

The DEEP-EST Project.

Viktor Khristenko, Maria Girone (CERN)

Outline

- What is the DEEP-EST / going to be?!
 - How does CERN/CMS participate?!
- What is currently available?!
- CERN/CMS, what benefits for us!?
- Current Ideas

What is the DEEP-EST Project?!

- Dynamical Exascale Entry Platform (DEEP) - Extreme Scale Technologies (EST)
- 3rd generation of the DEEP Project
- EU funded project, which aims to build Modular Supercomputer Architecture (MSA)
- **MSA = Heterogeneity of resources (processors, coprocessors, GPU, FPGA, etc...)**
- **Components:**
 - **Cluster Module - CPUs**
 - **Booster - CoProcessors Intel KNHs**
 - **The Data Analytics Module - CPUs, CPU+FPGA or coprocessors + “””huge amount of (DRAM+non-volatile) memory per core”””**

What is the DEEP-EST going to be?

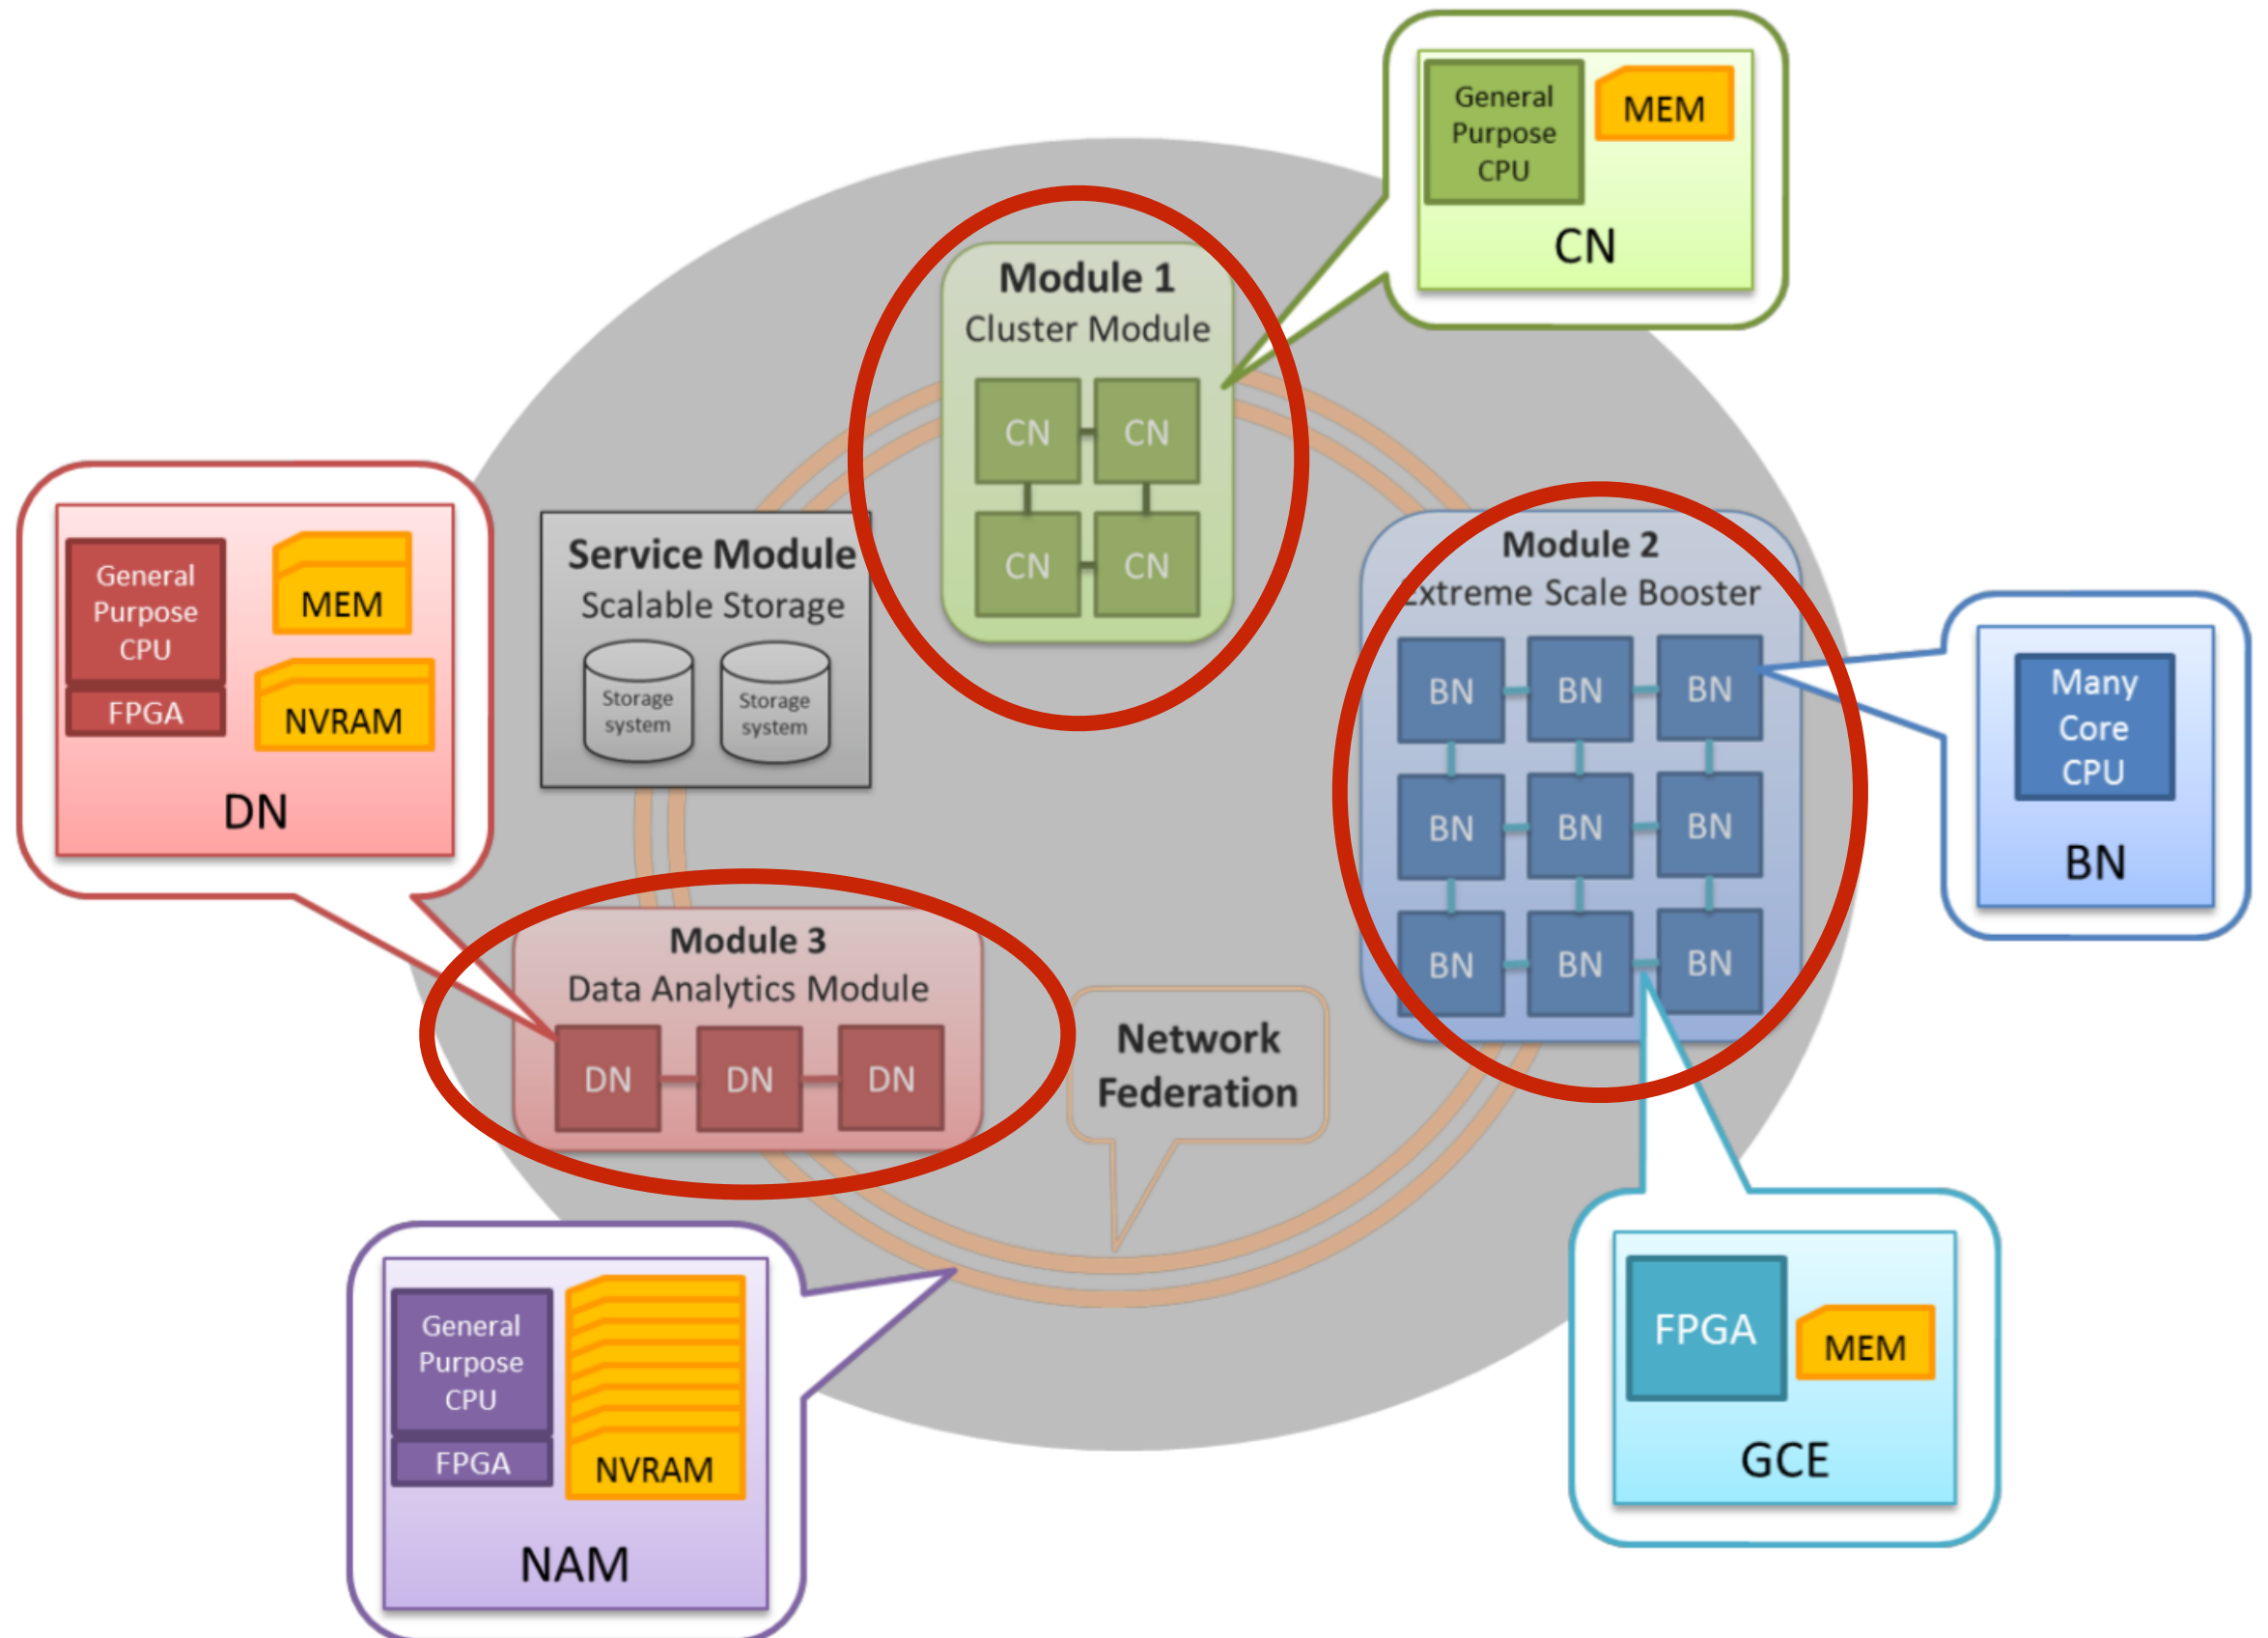


Figure 2: High level hardware configuration of the DEEP-EST prototype

How does CERN/CMS participate?!

- 6 applications participate in the DEEP-EST Project
 - Neuroscience / Weather Analytics / Biology Simulations / Astronomy / High Energy Physics
- CMS is a new member of this set!
- DEEP-EST wants CMS to provide various workflows to run on their system
- DEEP-EST does not know about what kind of physics workflows we run!
- Simulations -> Reconstruction -> Particle Flow -> Data Analysis -> anything!

What is currently available?!

- knc1:
 - 2 Xeon CPUs
 - 64 GB memory
 - 4 KNCs (named knc1-mic[0-3]) with 61 cores and 16 GB each
- knc2:
 - 2 Xeon CPUs
 - 64 GB memory
 - 2 KNCs (named knc2-mic[0-1]) with 57 cores and 6 GB each

Xeon CPUs +
KNC coProcessors

DEEP:

- Cluster:
 - 2 Xeon CPUs per node
 - 32 GB memory per node
 - 128 nodes
- Booster:
 - 2 KNCs per BNC (Booster Node Card)
 - 16 GB per KNC
 - 192 BNCs

Xeon CPUs +
KNC coProcessors

SDV: ↑

- Cluster:
 - 2 Xeon Haswell CPUs per node
 - 128 GB per node
 - plus 1 NVMe with 400 GB per node
 - 16 nodes
- KNLs:
 - 8 KNLs
 - 4 with EXTOLL
 - 2 with NVMe
 - 96 GB per KNL

Cluster: Xeon Haswell CPUs

Booster: KNLs
EXTOLL - network

DEEP-ER prototype (it is planned like this):

- Cluster Nodes:
 - The nodes from the SDV will be used as Cluster Nodes.
- Booster Nodes:
 - 3-4 chassis with 18 KNL boards each
 - 1 KNL per board
 - 96 GB per KNL
 - plus 1 NVMe with 400 GB per KNL

How does CERN/CMS benefit from the DEEP-EST Project

- We do not want to just benchmark CMSSW against the DEEP-EST system. We want to contribute!
- 1 Fellow working full-time on the R&D of the CMS Software eco-system targeting HL-LHC!
- DEEP-EST is agnostic about the HEP workflows -> advantage for us in terms of what we want to test out/run.
- A bunch of hardware (promised to be) that we can test against!
- A collaboration with sciences outside of HEP

Ideas

- CMS Big Data Project
 - Low weight ROOT I/O (spark-root) + Apache Spark for Physics Data Analytics
 - Physics ML Pipelines with Apache Spark + ...
- DEEP-EST is about heterogeneity => heterogeneity for CMS software framework