



**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

**ETH** zürich



# Future opportunities and challenges using OpenPOWER

European OpenPOWER Summit (ADG Session)

Sadaf Alam, CSCS

October 28, 2016

# Piz Daint (5272 Intel Xeon & Nvidia K20X GPU nodes)



# MeteoSwiss New Weather Supercomputer

World's First GPU-Accelerated  
Weather Forecasting System

<http://www.c2sm.ethz.ch/news/news/2015/09/c2sm-partner-meteoswiss-is-using-a-new-gpu-based-computer-architecture-to-produce-improved-weather-forecast.html>

2x Racks


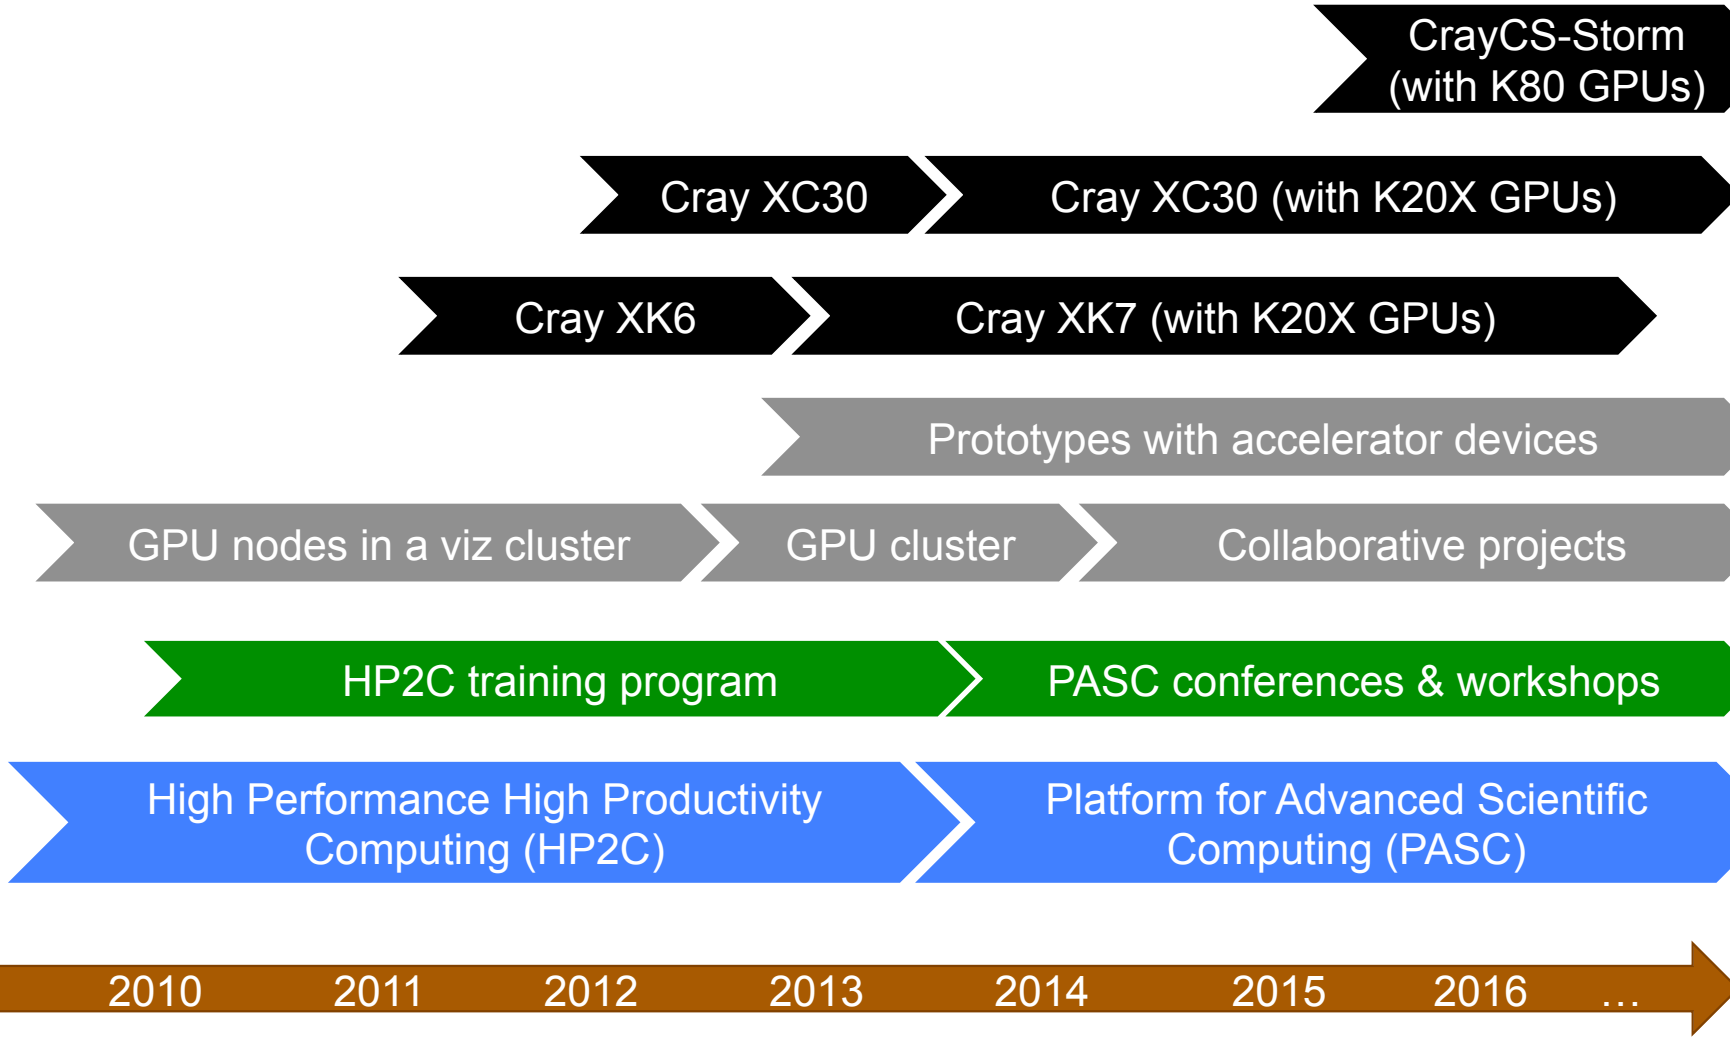
48 CPUs

192 Tesla K80 GPUs

> 90% of FLOPS from GPUs

Operational in 2016





\* Timelines & releases are not precise

Tesla Deployment Kit (TDK),  
NVML & healthmon Ganglia  
plugins

Cray PMDB,  
Node Health  
Check & RUR

GPU Data Centre  
Manager (DCGM),  
additional tools

*Custom solutions &  
integration at CSCS  
on case-by-case basis*

### GPU-enabled MPI & MPS

GPUDirect      GPUDirect-RDMA

OpenACC 1.0

OpenACC 2.0

OpenACC 2.5

OpenCL 1.0

OpenCL 1.1

OpenCL 1.2

OpenCL 2.0

OpenCL 2.x

CUDA 2.x  
CUDA 2.x  
CUDA 2.x  
CUDA 2.x

CUDA 2.x  
CUDA 2.x  
CUDA 3.x

CUDA 4.x  
CUDA 4.x  
CUDA 4.x

CUDA 5.x  
CUDA 5.x

CUDA 6.x

CUDA 7.x

CUDA b.x

Cray CS-Storm  
(dense K80)

Cray XK6

Cray XK7

Cray XC30 & hybrid XC30

X86 cluster with  
C2070, M2050,  
S1070

iDataPlex  
cluster  
M2090

Testbed with  
Kepler &  
Xeon Phi



2009      2010      2011      2012      2013      2014      2015      2016      ...

Requirements analysis

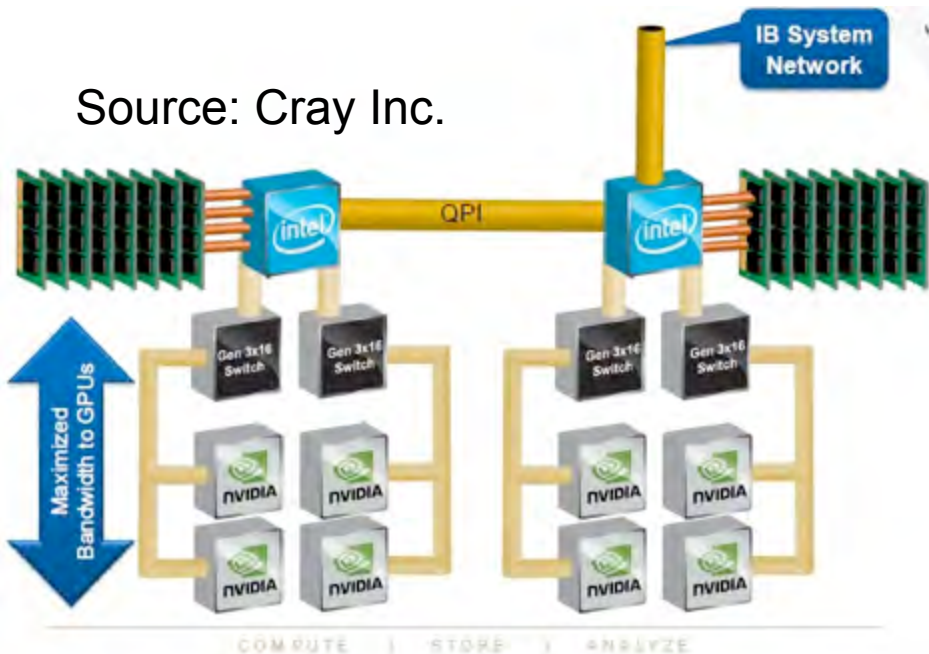
Applications development and tuning

24/7 monitoring & troubleshooting


\* Timelines & releases are not precise

# Opportunities (Performance)

Source: Cray Inc.

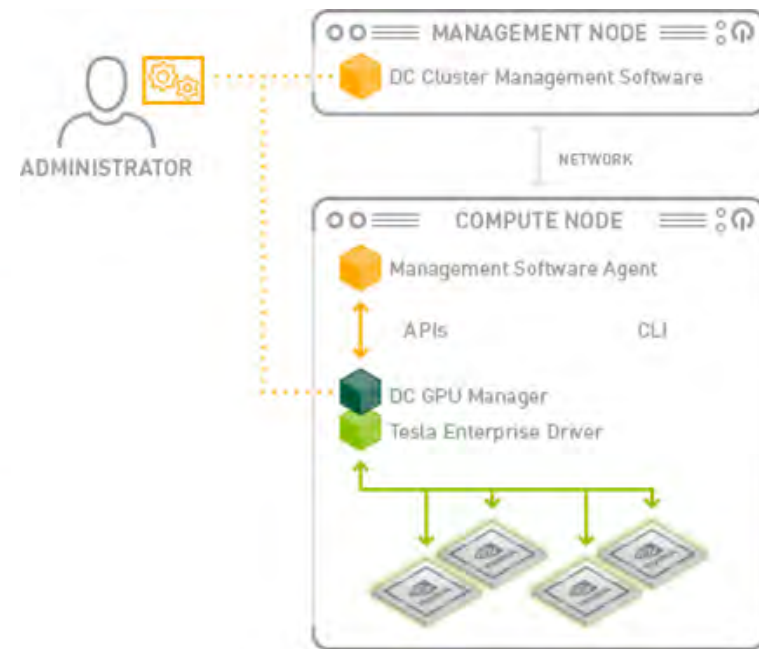


Server Block Diagram  
Microway OpenPOWER Server with NVIDIA Tesla P100 NVLink GPUs



Source: Microway

# Opportunities (Monitoring)



Detecting and Managing GPU Failures, Nick Cardo, CUG, 2015

<https://devblogs.nvidia.com/parallelforall/nvidia-data-center-gpu-manager-cluster-administration/>

# Challenges (accessibility, commodity & community)

- CSCS GPU based systems
  - X86 based (AMD & Intel)
  - Commodity (InfiniBand) plus proprietary (Cray)
  - Lustre
  - ... single node environment typically identical to a laptop
- Code development environment and tools
  - Similar on clusters, Cray, laptop, ...
  - Containers with GPUs
- Leveraging community solutions and connecting with widely accessible platforms
  - Connecting bottom up and top down



# Challenges (interoperability)

- There are many many logos on the foundation 😊
- A user concerns
  - Most already mentioned related to performance portability and interoperability
  - Resource management systems
- An operational concern aka operational management
  - Resource management
  - Scheduling systems
  - Log analysis
  - Performance analysis
  - Provisioning
  - ...



**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

**ETH** zürich



**Thank you for your attention.**