



montblanc-project.eu | @MontBlanc_EU

The Dibona prototype platform


*an industrial & highly scalable ThunderX2
integration*

Joël Wanza (Atos)



Motivations

- ▶ Increasing demand of alternative to x86 hegemonic position in HPC
- ▶ expecting a better energy efficiency (and pricing)
- ▶ strong and vibrant ecosystem inherited from embedded community



The Mont-Blanc prototypes

The Mont-Blanc prototype ecosystem

Prototypes are critical to accelerate software development
System software stack + applications

MONT-BLANC

PRACE prototypes

- Tibidabo
- Carma
- Pedraforca

Mini-clusters

- Arndale
- Odroid XU
- Odroid XU-3
- NVIDIA Jetson



Mont-Blanc prototype

- 1080 compute cards
- Dual Cortex-A15
- GPU Mali-T604
- 4 GB LPDDR3
- Up to 64 GB local storage
- USB-to-Eth network
- Fine grained power monitoring system
- Installed between Jan and May 2015
- Operational since May 2015 @ BSC



Bull
atos technologies


Mont-Blanc 3

- Bull Sequana™
- Cavium ThunderX2™




From Bull Sequana X1000 ...

- ▶ free water cooling system
(3 l/min via 4 liquid channels per blade)
- ▶ 2 * 48 compute blade
- (2 * 144 bi-socket nodes)
- ▶ 256GB/s sideplane (NIC ↔ L1)
- ▶ 256GB/s copper cable (L1 ↔ L2)
- ▶ ...



... to Mont-Blanc dibona platform

- ▶ 16 blades * 3 bi-socket compute nodes

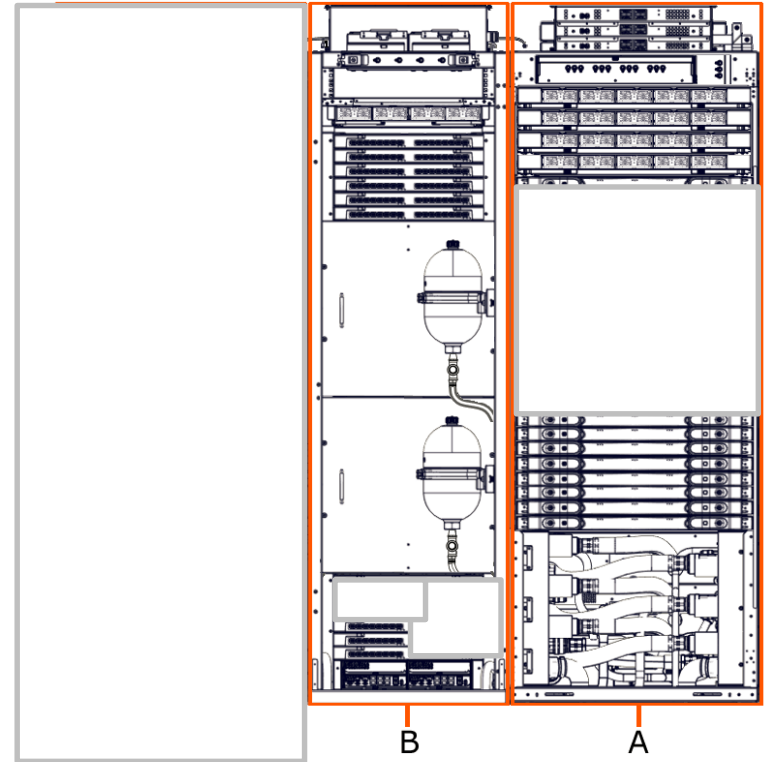


A Base compute rack cabinet

B Switch rack cabinet


CN: compute node

LX: level X switch

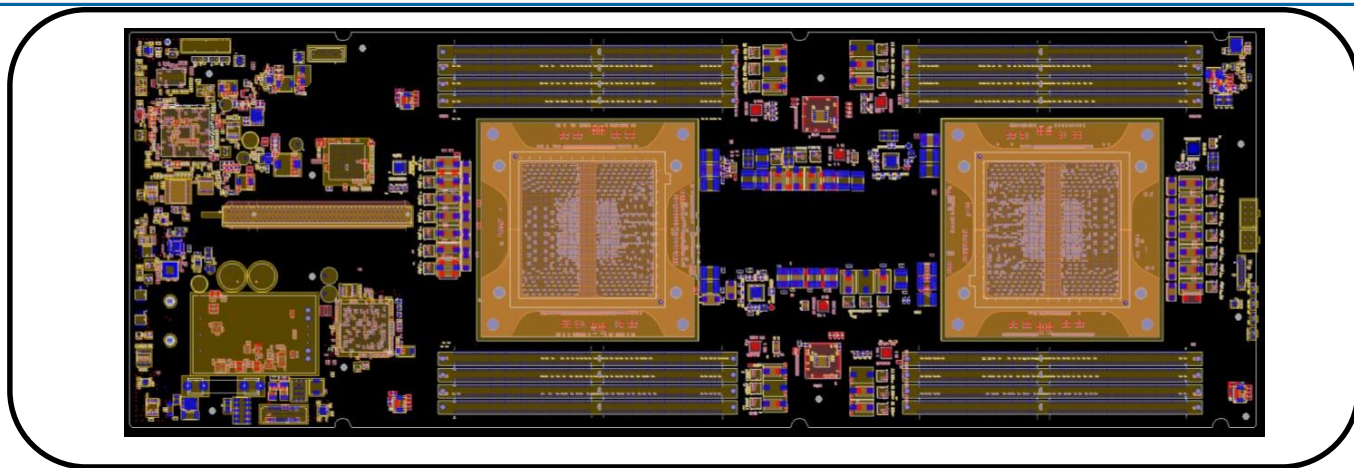


Cavium ThunderX2 based compute blade

- ▶ *Compute blade density* optimized for 2 sockets with up to 8 memory channels per socket in 1U
- ▶ *Interconnection technology* is Infiniband EDR using Mellanox 36-port switches:
 - One Mellanox IB EDR NIC is available per compute node
 - Passive copper cables are used for node connections inside the cell
 - Optical cables are used for cell connection to upper level, and IO/service node connection at L1 level
- ▶ *Interconnection topology* - Fat-tree with a pruning factor of $\frac{1}{2}$ at L1 level



Cavium ThunderX2 based compute node



- ▶ 2 ThunderX2 CPU (32 cores per CPU, 32MB L3 cache)
- ▶ 16 DDR4 DIMM slots (8 channels per CPU)
- ▶ Aspeed AST2500 BMC
- ▶ SATA drives (SSD only)
- ▶ FPGA for Reset & Power Sequencing and for sensors monitoring

Summary

→ The Mont-Blanc Dibona prototype:

- 48 nodes, each node includes 2 ThunderX2 processors SKU 155W
- 48 * 2 * 32 cores @ 2GHz
- 48 * 2 * 128 GB memory and 48 * 2 * 128 GB local storage (+ 8TB nfs)
- Fat tree interconnect topology with EDR 100Gb/s
- Theoretical peak performance of > 49 Tflops

→ Open platform for scientific research

→ BullSequana is ready to integrate further high end ARM based processors, from HW to the SW stack.



THANK Y



U !

joel.wanza-weloli@atos.net

Thanks

For more information please contact:

T+ 33 1 30 80 37 03

joel.wanza-weloli@atos.net

Atos, the Atos logo, Atos Codex, Atos Consulting, Atos Worldgrid, Bull, Canopy, equensWorldline, Unify, Worldline and Zero Email are registered trademarks of the Atos group. March 2017. © 2017 Atos.
Confidential information owned by Atos, to be used by the recipient only. This document, or any part of it, may not be reproduced, copied, circulated and/or distributed nor quoted without prior written approval from Atos.

Bull
atos technologies