جامعة الملك عبدالله للعلوم والتقنية
King Abdullah University of Science and Technology

Core Labs and Research Infrastructure

Shaheen 101 HPC Training

Supercomputing Core Laboratory (KSL)

Shaheen 101 HPC Training

Agenda

- Overview of Shaheen and Neser Architecture
- How to get an account and access to Shaheen
- Shaheen Storages: changes and policies storage
- Running environment and Job Scheduling
- Visualizations tools
- Running HPC workflow on Shaheen: VASP example
- Applications software example: CFD example
- Programming environment, debugging and profiling, Best practices and tips
- Questions and open discussions with KSL team

Follow us on Twitter: @KAUST_HPC

جامعة الملك عبدالله للعلوم والتقنية
King Abdullah University of Science and Technology

Core Labs and
Research Infrastructure

Shaheen 101 HPC Training

Introduction to KSL

KSL Mission and Goal

KSL Mission

- To provide state-of-the art supercomputing facilities, training and service to KAUST students, faculty, researchers, and to serve supercomputing needs in the Kingdom

Goal

- Become world-class reference supercomputing center in the Kingdom

KSL Supercomputing Services

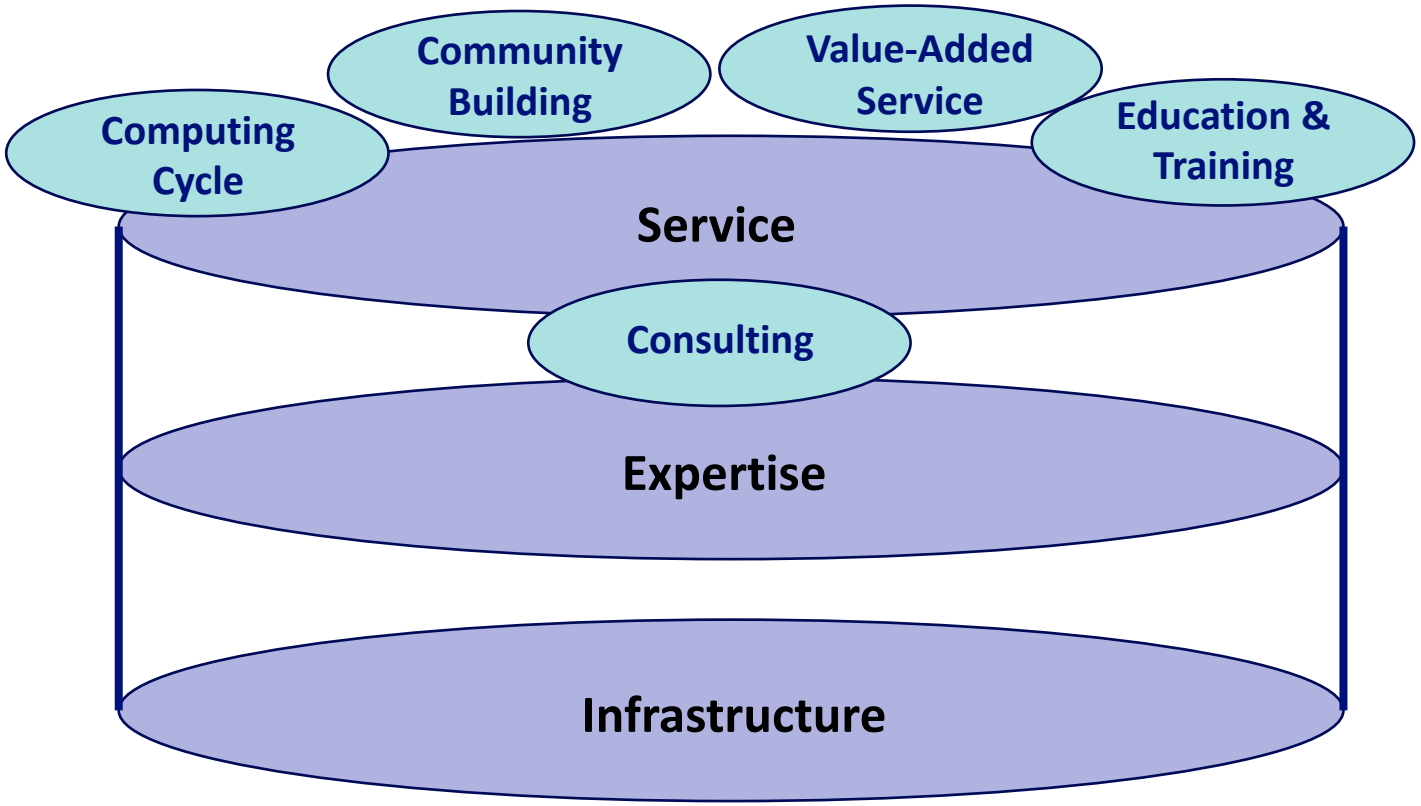

جامعة الملك عبدالله للعلوم والتقنية
King Abdullah University of Science and Technology
Core Labs and Research Infrastructure

Scientists Team

Dr. Saber Feki, Scientists Lead

Dr. Rooh Khurram

Dr. Bilel Hadri

Dr. Mohsin Shaikh

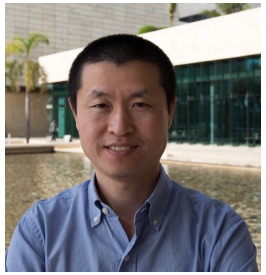

Dr. Zhiyong Zhu

Dr. Nagarajan Kathiresan

Dr. James Kress, KVL

Shaheen 101 Registrations

Position

193 responses

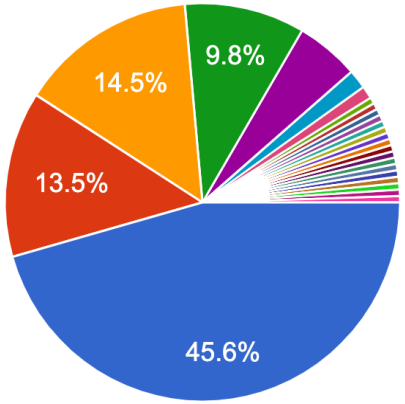

- ~90 in-Kingdom participants
- Most with no Shaheen account **YET**
- Diverse science areas

Do you have an account on Shaheen2?

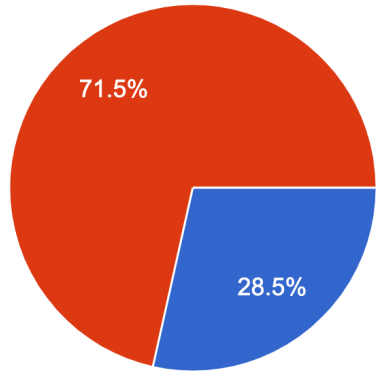

Field of research

193 responses

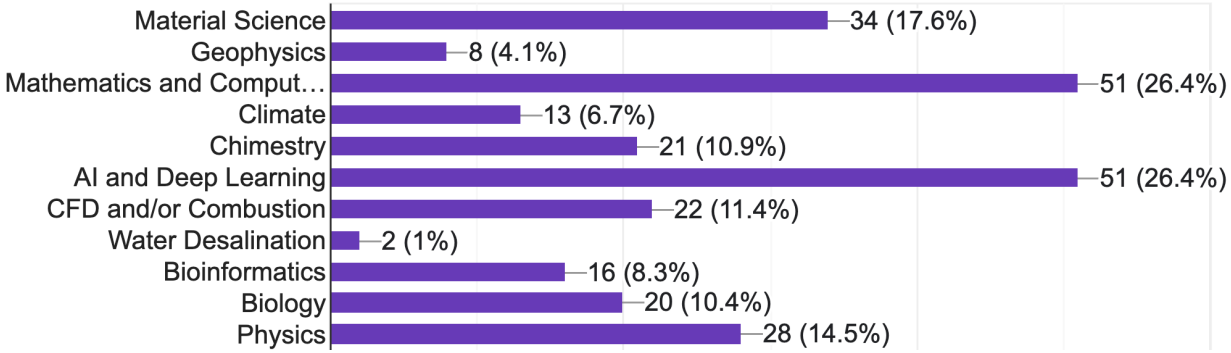

جامعة الملك عبدالله للعلوم والتقنية
King Abdullah University of Science and Technology

Core Labs and Research Infrastructure

Overview of Shaheen and Neser Architectures

Saber Feki, PhD

Senior Computational and Data Scientist Lead

Typical Supercomputer's Architecture

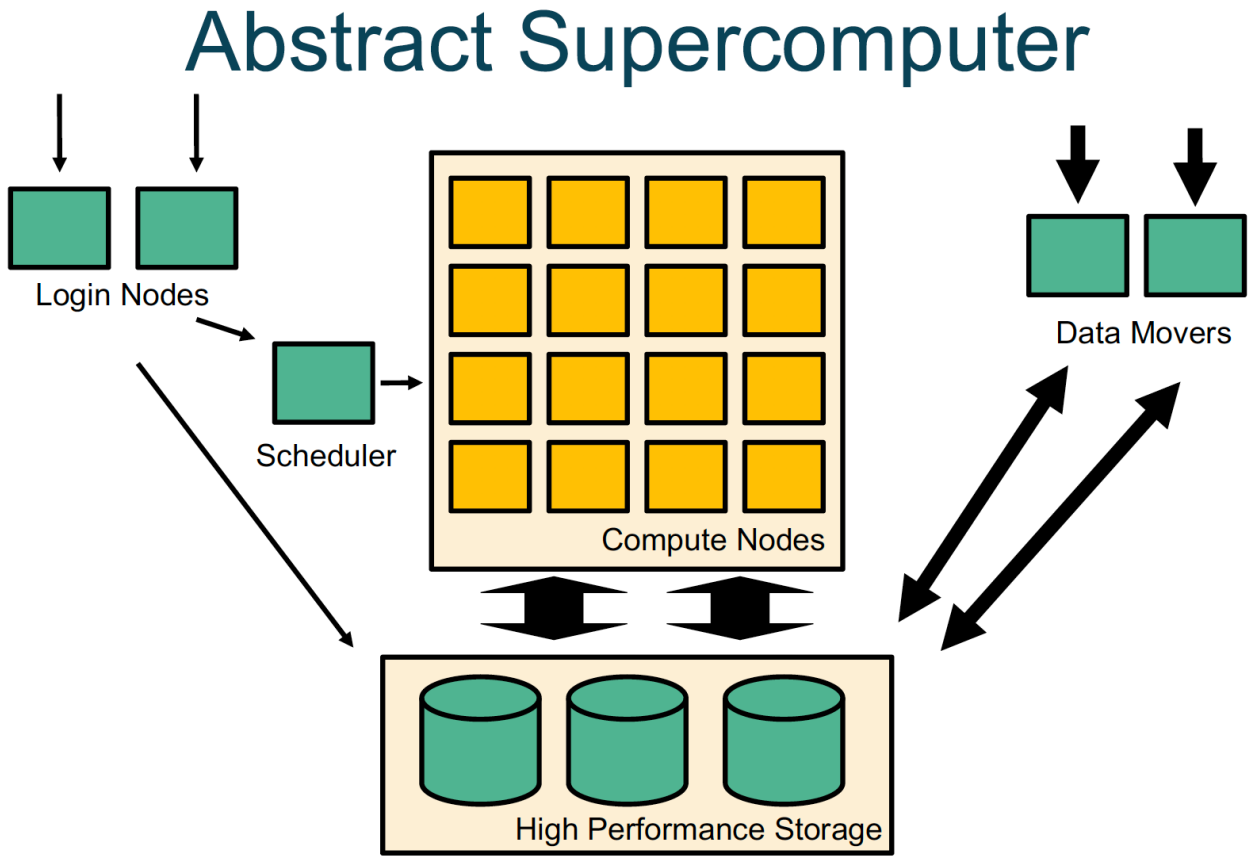

PAWSEY supercomputing centre

Shaheen II overview

COMPUTE	Node	Processor type: Intel Haswell	2 CPU sockets per node, 16 processors cores per CPU, 2.3GHz
		6174 Nodes	197,568 cores
		128 GB of memory per node	Over 790 TB total memory
	Power	Up to 3.1MW	Water Cooled
	Weight/Size	More than 100 metrics tons	36 XC40 Compute cabinets, plus disk, blowers, management , etc..
	Speed	7.2 Pflop/s speak theoretical performance	5.53 Pflop/s sustained LINPACK
	Network	Cray Aries interconnect with Dragonfly topology	57% of the maximum global bandwidth between the 18 groups of two cabinets.
STORE	Scratch	Sonexion 2000 Lustre appliance	17.6 petabytes of usable storage. Over 500 GB/s bandwidth
	Project	E1000 Lustre appliance (2022)	37 petabytes of usable storage.
	Burst Buffer	DataWarp	Solid Sate Devices (SSD) fast data cache. Over 1.5 TB/s bandwidth
	Archiving	Tiered Adaptive Storage (TAS)	Hierarchical storage with NetApp disk cache and ~90 PB of tape storage, using a spectra logic tape library.

Shaheen II XC40 Compute Blade

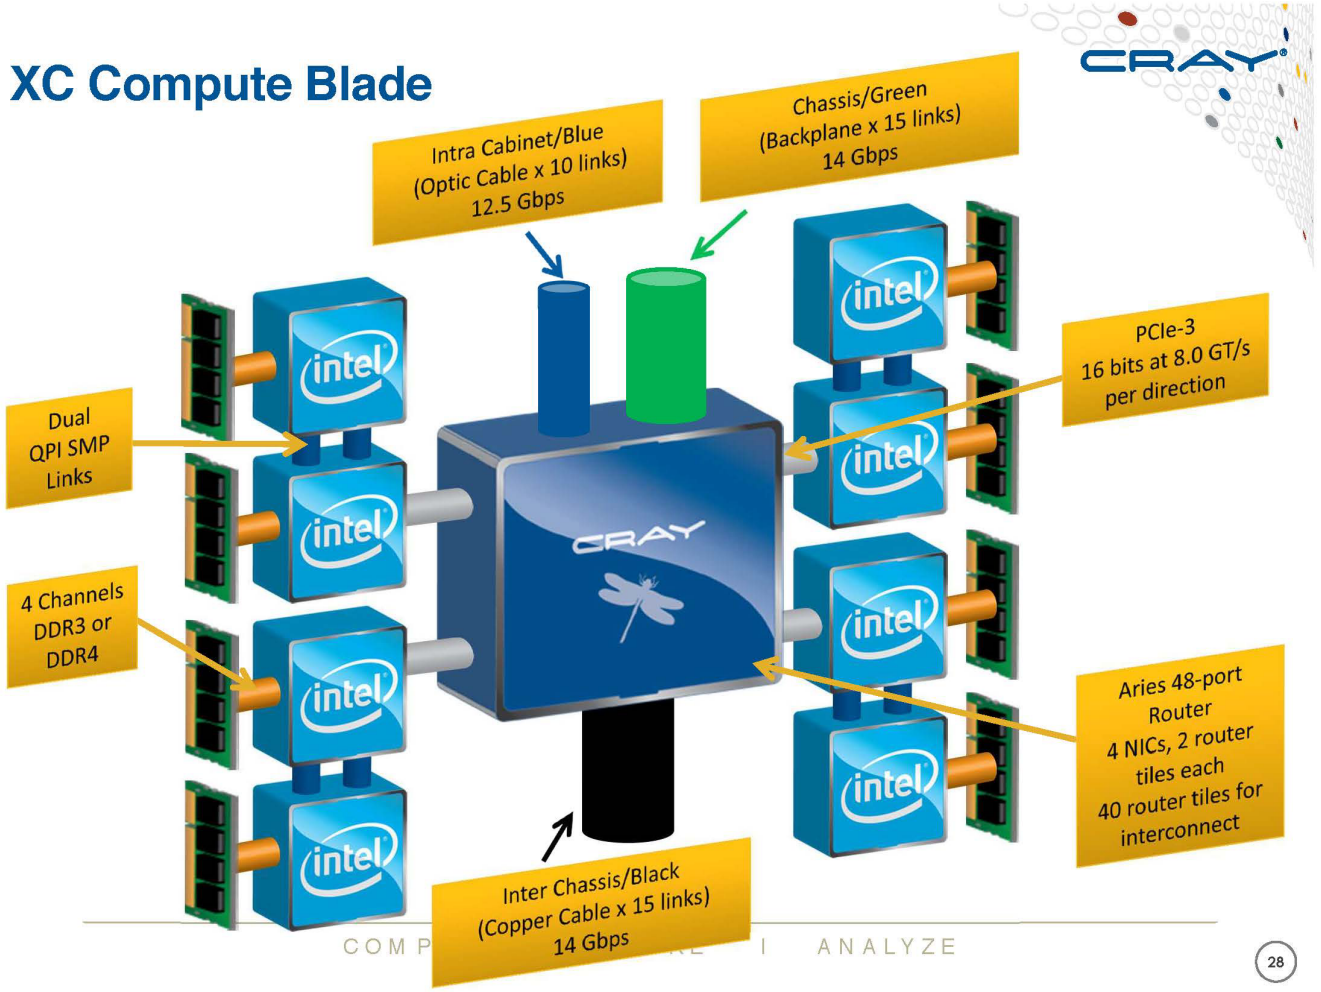

Shaheen II XC40 Interconnect

Cray XC Rank1 Network

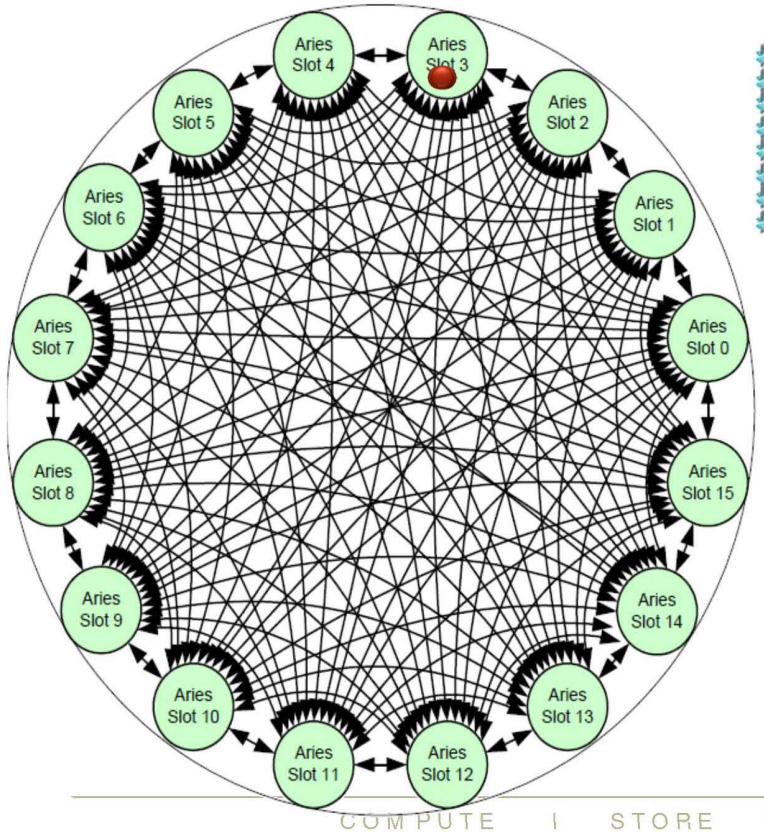

- Chassis with 16 compute blades
- 128 Sockets
- Inter-Aries communication over backplane
- Per-Packet adaptive Routing

Shaheen II XC40 Interconnect

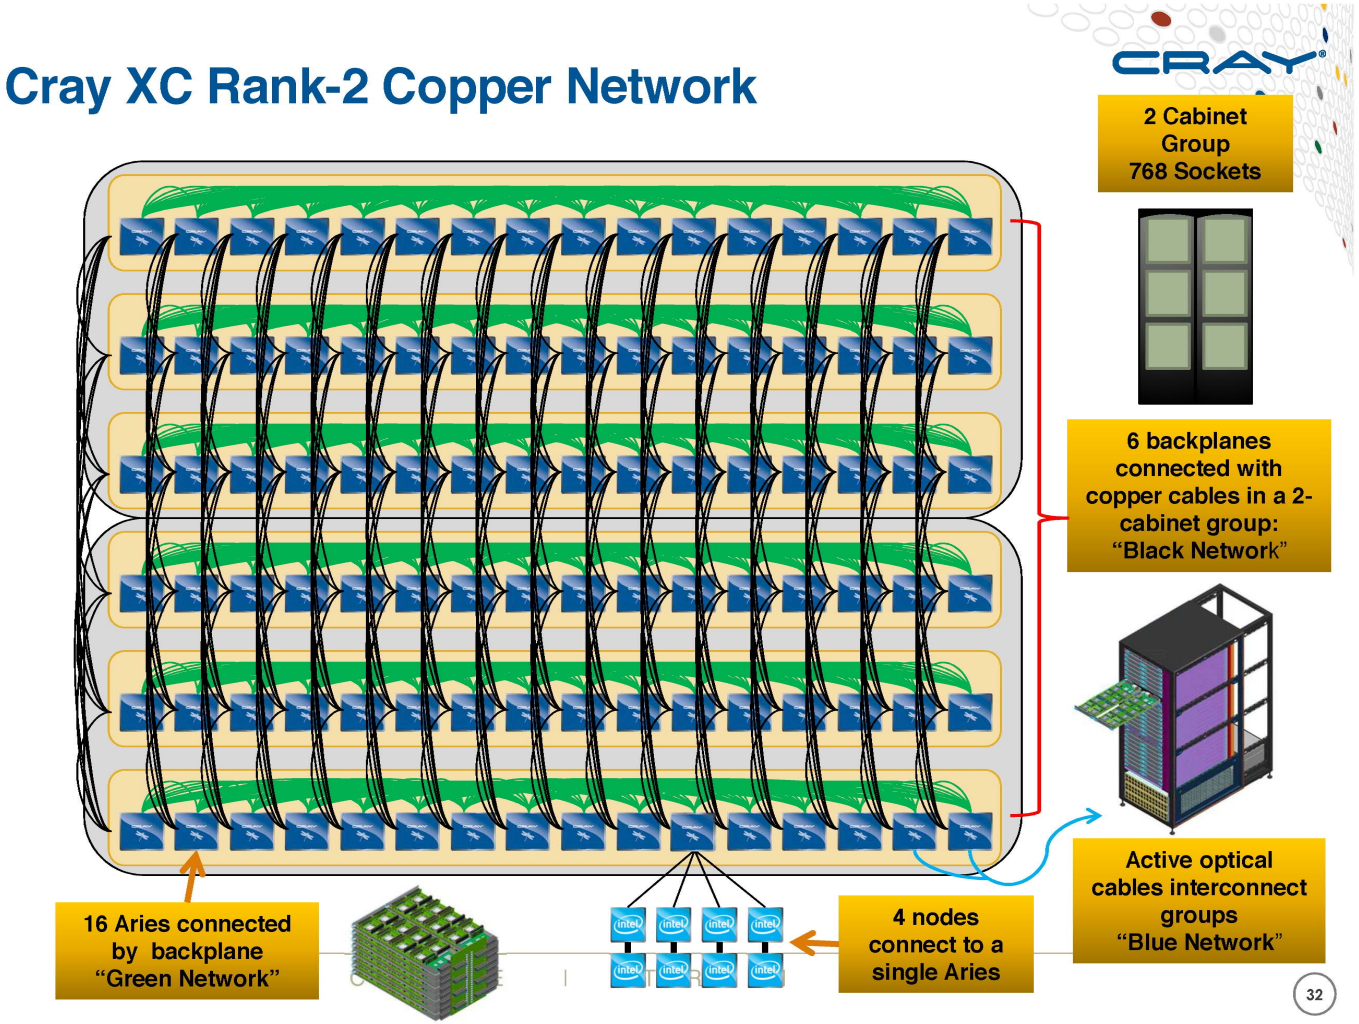

Shaheen II XC40 Interconnect

Cray XC Rank-2 Cabling

- Cray XC40 two-cabinet group
  - 768 Sockets
  - 96 Aries Chips
- All copper and backplanes signals running at 14 Gbps

Shaheen II XC40 Interconnect

Cray XC Network Overview – Rank-3 Network

- An all-to-all pattern is wired between the groups using optical cables (blue network)
- Up to 240 ports are available per 2-cabinet group
- The global bandwidth can be tuned by varying the number of optical cables in the group-to-group connections

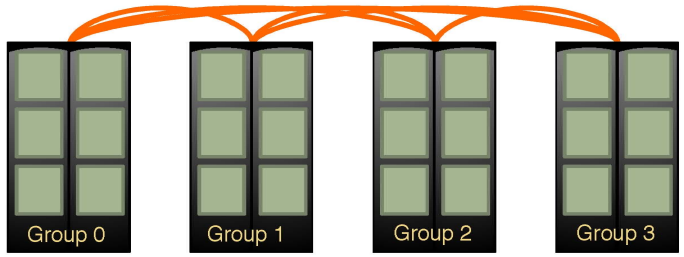

Example: An 4-group system is interconnected with 6 optical "bundles". The "bundles" can be configured between 20 and 80 cables wide

Shaheen II XC40 Interconnect

Cray XC Routing

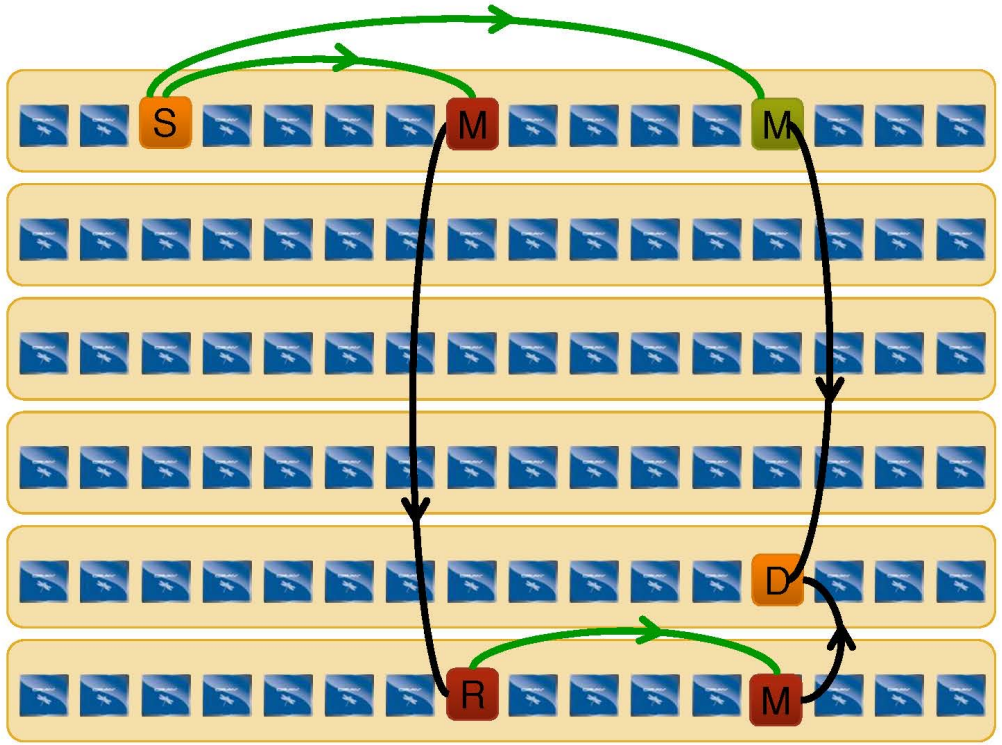

Minimal routes between any two nodes in a group are just two hops

Non-minimal route requires four hops.

With adaptive routing we select between minimal and non-minimal paths based on load

The Cray XC40 Class-2 Group has sufficient bandwidth to support full injection rate for all 384 nodes with non-minimal routing

Shaheen II Scratch Parallel Filesystem

- Cray Sonexion 2000 Storage System consisting of 12 cabinets containing a total of 5988 4TB SAS disk drives.
- The cabinets are interconnected by FDR InfiniBand Fabric .
- Each cabinet can contain up to 6 Scalable Storage Units (SSU); Shaheen II has a total of 72 SSUs.
- As there are 2 OSS/OSTs for each SSU, this means that there are 144 OSTs in total
- /scratch is in the lustre Parallel Filesystem

Shaheen II Project Parallel Filesystem

- HPE Cray ClusterStor E1000 Storage System consisting of 5 cabinets containing a total of 3392 16TB hard drives.
- The cabinets are interconnected by InfiniBand Fabric to Shaheen II and Ibex.
- The project storage is backed up in KSL tape library with a 4PB zero Watt storage as a cache
- /project is now in the lustre2 Parallel Filesystem and read only from compute nodes

Neser pre-post processing cluster

- Cray CS500 Cluster Introduced in September 2018
- 792 physical compute cores, achieving a peak performance of about 50TFlops/s
- 20+ compute nodes:
  - 19 CPU nodes with 40 Cores each, i.e two Skylake Intel Xeon(R) Gold CPU 2.0 GHz, 192 GB of memory and 12TB of local disk
  - 2 nodes are equipped with 768 GB of memory
  - 1 GPU node with 16 Nvidia Tesla K80 (to be decommissioned)
  - 6 nodes with ARM CPUs from Fujitsu and HBM memory (New)
- Connected with FDR InfiniBand to both Shaheen Lustre parallel file systems for /project and /scratch

Shaheen 101 HPC Training

Agenda

- Overview of Shaheen and Neser Architecture
- **How to get an account and access to Shaheen**
- Shaheen Storages: changes and policies storage
- Running environment and Job Scheduling
- Visualizations tools
- Running HPC workflow on Shaheen: VASP example
- Applications software example: CFD example
- Programming environment, debugging and profiling, Best practices and tips
- Questions and open discussions with KSL team

Follow us on Twitter: @KAUST_HPC