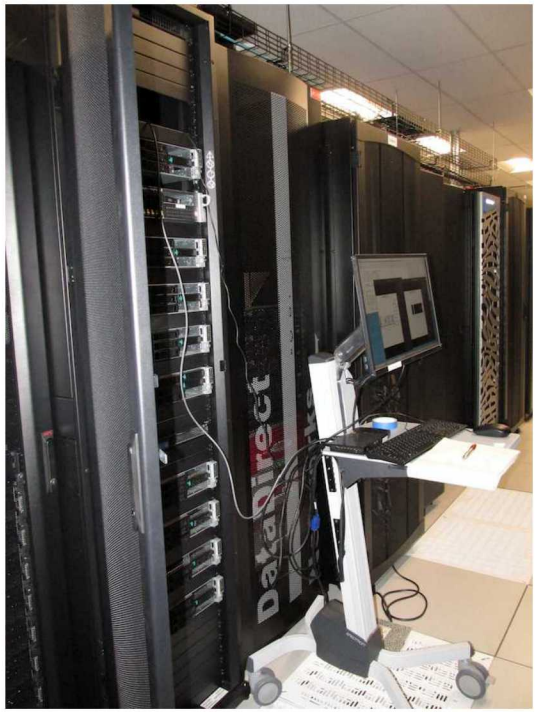

CTH – Shape Charge

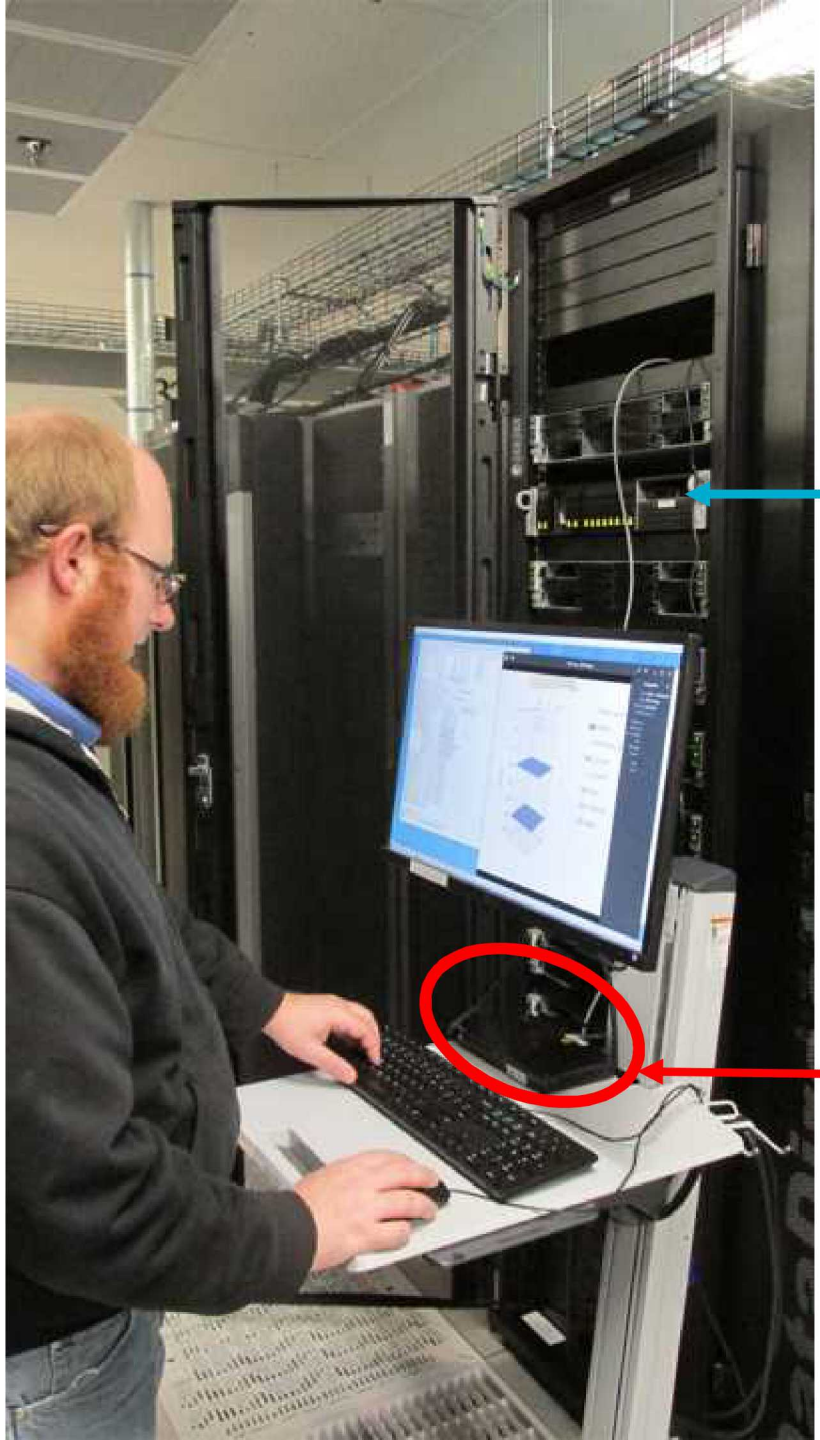

HPC Cabinet – CTH Performance comparison

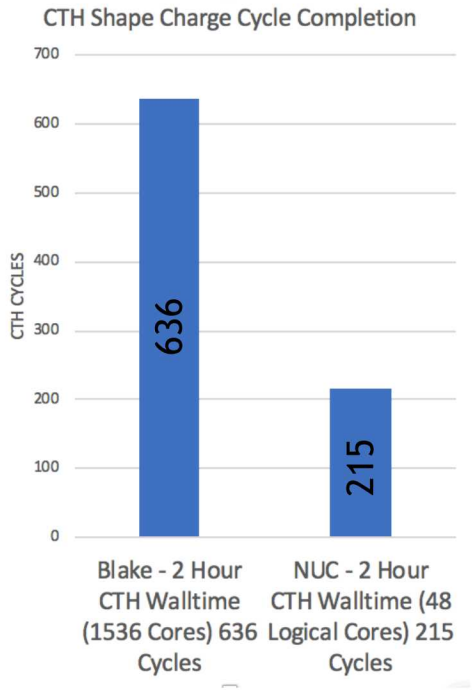

Blake – 1536 Skylake HPC Cabinet

- $2\mu s$ simulation output (jpg)

NUC – 2 Hour

- $0\mu s$ simulation output (jpg)

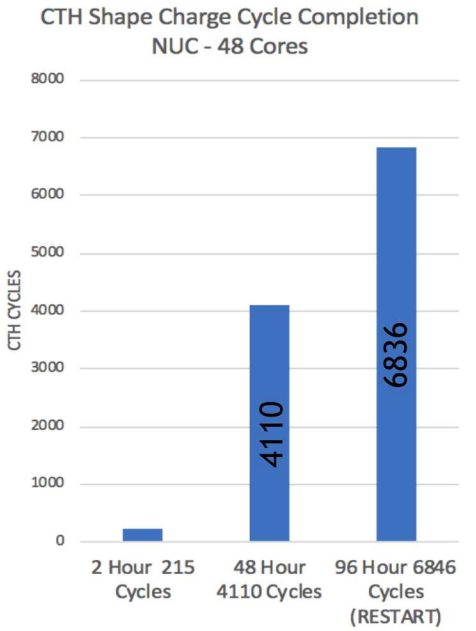

NUC – 48 Hour

- $11\mu s$ simulation output (jpg)

NUC – 96 Hour (Restart)

- $24\mu s$ simulation output (jpg)

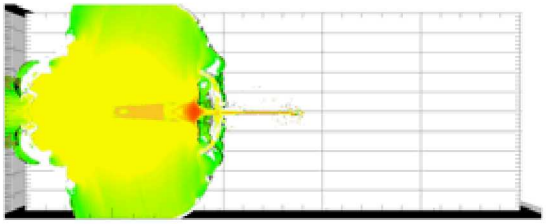

HPC Cabinet – 2 Node Serrano (CTS) Compute vs NUC pHPC

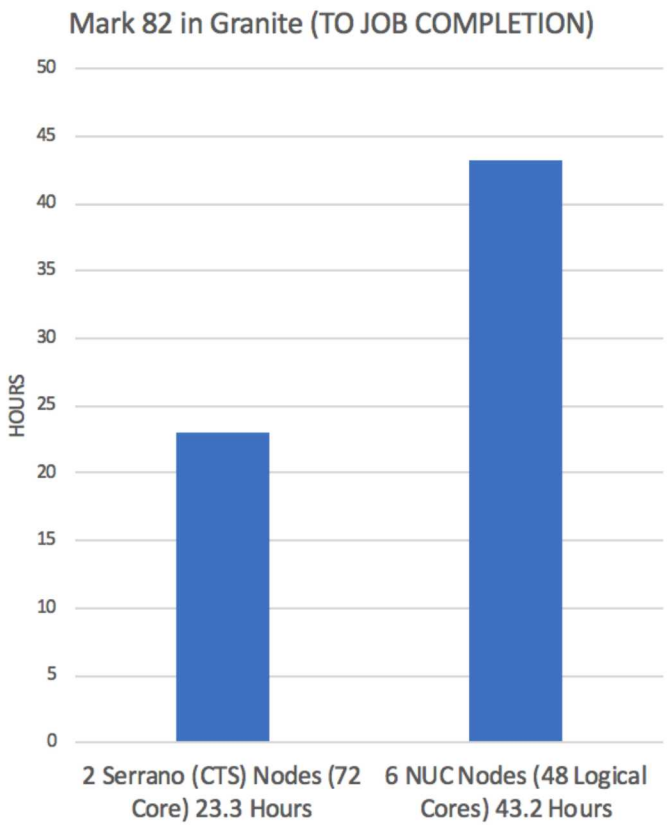

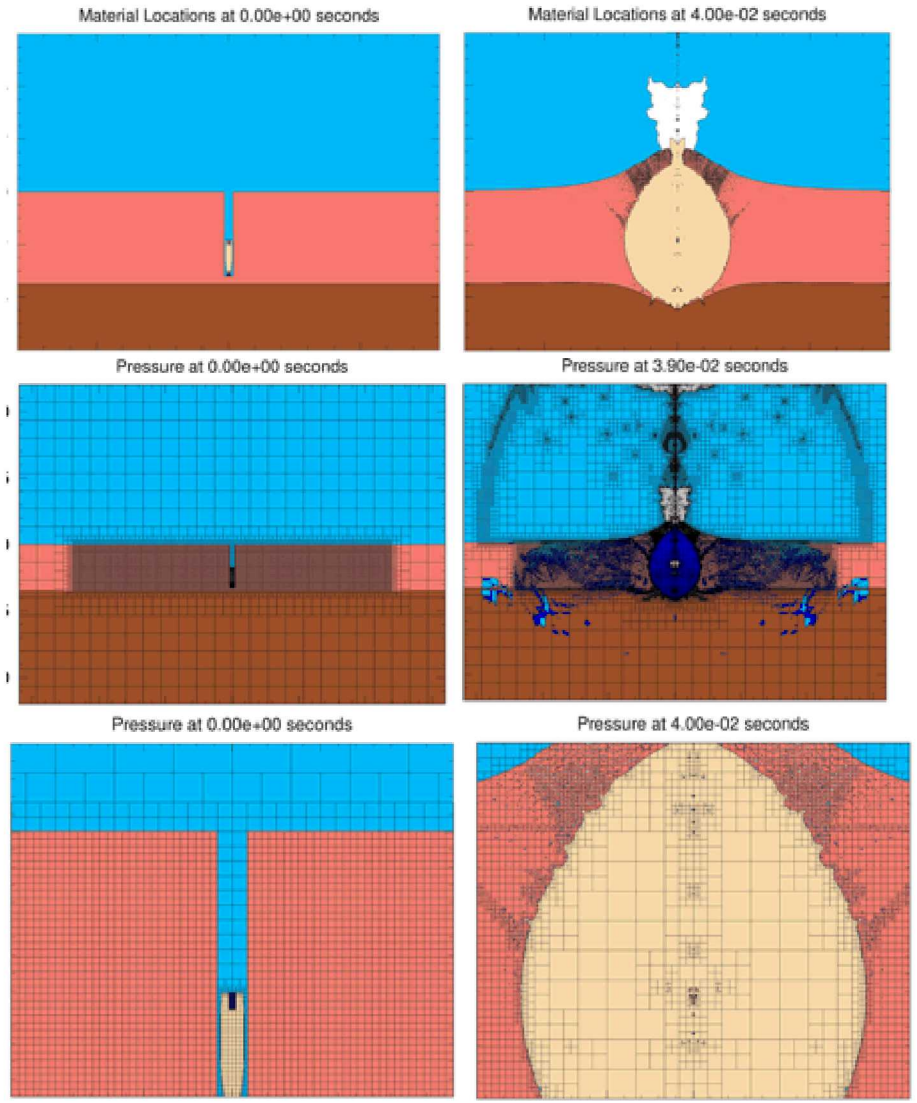

Serrano

- 1.2 TFLOPs/Node
- 2.1 GHz Intel Broadwell
- ~380 Watts/Node

Progress to date

Fully Operational HPC Desktop environment

- Mirrors functionality of production clusters
- Batch queue submission capability
- Compiler Environment

CTH 12.1 Built

- Currently compiled and shown to be operational on the desktop
  - Various OUO problems have been tested
- Capable of cold booting HPC cabinet for high fidelity runs
  - Tested on 1536 cores, 32 Nodes (Intel Skylake / Omnipath)

Desktop Workload

- Capable of running unmodified input decks used on production clusters

CTH input deck supplied by: John Korbin, Jason Wilke, David Peterson

Future - Production capability for desktop and Secure Standalone HPC

Current Capabilities

- Under $10k
- Wall Power (Minimal Facilities)
- Proven standalone system that never touches internet or exposes existence of a project
- Ability for single individual to control all data from start to finish

Path forward

- Review architectures in testbed environment
- Review security requirements in non networked environments
- Continued testing of Sandia Mission Codes
- Engage with manufacturer to drive hardware development

Staffing

TBD

Funding

DOE/NNSA