HPC2N @ Umeå University

Introduction to HPC2N and Kebnekaise

Jerry Eriksson, Pedro Ojeda-May, and Birgitte Brydsö

Outline

- Short presentation of HPC2N

- HPC at a glance.

- HPC2N – Abisko, **Kebnekaise**

- HPC Programming models – how to develope your own code
  (separate slides packages)
  - Nvidia,GPU: OpenAcc, Cuda
 - Intel OpenMP

HPC2N "HPC to North"

- A national center for Parallel and Scientific Computing

- Five partners:
  - Luleå University of Technology
  - Mid Sweden University
  - Swedish Institute of Space Physics
  - Swedish University of Agricultural Sciences - SLU
  - Umeå University
- Funded by the **Swedish Research Council (VR)** and its Meta-Center **SNIC** togheter with the **partners.**

From macro scale to micro scale

- Provides state-of-the-art resources and expertise for Swedish eScience

  - Scalable and parallel HPC

  - Large-scale storage facilities

  - Grid and cloud computing

  - Software and advanced support for eScience applications

  - International network for research and development
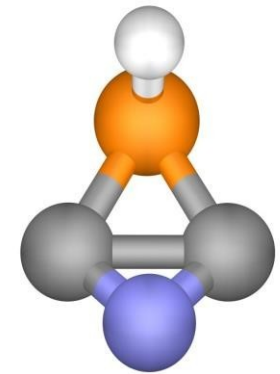
DFT computation, semi-stable, binding energy 15eV; Sven Öberg, LTU

Main areas of HPC2N users

- Biosciences and medicine
- Chemistry
- Computing science
- Engineering
- Materials science
- Mathematics and statistics
- Physics including space physics
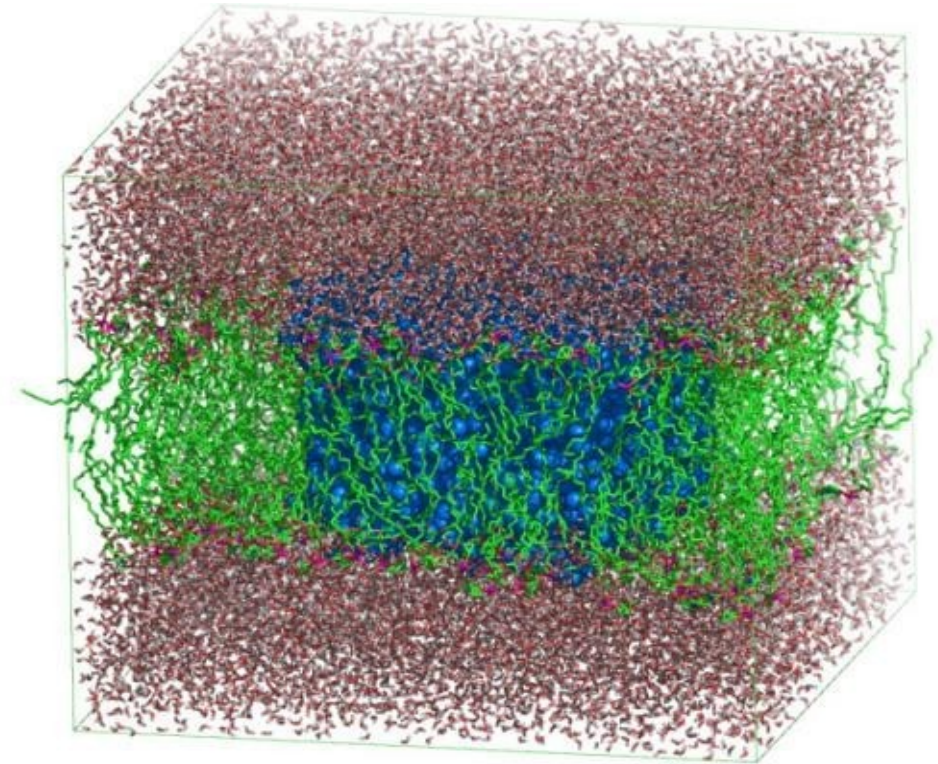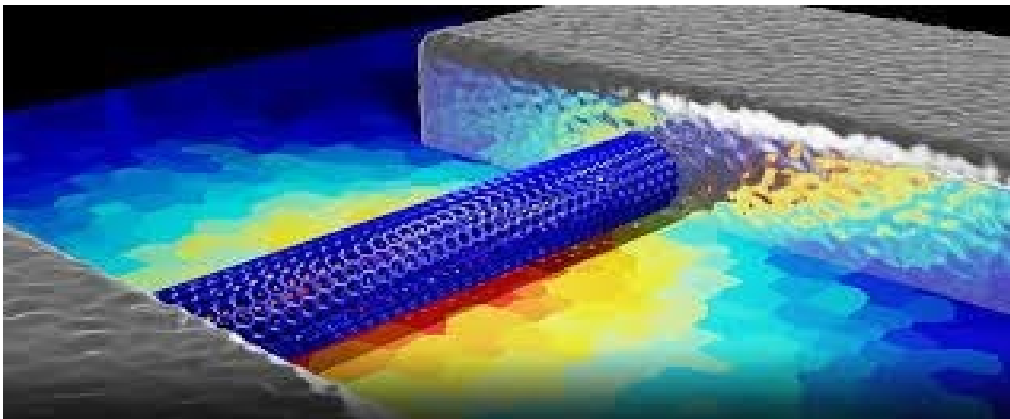
Storage Levels @ HPC2N

Basically three types of storage are available at HPC2N:

- Center Storage - Parallel file system (fast discs)
  - Closely connected to our computing resources; Abisko and Kebnekaise

- SweStore - disk based (dCache)
  - part of SNIC Storage, responsible for national accessible storage

- Tape Storage
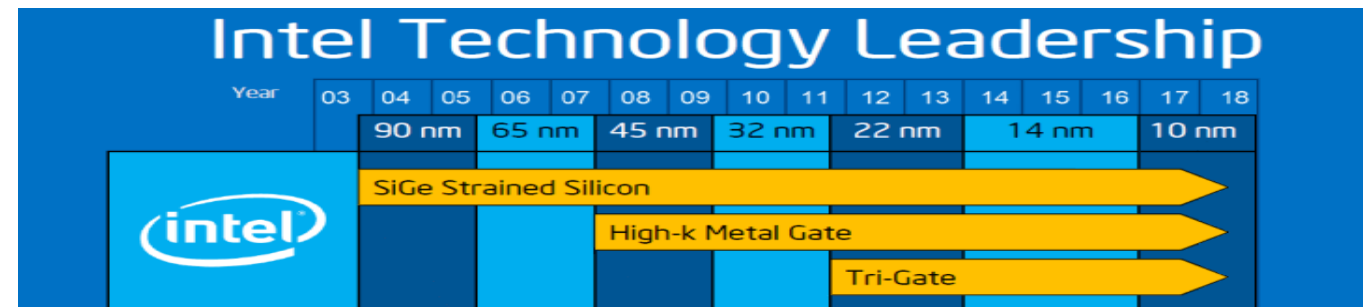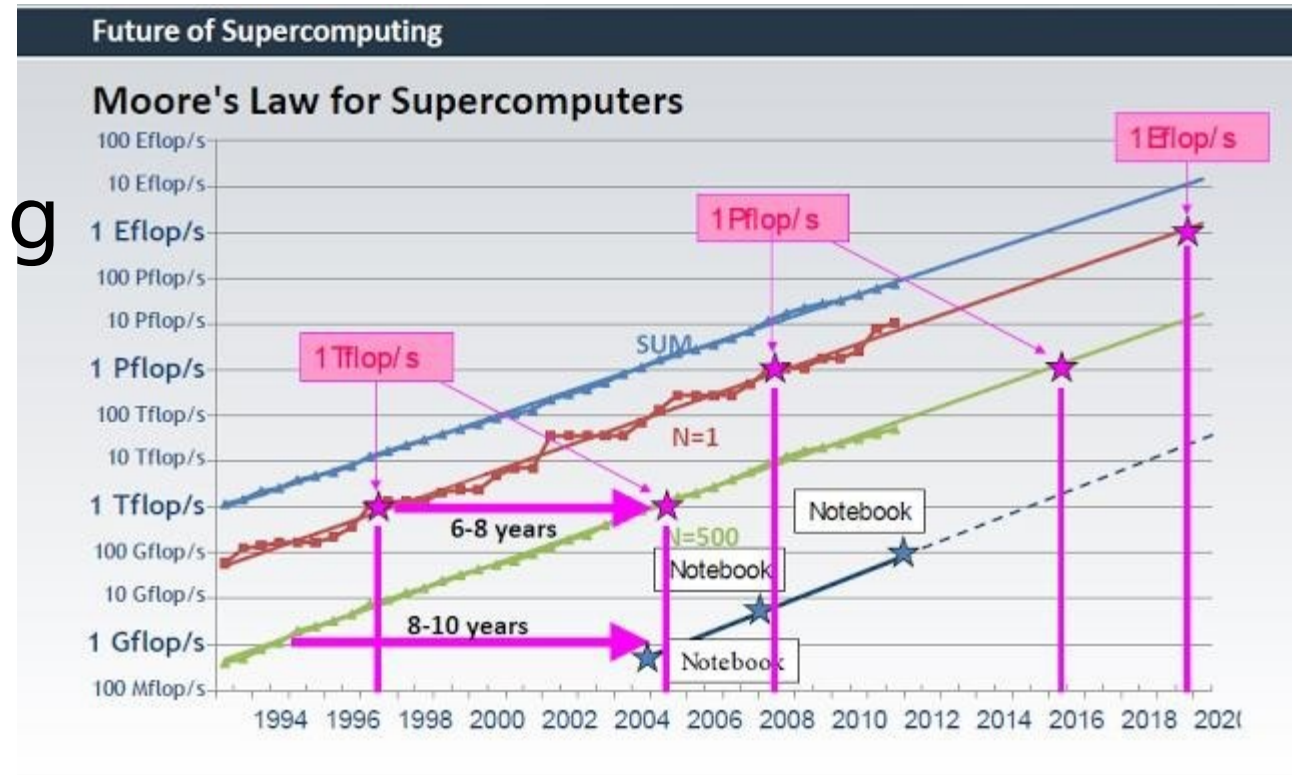 - Backup
  - Long term storage

HPC2N Think Tank!

- User support (primary, advanced, tailored)

    - Research group meetings @ UmU

- User training and education program

- Workshops & Colloquia

- Research & Development - Technology transfer

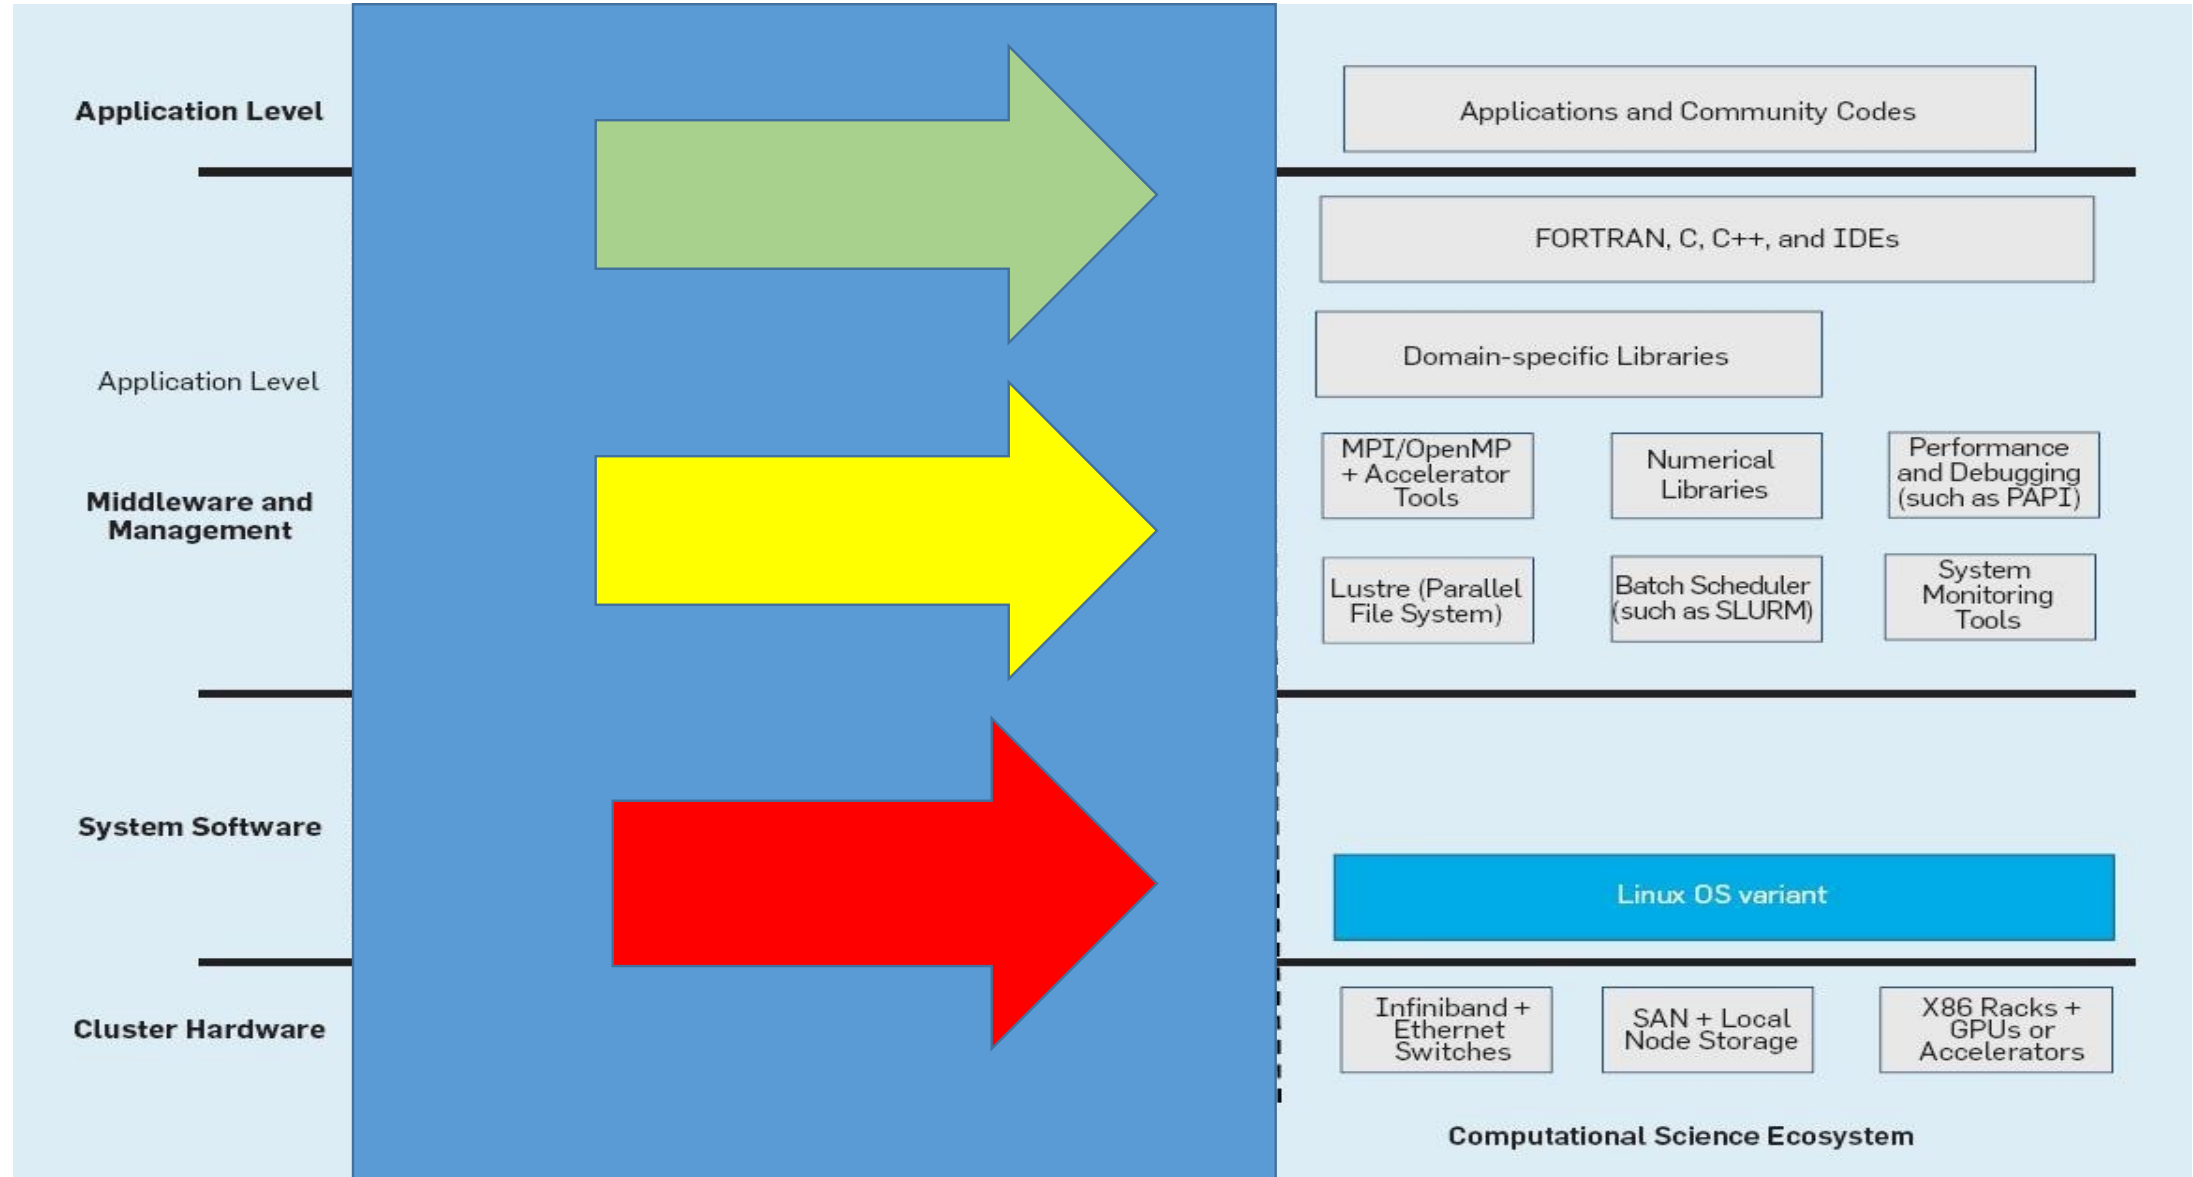- Provide various state-of-the-art HPC resources

HPC– Towards Exascale Computing

- **Moore's law**: the number of transistors in a chip doubles every second years.
- Parallel Computing:
  - Increase number of cores.
- Heterogenous clusters
  - Different processors and memories.
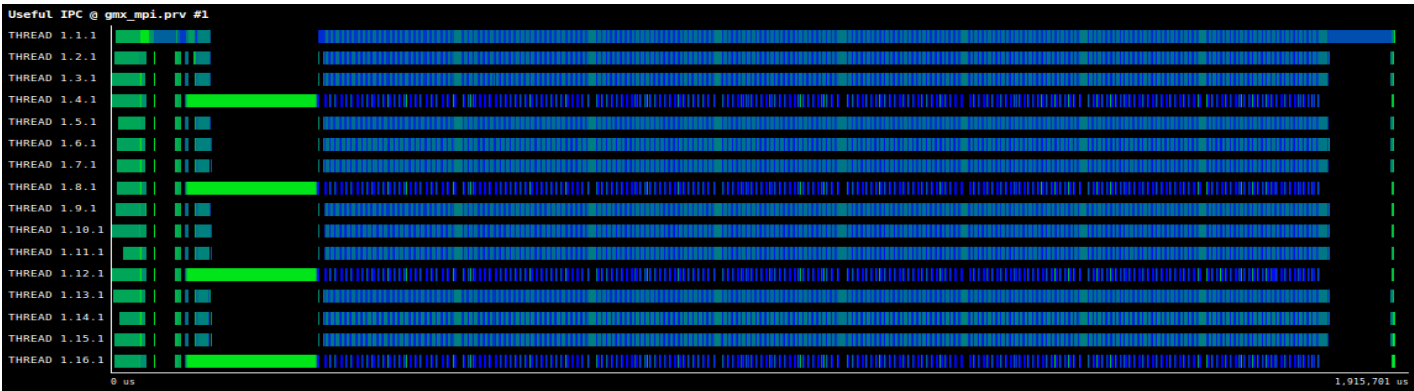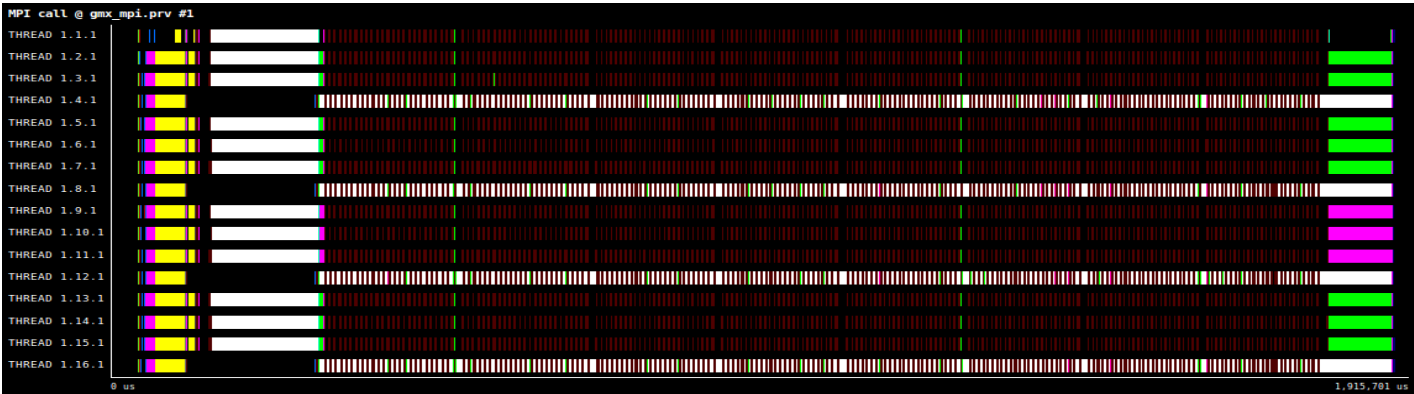- Power efficiency !

Moore's Law for Supercomputers

Intel Technology Leadership

HPC EcoSystems

Application Level

Application Level

Middleware and Management

System Software

Cluster Hardware

Applications and Community Codes

FORTRAN, C, C++, and IDEs

Domain-specific Libraries

MPI/OpenMP + Accelerator Tools

Numerical Libraries

Performance and Debugging (such as PAPI)

Lustre (Parallel File System)

Batch Scheduler (such as SLURM)

System Monitoring Tools

Linux OS variant

Infiniband + Ethernet Switches

SAN + Local Node Storage

X86 Racks + GPUs or Accelerators

Computational Science Ecosystem

PRACE - Partnership for Advanced Computing in Europe

Tracing tools (GROMACS, 16 Cores)

Now to the clusters and programming m

A large amount of numbers and technical information will follow!!

Relax, you do not need to now everything in detail, and we offer training for those things you should know.

Abisko

- 332 nodes with a total of 15936 CPU cores.
- AMD Opteron 6238 (Interlagos)
- The 10 'fat' nodes have 512 GB RAM each, and the 322 'thin' nodes have 128 GB RAM each.
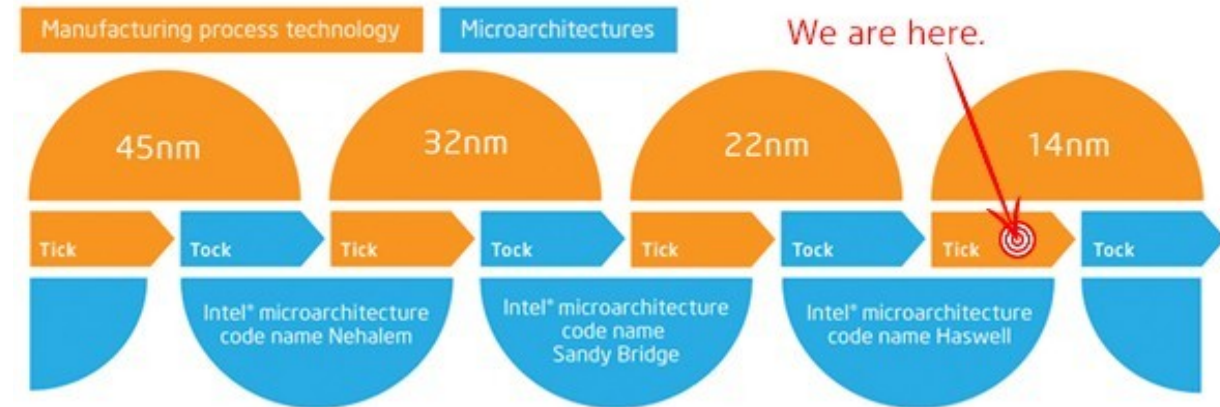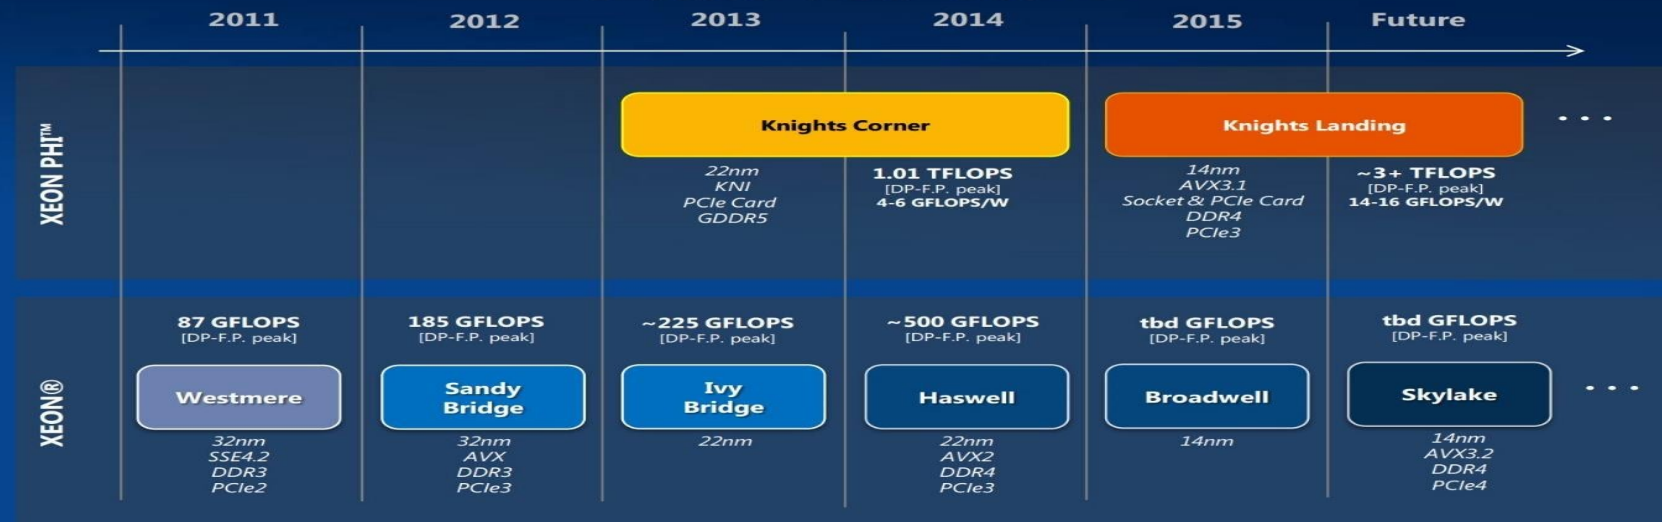- (More details can be found on our web-pages)

Kebnekaise

Intels processors

Compute nodes

- 432 nodes
- Intel Broadwell   (E5-2690v4)
- 2x14 cores/node
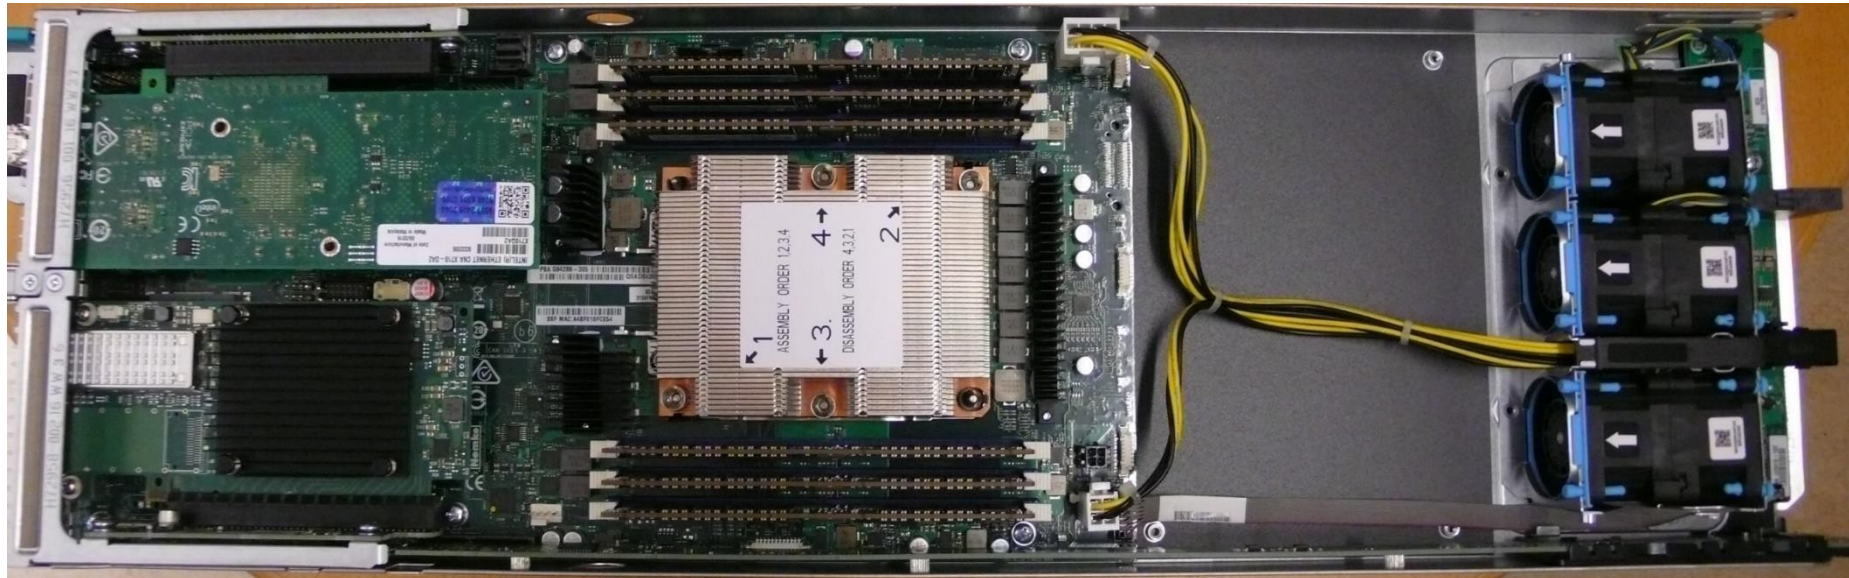- 128GB memory
- Infiniband FDR

Large memory nodes

- 20 nodes

- Intel Broadwell (E7-8860v4)

- 4x18 cores/node

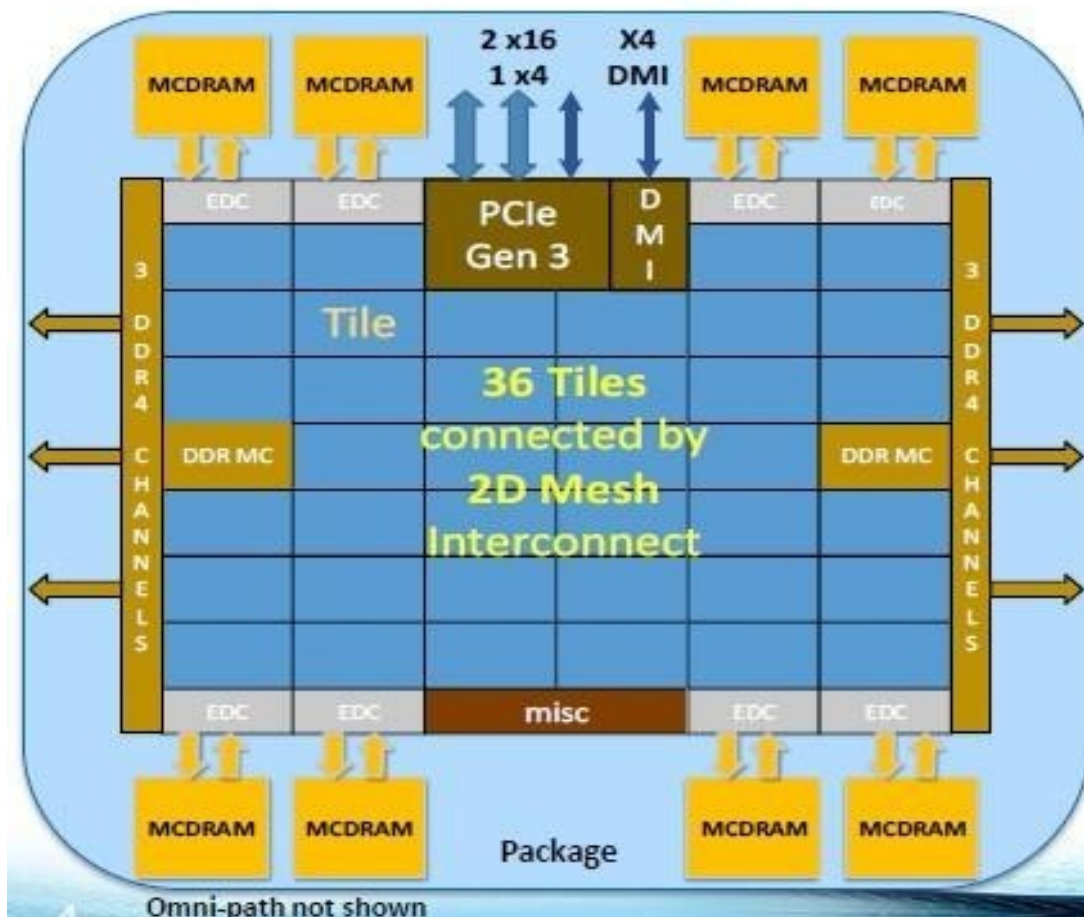- 3TB memory

- Infiniband EDR

KNL - Intel Knights Landing

- 36 nodes
  - 68 cores
  - 1.4GHz (1.2GHz AVX)
- 192 GB memory - 16 GB MCDRAM
- Infiniband FDR
- *Installation in February*

Knights Landing Overview

TILE

2 VPU	CHA	2 VPU
Core	1MB L2	Core

2 x16
1 x4

X4 DMI

MCDRAM MCDRAM   PCIe Gen 3   D M I   MCDRAM MCDRAM

EDC EDC EDC EDC

3 D D R 4 C H A N N E L S

DDR MC

Tile

36 Tiles connected by 2D Mesh Interconnect

DDR MC

3 D D R 4 C H A N N E L S

EDC EDC misc EDC EDC

MCDRAM MCDRAM   Package   MCDRAM MCDRAM

Omni-path not shown

Chip: 36 Tiles interconnected by **2D Mesh**

Tile: 2 Cores + 2 VPU/core + 1 MB L2

Memory: MCDRAM: 16 GB on-package; High BW
   DDR4: 6 channels @ 2400 up to 384GB

IO: 36 lanes PCIe Gen3. 4 lanes of DMI for chipset

Node: 1-Socket only

Fabric: Omni-Path on-package (not shown)

Vector Peak Perf: 3+TF DP and 6+TF SP Flops

Scalar Perf: ~3x over Knights Corner

**Streams Triad (GB/s):** MCDRAM : 400+; DDR: 90+

Intel Xeon Phi

Xeon Phi	Clock Speed	Cores / Threads	Peak DP TFLOPS	DDR4 Memory	MCDRAM Memory		TDP (Watts)	1K Tray Unit Price	$ / TFLOPS
					Capacity	Speed			
Knights Landing									
7290	1.5 GHz	72 / 288	3.46	384 GB	16 GB	7.2 GT/sec	245	$6,254	$1,810
7250	1.4 GHz	68 / 272	3.05	384 GB	16 GB	7.2 GT/sec	215	$4,876	$1,601
7230	1.3 GHz	64 / 256	2.66	384 GB	16 GB	7.2 GT/sec	215	$3,710	$1,393
7210	1.3 GHz	64 / 256	2.66	384 GB	16 GB	6.4 GT/sec	215	$2,438	$916

Xeon Phi	Clock Speed	Cores / Threads	Peak DP TFLOPS	Cache Memory	GDDR5 Memory		TDP (Watts)	1K Tray Unit Price	$ / TFLOPS
					Capacity	Speed			
Knights Corner									
7120P	1.24 GHz	61 / 61	1.21	30.5 MB	16 GB	5.5 GT/sec	300	$4,129	$3,412
7120X	1.24 GHz	61 / 61	1.21	30.5 MB	16 GB	5.5 GT/sec	300	$4,129	$3,412
5110P	1.05 GHz	60 / 60	1.01	30 MB	8 GB	5.0 GT/sec	225	$2,649	$2,623
5120D	1.05 GHz	60 / 60	1.01	30 MB	8 GB	5.5 GT/sec	245	$2,759	$2,732
3120A	1.10 GHz	57 / 57	1.0	28.5 MB	6 GB	5.0 GT/sec	300	$1,695	$1,695
3120P	1.10 GHz	57 / 57	1.0	28.5 MB	6 GB	5.0 GT/sec	300	$1,695	$1,695

General or special-purpose processor ?

GPU nodes

- 32 nodes with 2x NVidia K80

- 4 nodes with 4x NVidia K80

- Intel Broadwell 2x14 cores (E5-2690v4)

- 128 GB memory

- Infiniband FDR
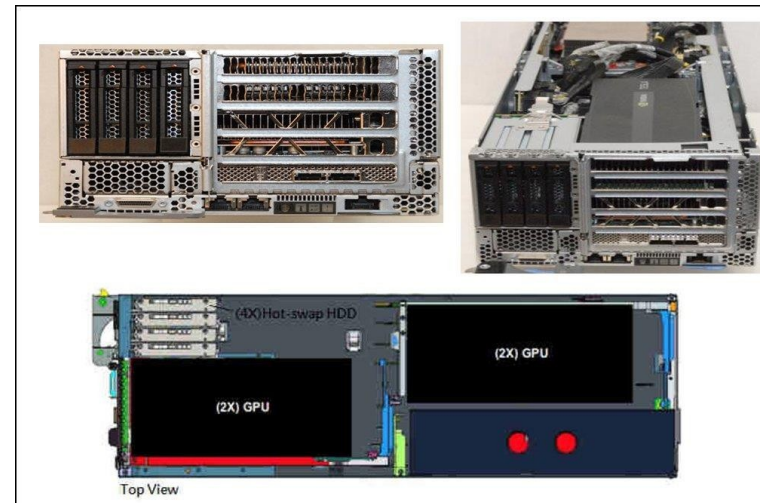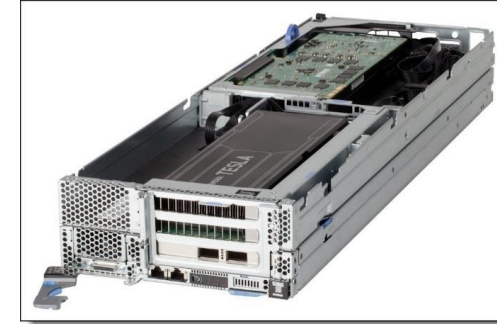
High Speed Interconnect

- Infiniband
- Three level fat tree structure
- FDR cards in nodes (leafs)
- EDR cards in large memory nodes
- EDR in switches

Kebnekaise in numbers

- 13 racks
- 544 nodes
- 17552 cores (of which 2448 cores are KNL-cores)
- 399360 CUDA cores (80 * 4992 cores/K80)
- More than 125TB memory (20*3TB + (432 + 36) * 128GB + 36*192GB)
- 66 switches (Infiniband, Access network, Management network)

Kebnekaise in numbers

- 83% of the system are standard and Large Memory nodes
- 7% GPU-nodes
- 7% KNL-nodes
- 4% Other nodes (login and management nodes, LNET-routers etc)
- 728 TFlops/s Peak performance
- 629 TFLops/s HPL (all parts)
- HPL: 86% of Peak performance

Standard Nodes	374 TFlops/s
Large Memory Nodes	34 TFlops/s
2xGPU Nodes	129 TFlops/s
4xGPU Nodes	30 TFlops/s
KNL Nodes	62 TFlops/s
Total (All parts)	**629 Flops/s**