Welcome!

European Exascale Processor & Memory Node Design

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 671578

Paving the way towards a highly energy-efficient and highly integrated compute node for the Exascale revolution: the ExaNoDe approach.

Kevin Pouget

Virtualization Engineer
Virtual Open Systems (ExaNoDe partner)

August 30th, 2017

Disclaimer: This presentation does not represent the opinion of the EC and the EC is not responsible for any use that might be made of information appearing herein.

 Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

Paving the way towards a highly energy-efficient and highly integrated **compute node** for **the Exascale revolution**: the ExaNoDe approach

Outline:

1. **ExaNoDe Overview**
2. **System Architecture and Integration**
3. **ExaNoDe Prototype**
4. **Application and Software**

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

Paving the way towards a highly **energy-efficient** and highly **integrated** **compute node** for **the Exascale revolution**: the ExaNoDe approach

Outline:

1. **ExaNoDe Overview**
2. **System Architecture and Integration**
3. **ExaNoDe Prototype**
4. **Application and Software**

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

The new European HPC research landscape

European Commission

PRACE

ETP 4 HPC

THE EUROPEAN TECHNOLOGY PLATFORM FOR HIGH PERFORMANCE COMPUTING

Source: Panagiotis TSARCHOPOULOS SC'15

The new European HPC research landscape

European Commission

PRACE

ETP 4 HPC

Source: Panagiotis TSARCHOPOULOS SC'15

The new European HPC research landscape

European Commission

PRACE

ETP 4 HPC

Source: Panagiotis TSARCHOPOULOS SC'15

ExaNoDe technological context

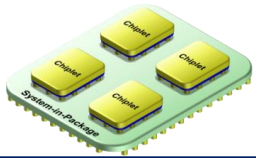

EURO SERVER

Architecture

(ended in Jan. 17)

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

ExaNoDe technological context

ExaNeSt

ECOSCALE

ExaNoDe

Compute node (HW + SW)

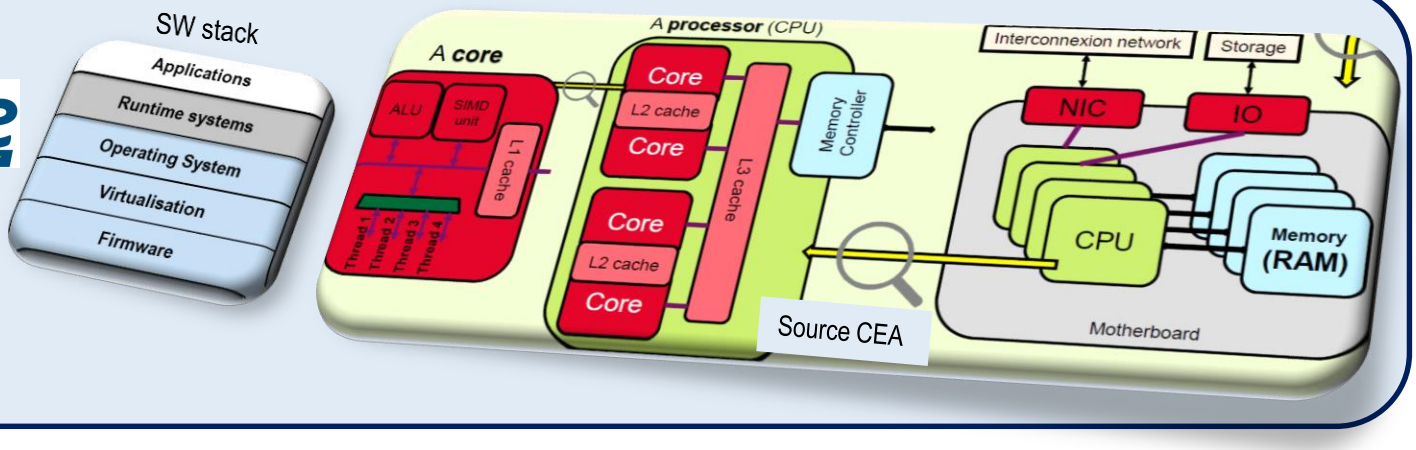

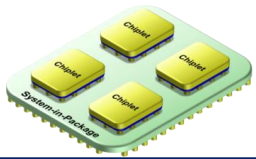

EURO SERVER

Architecture

(ended in Jan. 17)

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

ExaNoDe technological context

ExaNeSt

Interconnect
Storage

ECOSCALE

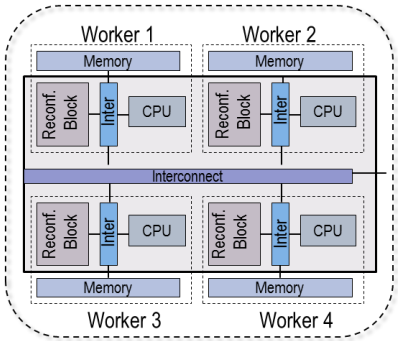

FPGA acceleration

ExaNoDe

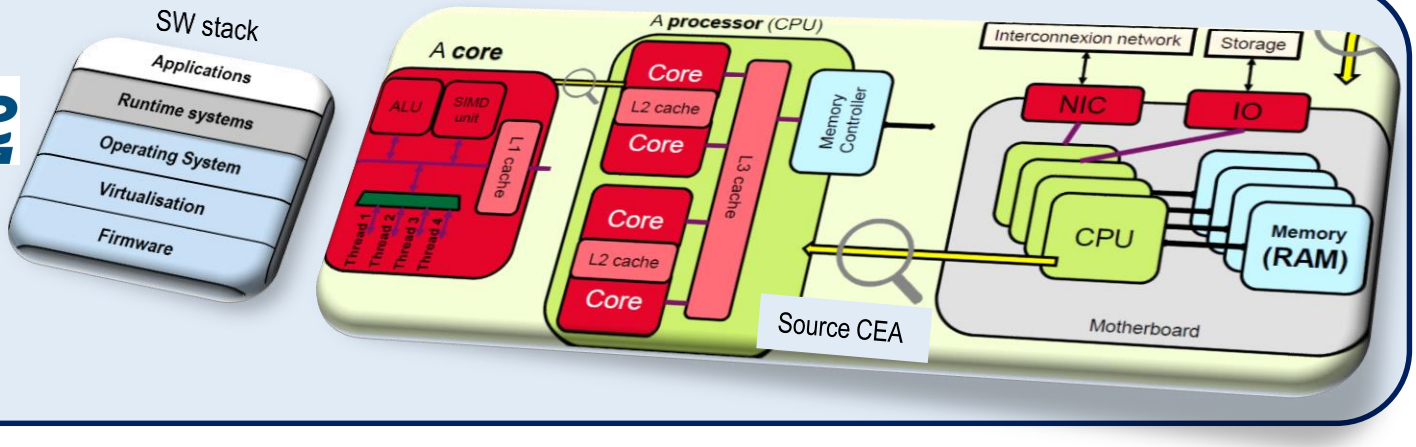

Compute node
(HW + SW)

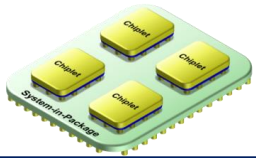

EURO SERVER

Architecture

(ended in Jan. 17)

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

Ecosystem from European Projects

ARMv8, UNIMEM for micro-servers

EURO SERVER

Architecture

Compute node

ExaNoDe

Compute node (HW & SW)

ECOSCALE

Acceleration with FPGA

ExaNeSt

Interconnect, storage & cooling

ExaNoDe as part of a global strategy

EURO SERVER

redesigns the entreprise server:

Lower cost through system integration

Energy efficiency : low-power 64 bit processor and more efficient software

Mutualization of I/O resources

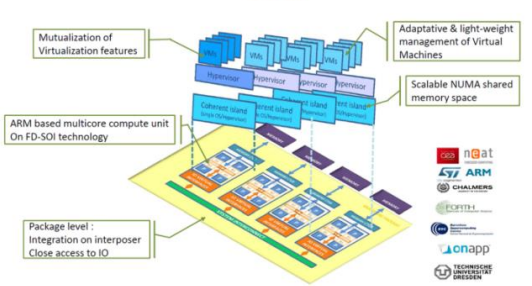

Source: Isabelle Dor

ECOSCALE

focuses on acceleration

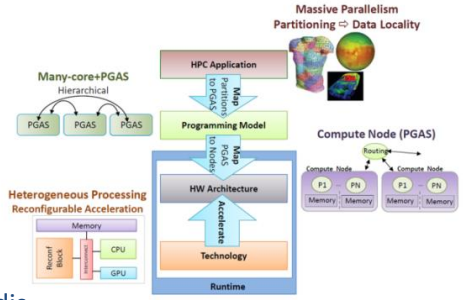

Source: Iakovos Mavroidis

European Exascale System Interconnect and Storage - www.exanest.eu

ExaNeSt

Storage, Interconnect, Cooling

Storage: fast, distributed in-node non-volatile memory
Interconnect: low-latency, unified compute & storage traffic
Packaging: advanced, liquid-cooled
App's: real, scientific and datacenter
Prototype: 1000+ ARM cores
from *EuroServer*: ARM nodes with UNIMEM address space & shared I/O
from *ExaNoDe*: Chiplets, Si Interposer
with *ECOSCALE*: Heterog. ARM+FPGA's

Iceotope Ltd: *Fully Immersed Cooling Technology*

Source: Manolis Marazakis

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

Ecosystem from European Projects

ARMv8, UNIMEM for micro-servers

EURO SERVER

Architecture

Compute node

ExaNoDe

Access to early prototype,

Compute node (HW & SW)

Rack-level testbed

ECOSCALE

Acceleration with FPGA

ExaNeSt

Interconnect, storage & cooling

ExaNoDe as part of a global strategy

EURO SERVER

redesigns the entreprise server:

Lower cost through system integration

Energy efficiency : low-power 64 bit processor and more efficient software

Mutualization of I/O resources

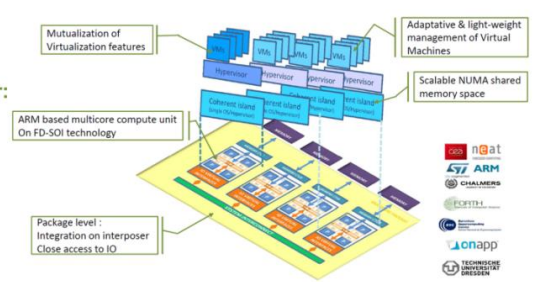

Source: Isabelle Dor

ECOSCALE

focuses on acceleration

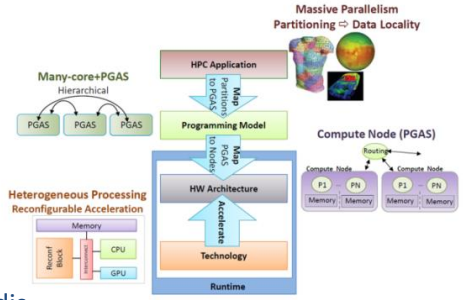

Source: Iakovos Mavroidis

European Exascale System Interconnect and Storage - www.exanest.eu

ExaNeSt

Storage, Interconnect, Cooling

Storage: fast, distributed in-node non-volatile memory
Interconnect: low-latency, unified compute & storage traffic
Packaging: advanced, liquid-cooled
App's: real, scientific and datacenter
Prototype: 1000+ ARM cores
from *EuroServer*: ARM nodes with UNIMEM address space & shared I/O
from *ExaNoDe*: Chiplets, Si Interposer
with *ECOSCALE*: Heterog. ARM+FPGA's

Iceotope Ltd: *Fully Immersed Cooling Technology*

Source: Manolis Marazakis

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

Ecosystem from European Projects

ARMv8, UNIMEM for micro-servers

EURO SERVER

Architecture

Compute node

ExaNoDe

FPGA toolchain

Compute node (HW & SW)

Access to early prototype,

Rack-level testbed

ECOSCALE

Acceleration with FPGA

ExaNeSt

Interconnect, storage & cooling

ExaNoDe as part of a global strategy

EURO SERVER

redesigns the entreprise server:

Lower cost through system integration

Energy efficiency : low-power 64 bit processor and more efficient software

Mutualization of I/O resources

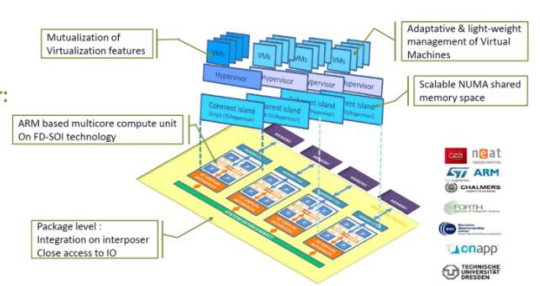

Source: Isabelle Dor

ECOSCALE

focuses on acceleration

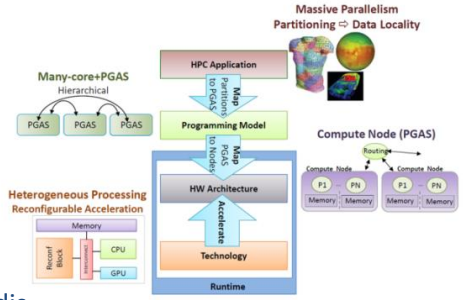

Source: Iakovos Mavroidis

European Exascale System Interconnect and Storage - www.exanest.eu

ExaNeSt

Storage, Interconnect, Cooling

Storage: fast, distributed in-node non-volatile memory
Interconnect: low-latency, unified compute & storage traffic
Packaging: advanced, liquid-cooled
App's: real, scientific and datacenter
Prototype: 1000+ ARM cores
from *EuroServer*: ARM nodes with UNIMEM address space & shared I/O
from *ExaNoDe*: Chiplets, Si Interposer
with *ECOSCALE*: Heterog. ARM+FPGA's

Iceotope Ltd: *Fully Immersed Cooling Technology*

Source: Manolis Marazakis

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

Ecosystem from European Projects

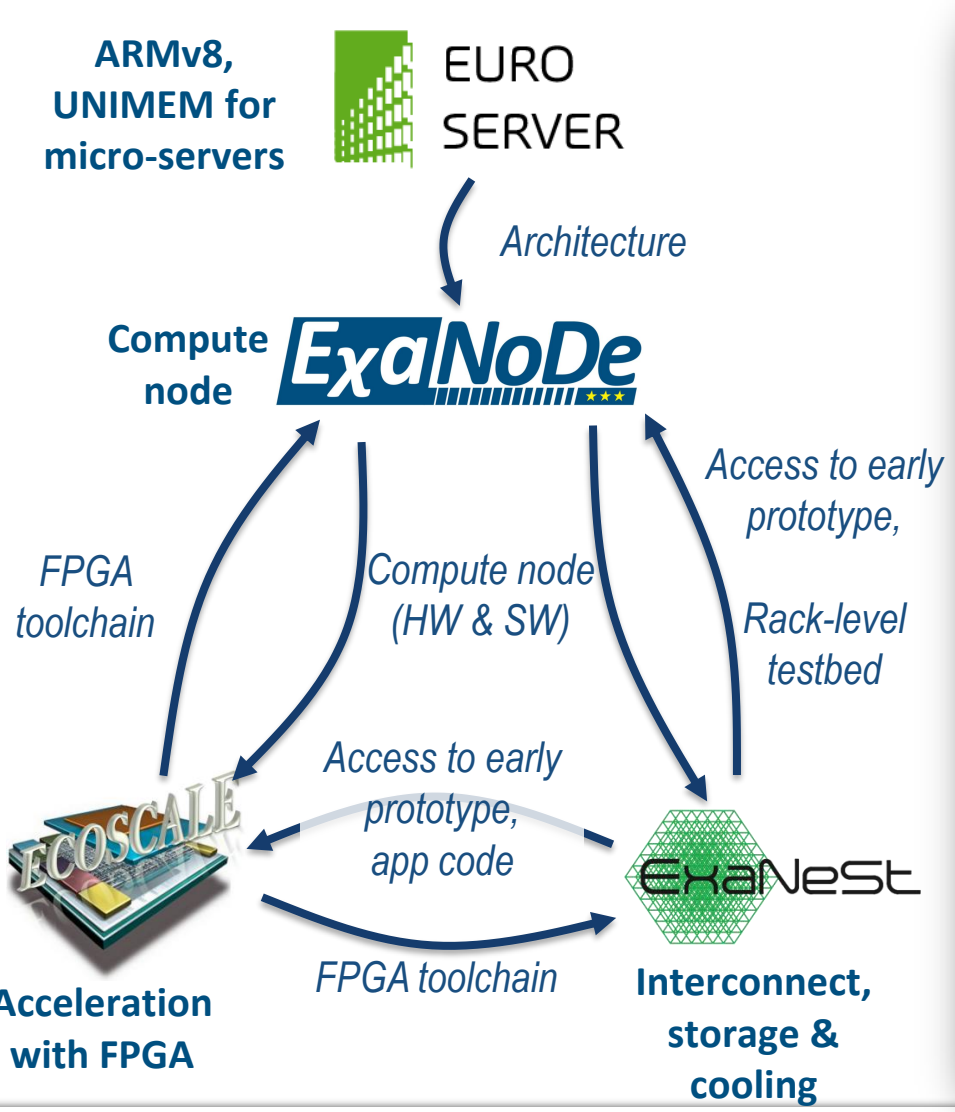

ExaNoDe as part of a global strategy

EURO SERVER

redesigns the entreprise server:

Lower cost through system integration

Energy efficiency : low-power 64 bit processor and more efficient software

Mutualization of I/O resources

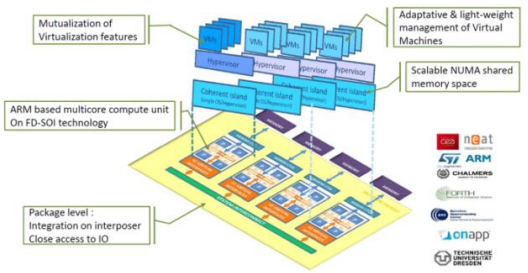

Source: Isabelle Dor

focuses on acceleration

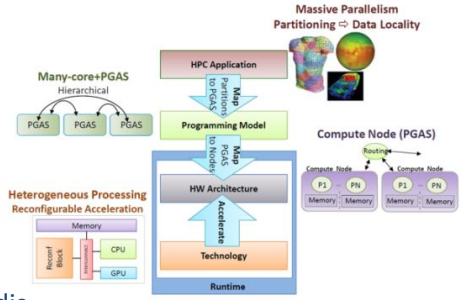

Source: Iakovos Mavroidis

European Exascale System Interconnect and Storage - www.exanest.eu

Storage, Interconnect, Cooling

Storage: fast, distributed in-node non-volatile memory
Interconnect: low-latency, unified compute & storage traffic
Packaging: advanced, liquid-cooled
App's: real, scientific and datacenter
Prototype: 1000+ ARM cores
from *EuroServer*: ARM nodes with UNIMEM address space & shared I/O
from *ExaNoDe*: Chiplets, Si Interposer
with *ECOSCALE*: Heterog. ARM+FPGA's

Iceotope Ltd: *Fully Immersed Cooling Technology*

Source: Manolis Marazakis

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

ExaNoDe Project Implementation

Start: ***October 1st, 2015***

Duration: ***3 years***

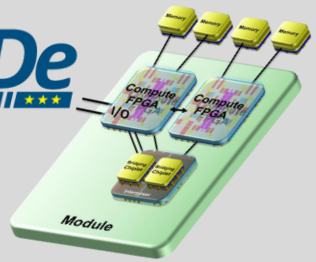

8.6 M€

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

ExaNoDe Project Implementation

Start: ***October 1st, 2015***
Duration: ***3 years***

Coordinator

CONSORTIUM

ExaNoDe

8.6 M€

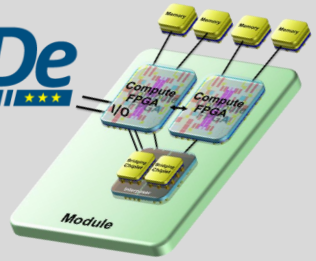

HORIZON 2020

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

ExaNoDe Project Implementation

Start: ***October 1st, 2015***
Duration: ***3 years***

Coordinator

CONSORTIUM

IMPLEMENTATION

ExaNoDe

8.6 M€

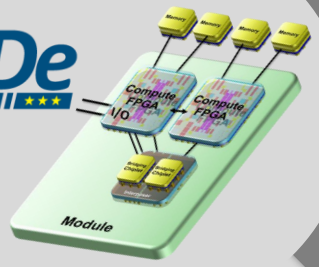

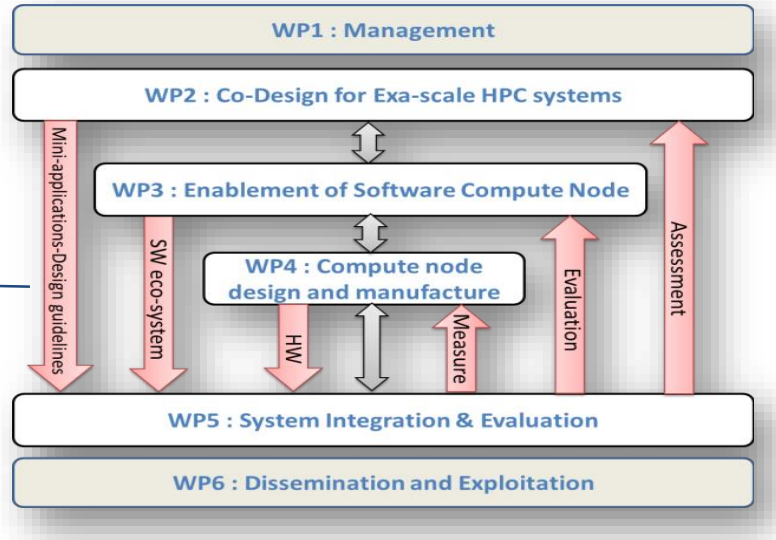

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

ExaNoDe Project Implementation

Start: ***October 1st, 2015***
Duration: ***3 years***

Coordinator

CONSORTIUM

IMPLEMENTATION

PROTOTYPE

ExaNoDe

8.6 M€

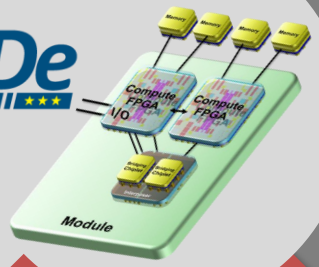

HORIZON 2020

- WP1 : Management
- WP2 : Co-Design for Exa-scale HPC systems
- WP3 : Enablement of Software Compute Node
- WP4 : Compute node design and manufacture
- WP5 : System Integration & Evaluation
- WP6 : Dissemination and Exploitation

Mini-applications-Design guidelines · SW eco-system · HW · Measure · Evaluation · Assessment

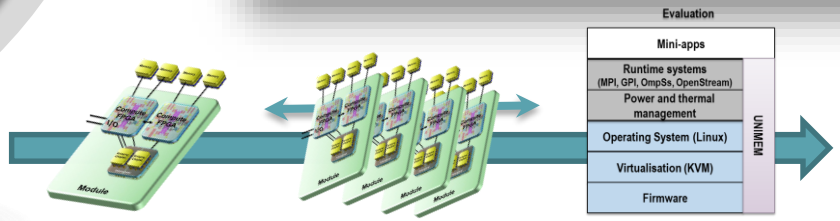

Compute Node · HW Prototype · SW Stack

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

ExaNoDe at a glance

Key technologies for compute nodes
towards a future **Exascale capability**

Key technologies

Exascale requirements

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

ExaNoDe at a glance

Key technologies for compute nodes
towards a future **Exascale capability**

	System Architecture
Key technologies	• ARMv8 • Coherent islands • Global Address Space
Exascale requirements	👍 Energy efficiency 👍 Scalability

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

ExaNoDe at a glance

Key technologies for compute nodes towards a future **Exascale capability**

	System Architecture	Silicon Integration

Key technologies	• ARMv8
	• Coherent islands
	• Global Address Space
		• 3D Integration:
		  • Chiplet
		  • Active Interposer
		• Multi-Chip-Module:
		  • FPGA, Memory

Exascale requirements	👍 Energy efficiency 👍 Scalability	👍 Design/manufacturing costs 👍 Heterogeneity/Specialization

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

ExaNoDe at a glance

Key technologies for compute nodes
towards a future **Exascale capability**

	System Architecture	Silicon Integration	Software

Key technologies	• ARMv8
	• Coherent islands
	• Global Address Space
		• 3D Integration:
		  • Chiplet
		  • Active Interposer
		• Multi-Chip-Module:
		  • FPGA, Memory
			• FW, OS
			• Virtualization
			• Programming models
			• Runtimes
			• Mini-apps

Exascale requirements	👍 Energy efficiency 👍 Scalability	👍 Design/manufacturing costs 👍 Heterogeneity/Specialization	👍 Co-design 👍 Scalability

All combined in an integrated prototype

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

2. System Architecture and Integration

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

ExaNoDe Architectural Approach

- **Compute unit (coherent island)**
  - Several ARMv8 cores
  - Local coherent interconnect
  - Path to local memory
  - Path “to remote” memory locations
  - Path “from remote” compute units

Source: John Goodacre – DATE’13

EURO SERVER

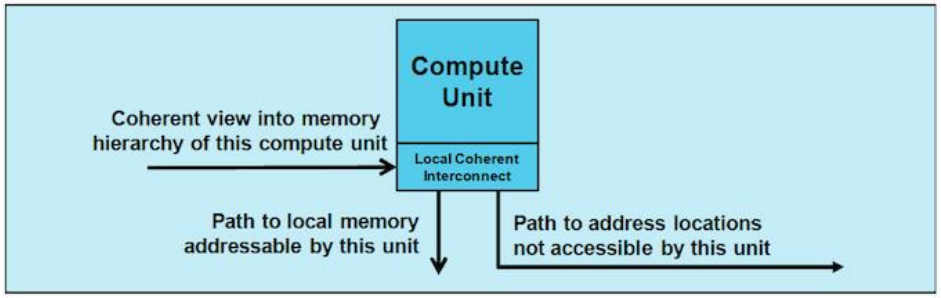

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

ExaNoDe Architectural Approach

- **Compute unit (coherent island)**
  - Several ARMv8 cores
  - Local coherent interconnect
  - Path to local memory
  - Path “to remote” memory locations
  - Path “from remote” compute units

Source: John Goodacre – DATE’13

EURO SERVER

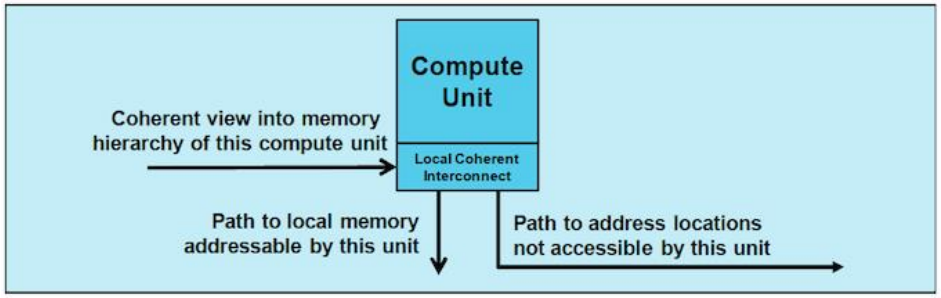

- **Compute Node**
  - Scale-out …

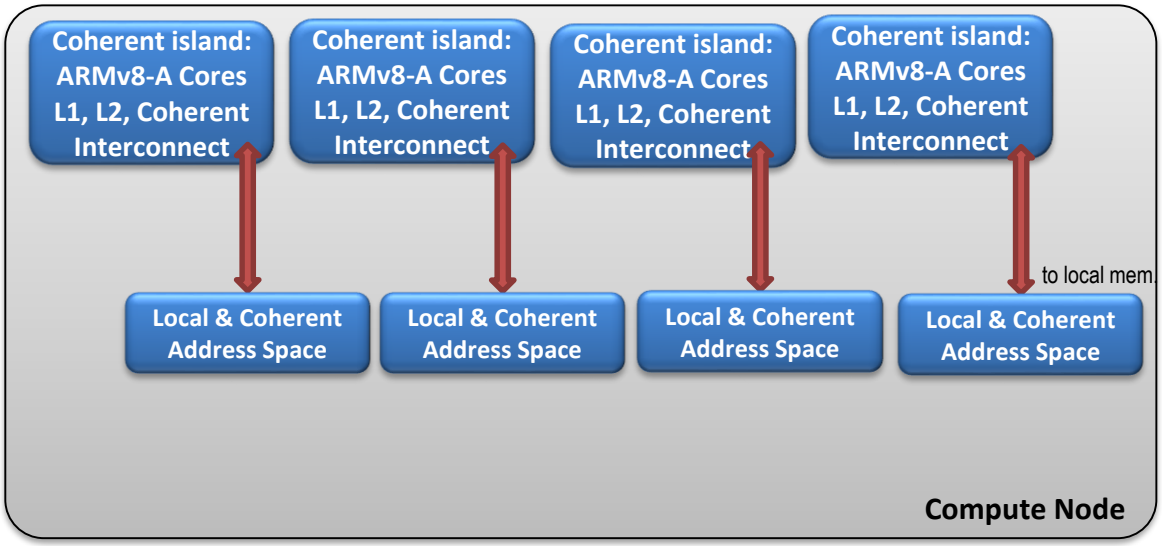

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

ExaNoDe Architectural Approach

- **Compute unit (coherent island)**
  - Several ARMv8 cores
  - Local coherent interconnect
  - Path to local memory
  - Path “to remote” memory locations
  - Path “from remote” compute units

Source: John Goodacre – DATE’13

EURO SERVER

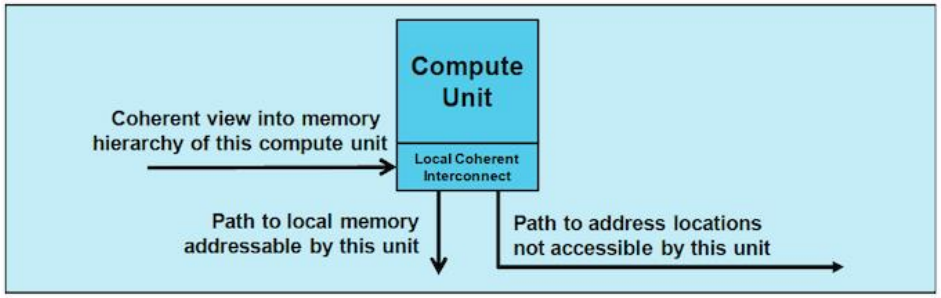

- **Compute Node**
  - Scale-out …
  - … with “from” and “to” remote port connections to the global network

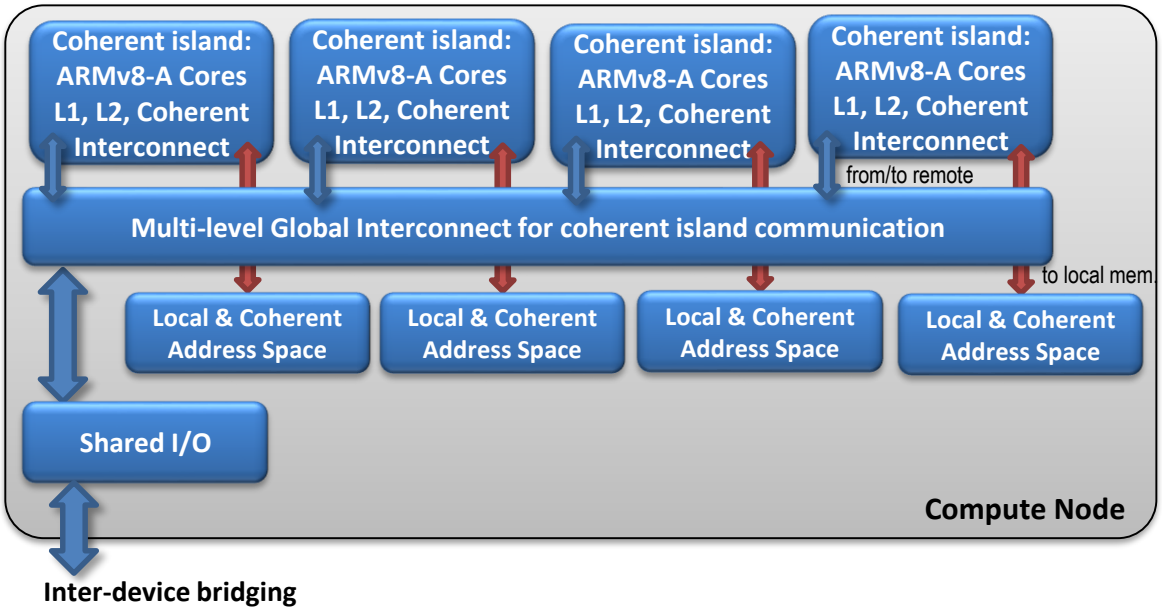

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

ExaNoDe Architectural Approach

- **Compute unit (coherent island)**
  - Several ARMv8 cores
  - Local coherent interconnect
  - Path to local memory
  - Path “to remote” memory locations
  - Path “from remote” compute units

Source: John Goodacre – DATE’13

EURO SERVER

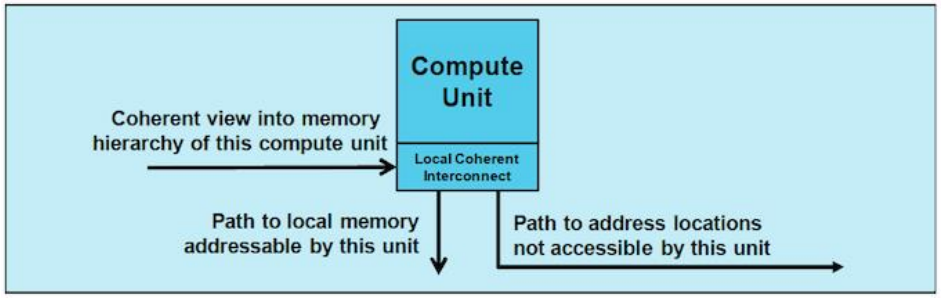

- **Compute Node**
  - Scale-out …
  - … with “from” and “to” remote port connections to the global network

- **Memory scheme**
  - Global Address Space and remote coherent access with UNIMEM

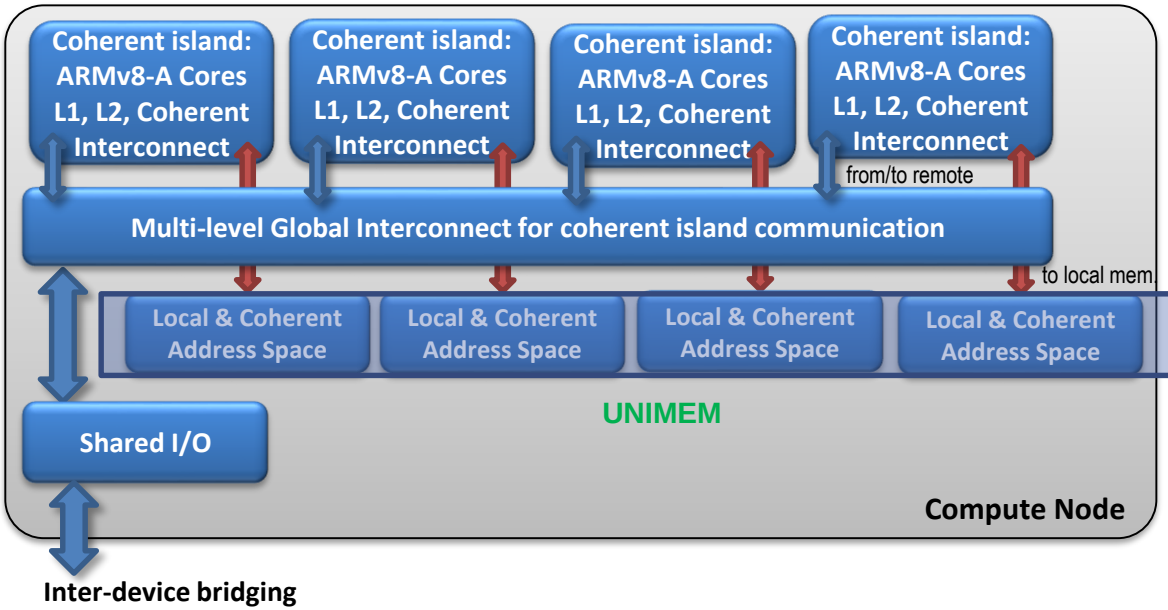

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

UNIMEM: Partitioned Global Address Space

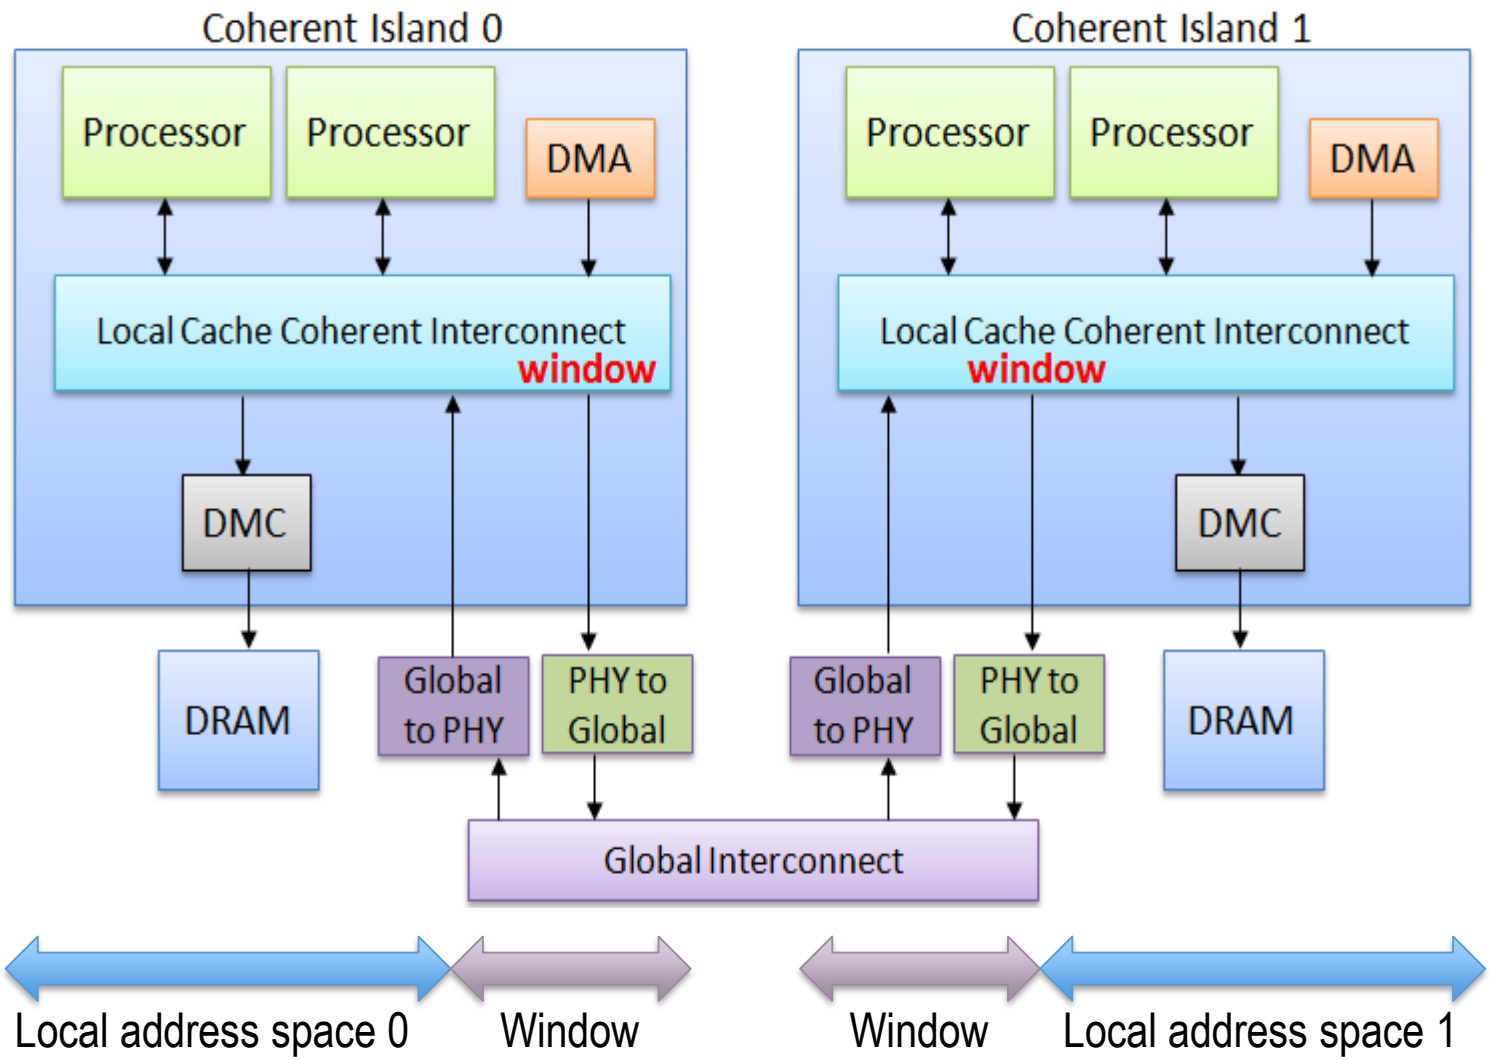

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

Combined Architecture/Integration

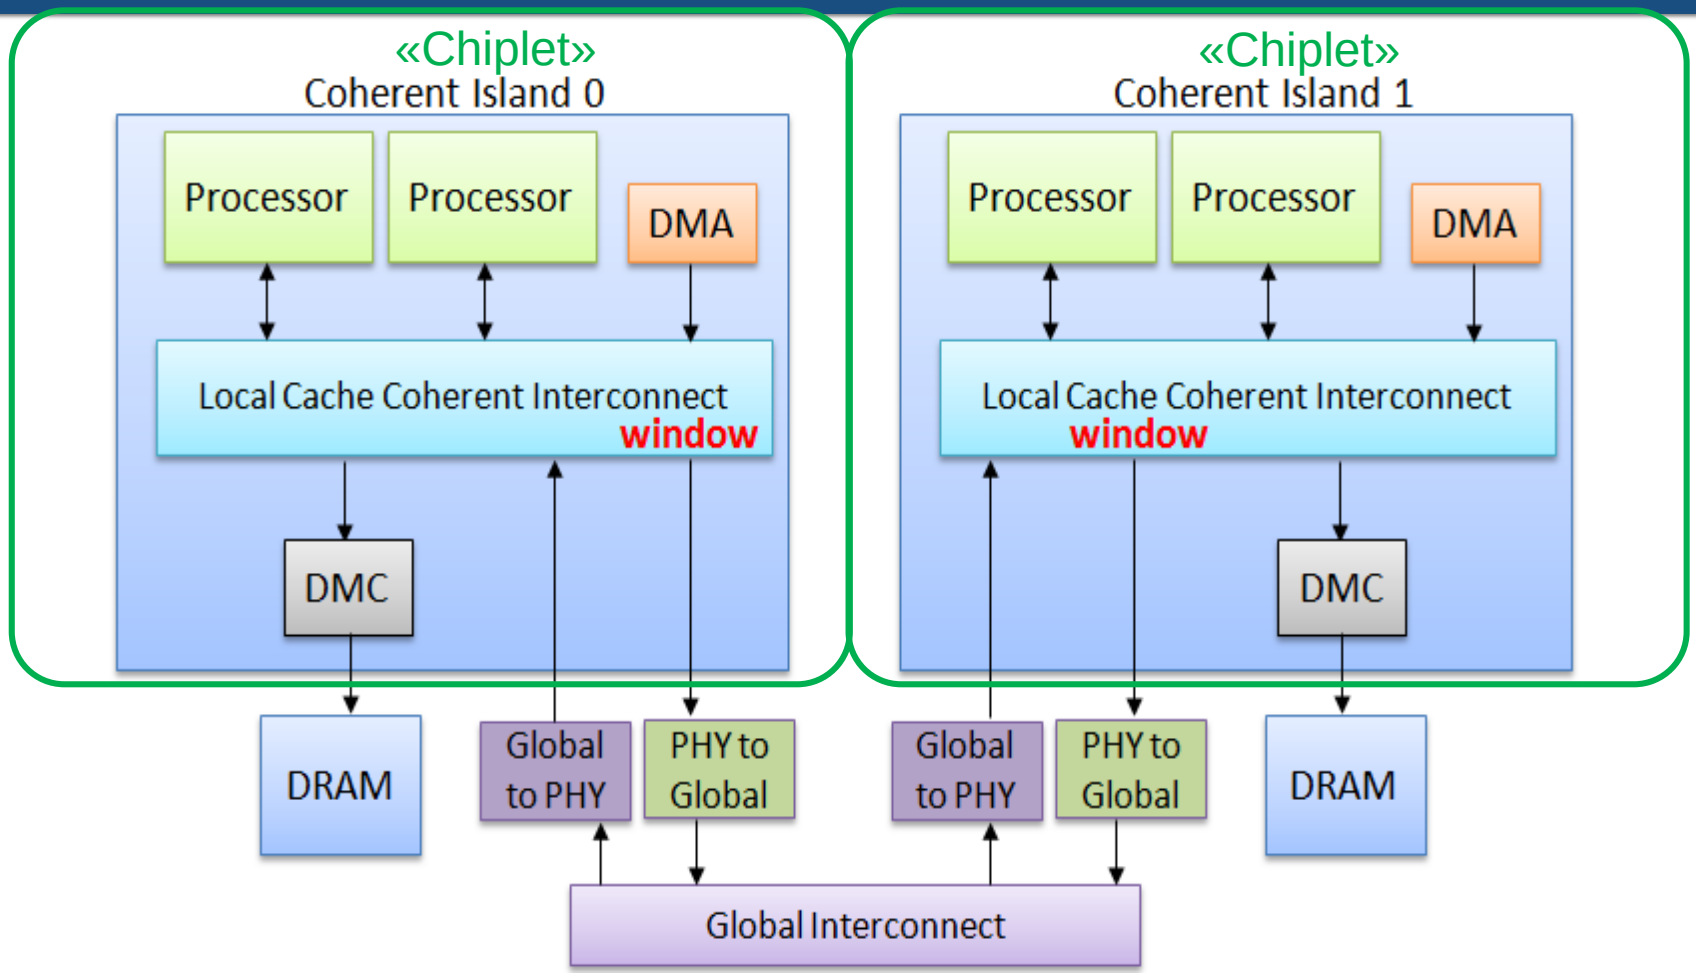

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

Combined Architecture/Integration

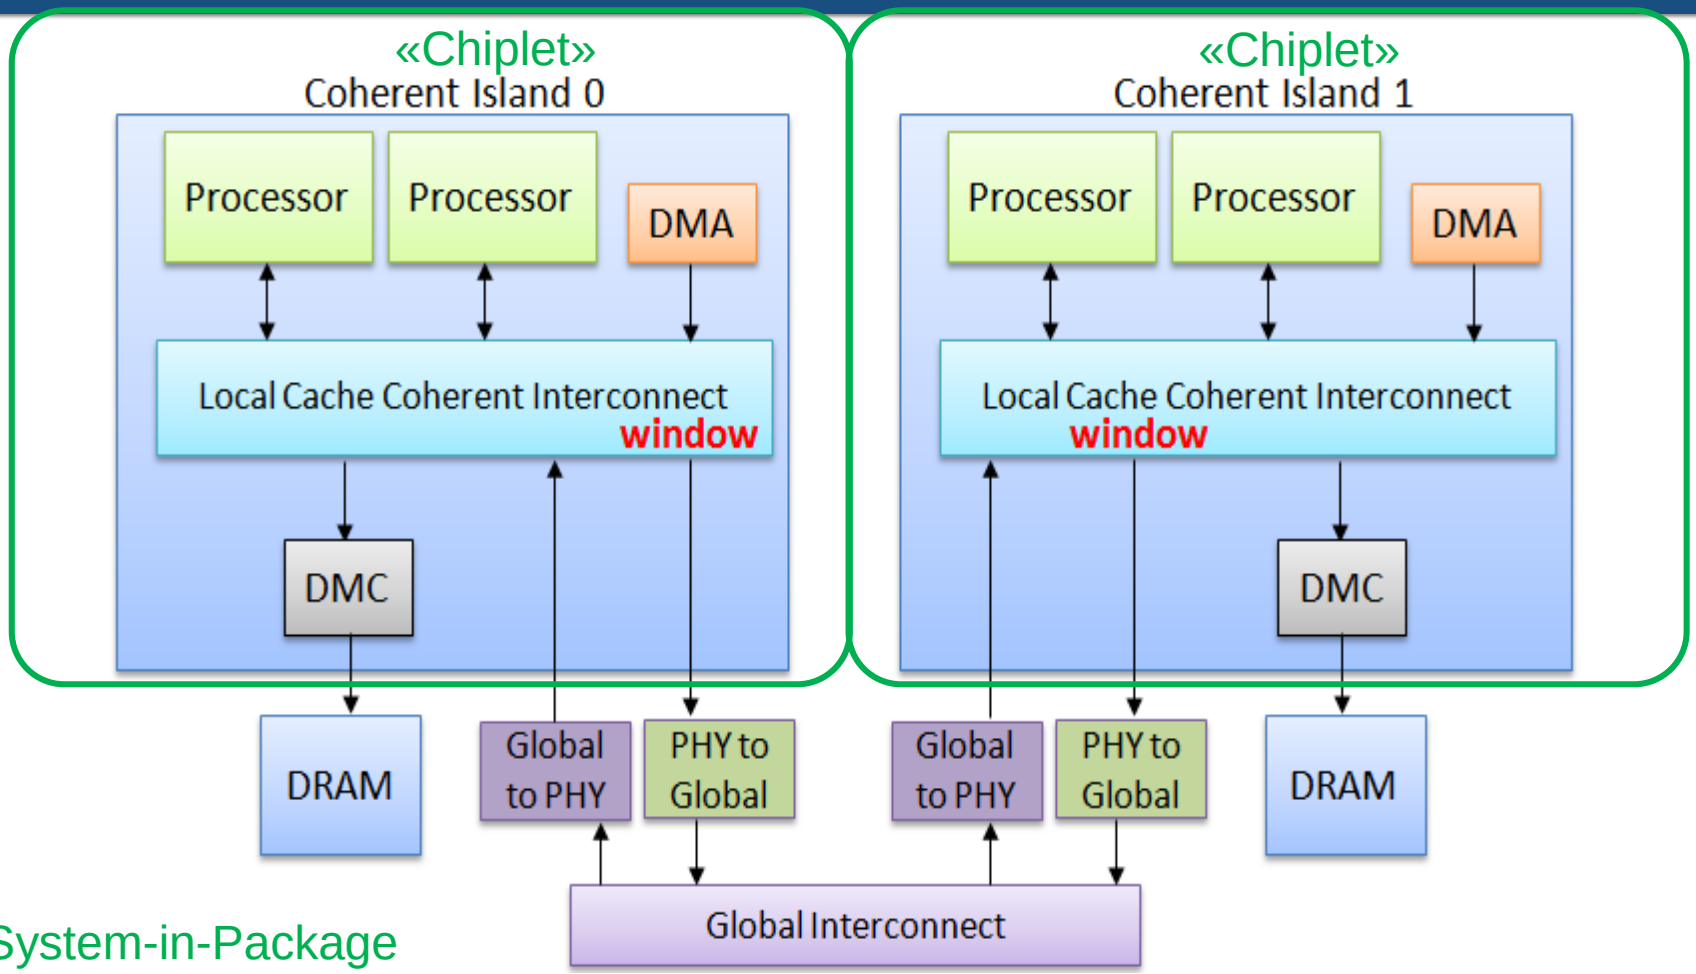

System-in-Package

EURO SERVER

ExaNoDe

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

Combined Architecture/Integration

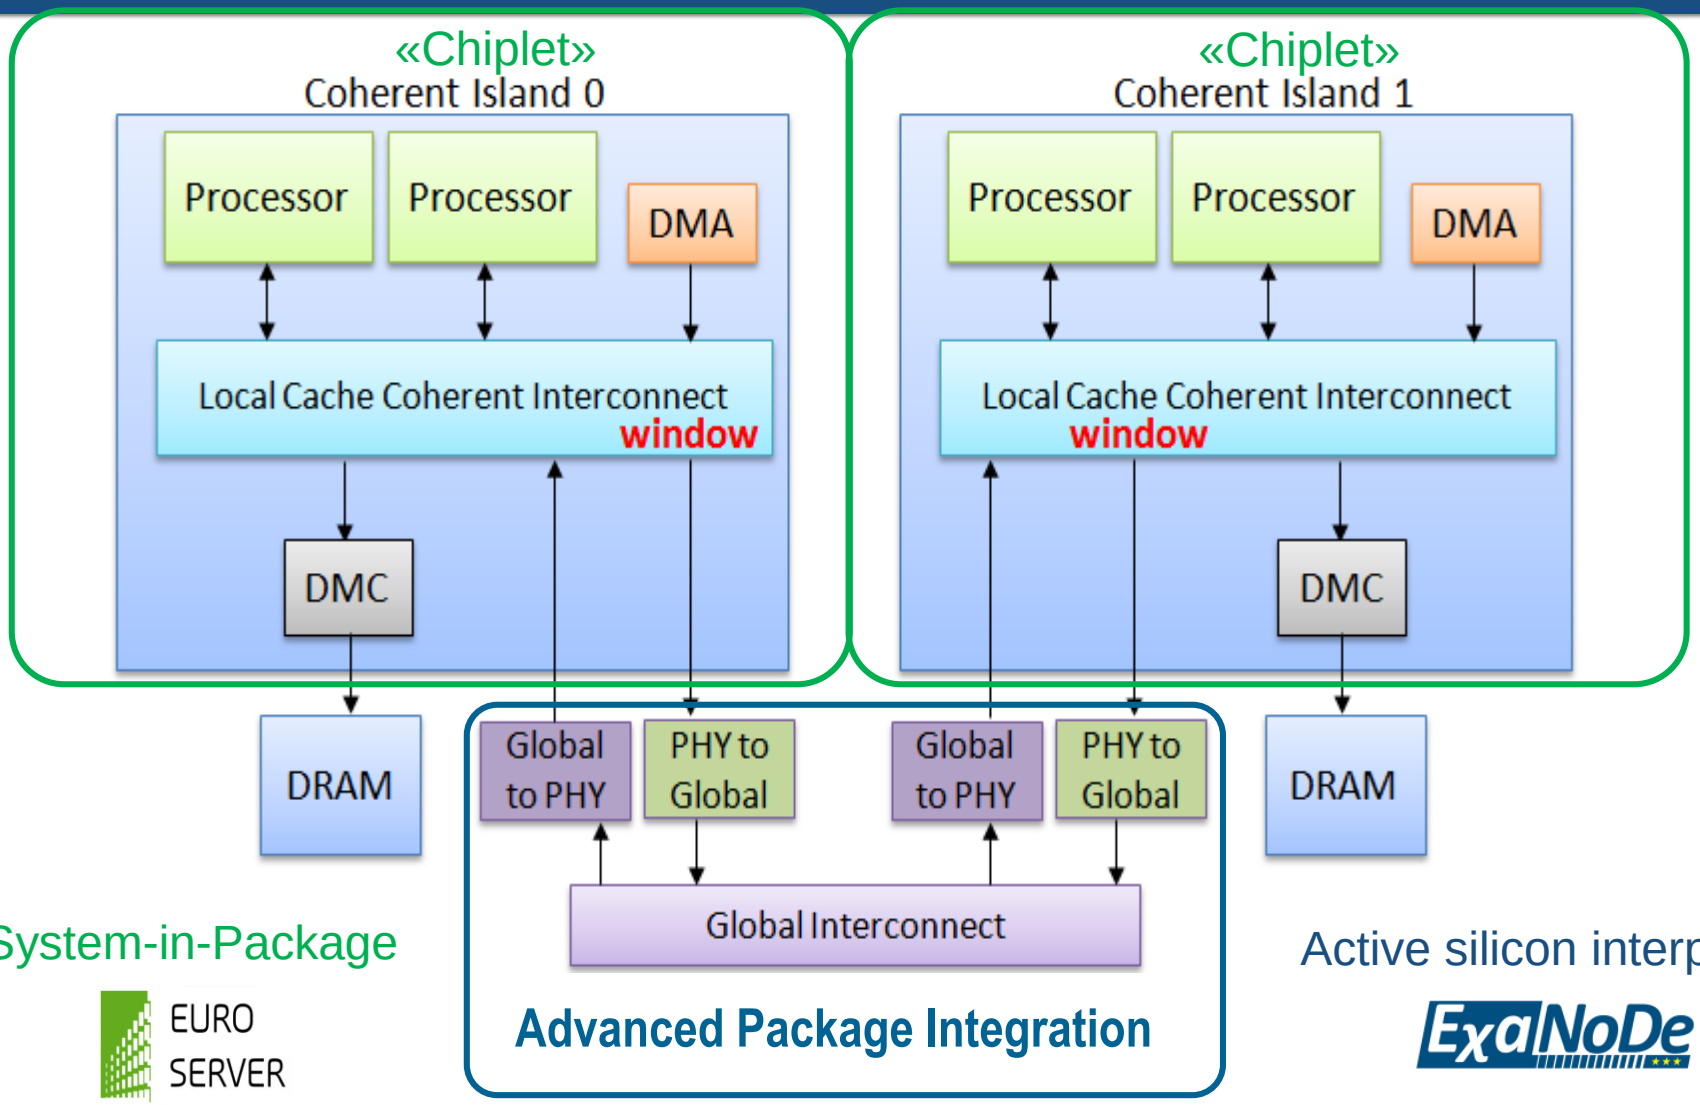

System-in-Package

EURO SERVER

Active silicon interposer

ExaNoDe

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

ExaNoDe Integration Approach

- **3D integration with active silicon interposer and chiplets**

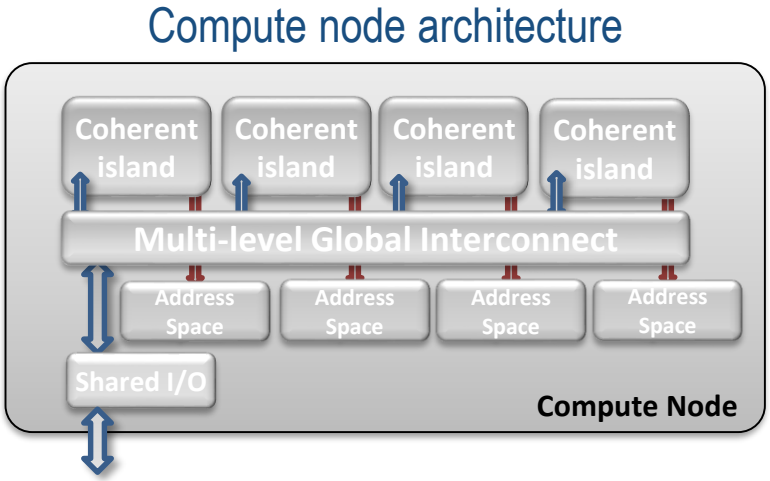

Copyright © 2017 Members of the EUROSERVER Consortium

contact: denis.dutoit@cea.fr, project coordinator

ExaNoDe Integration Approach

- 3D integration with active silicon interposer and chiplets

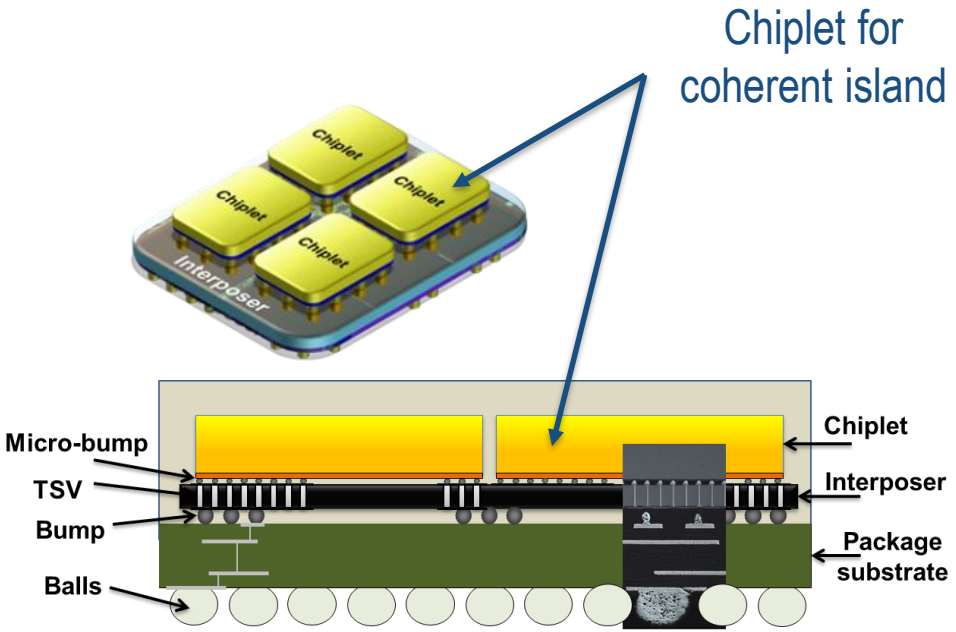

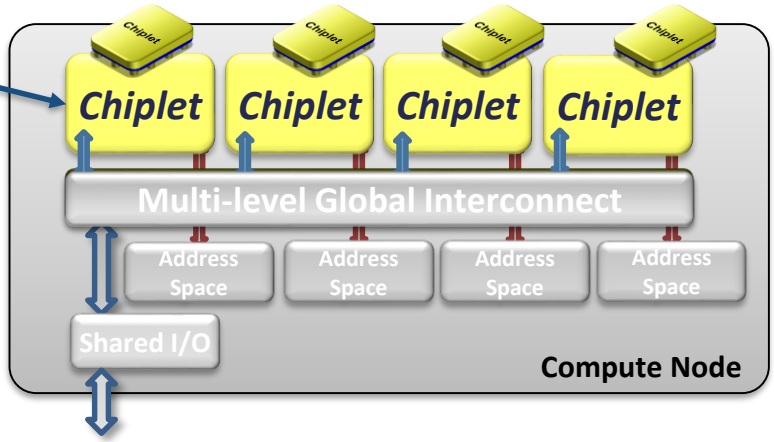

Copyright © 2017 Members of the EUROSERVER Consortium

contact: denis.dutoit@cea.fr, project coordinator

ExaNoDe Integration Approach

- 3D integration with active silicon interposer and chiplets

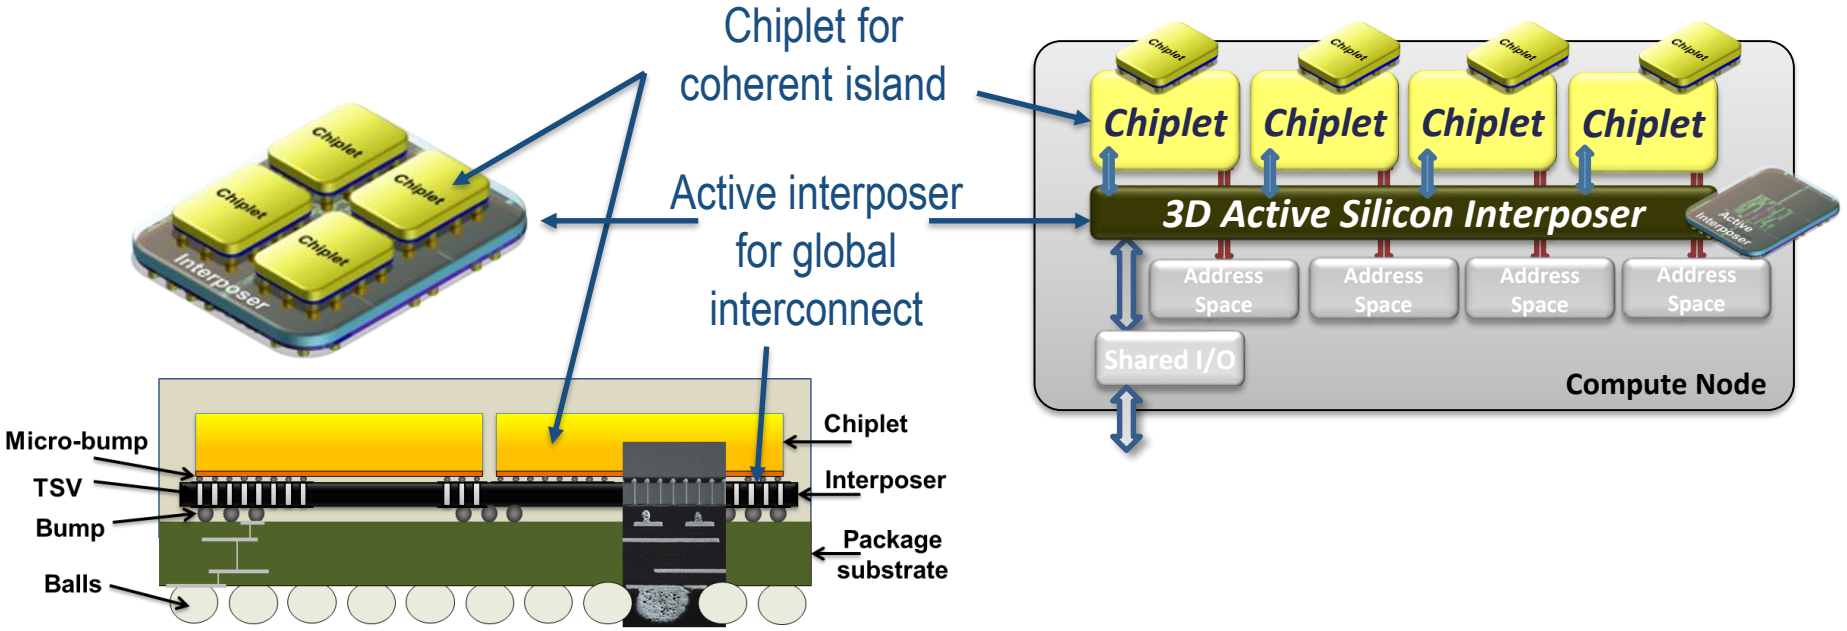

Copyright © 2017 Members of the EUROSERVER Consortium

ExaNoDe Integration Approach

- 3D integration with active silicon interposer and chiplets

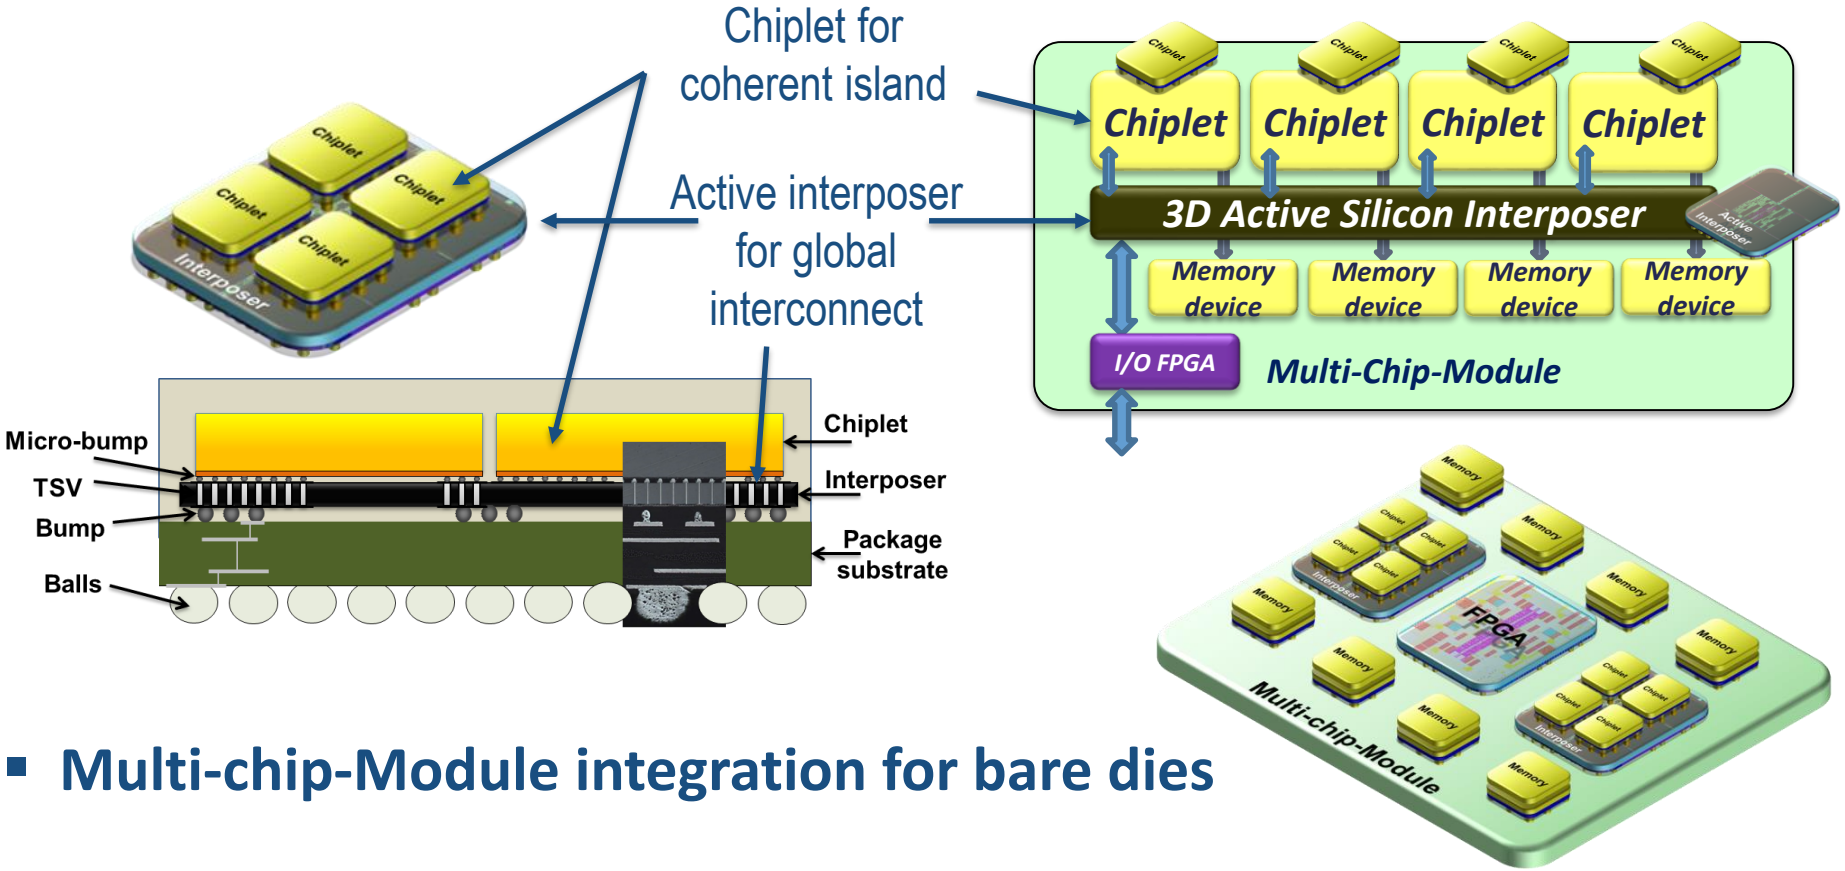

- Multi-chip-Module integration for bare dies

Copyright © 2017 Members of the EUROSERVER Consortium

EUROSERVER: Package

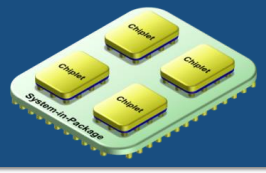

- **Multi-Core System with four chiplets in a 40x40 package**

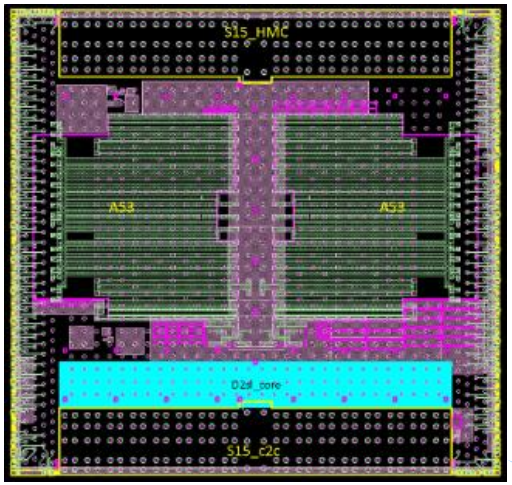

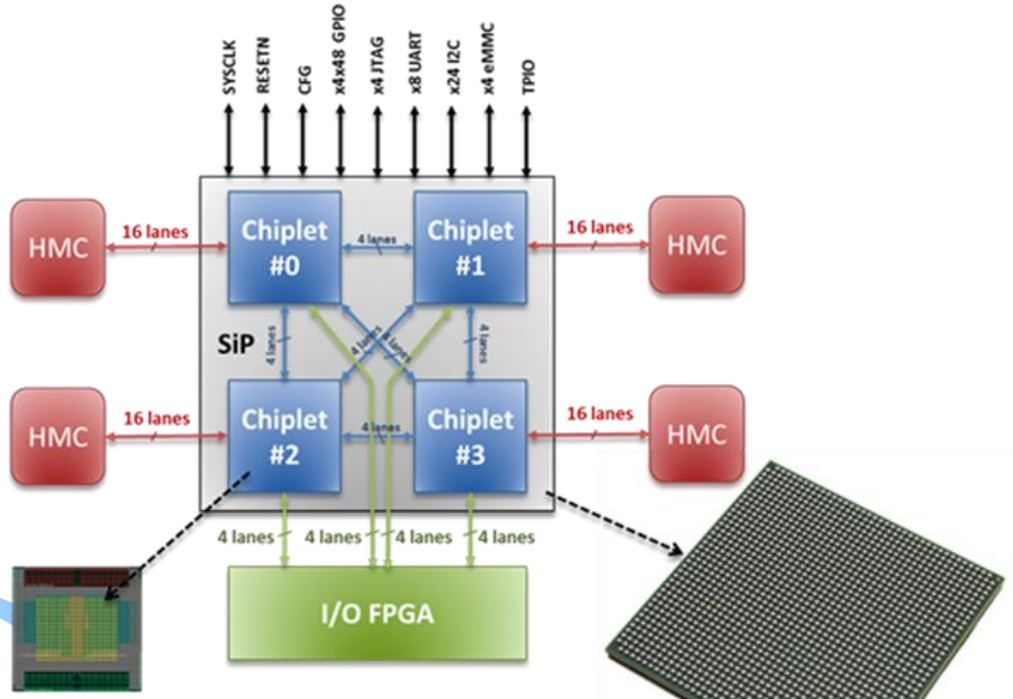

- A package contains 4 chiplets
- Each chiplet contains 2 quad-core ARMv8 A53
- System of 32 A53 cores in a Package

Copyright © 2017 Members of the EUROSERVER Consortium

3. ExaNoDe Prototype

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

3. ExaNoDe Prototype

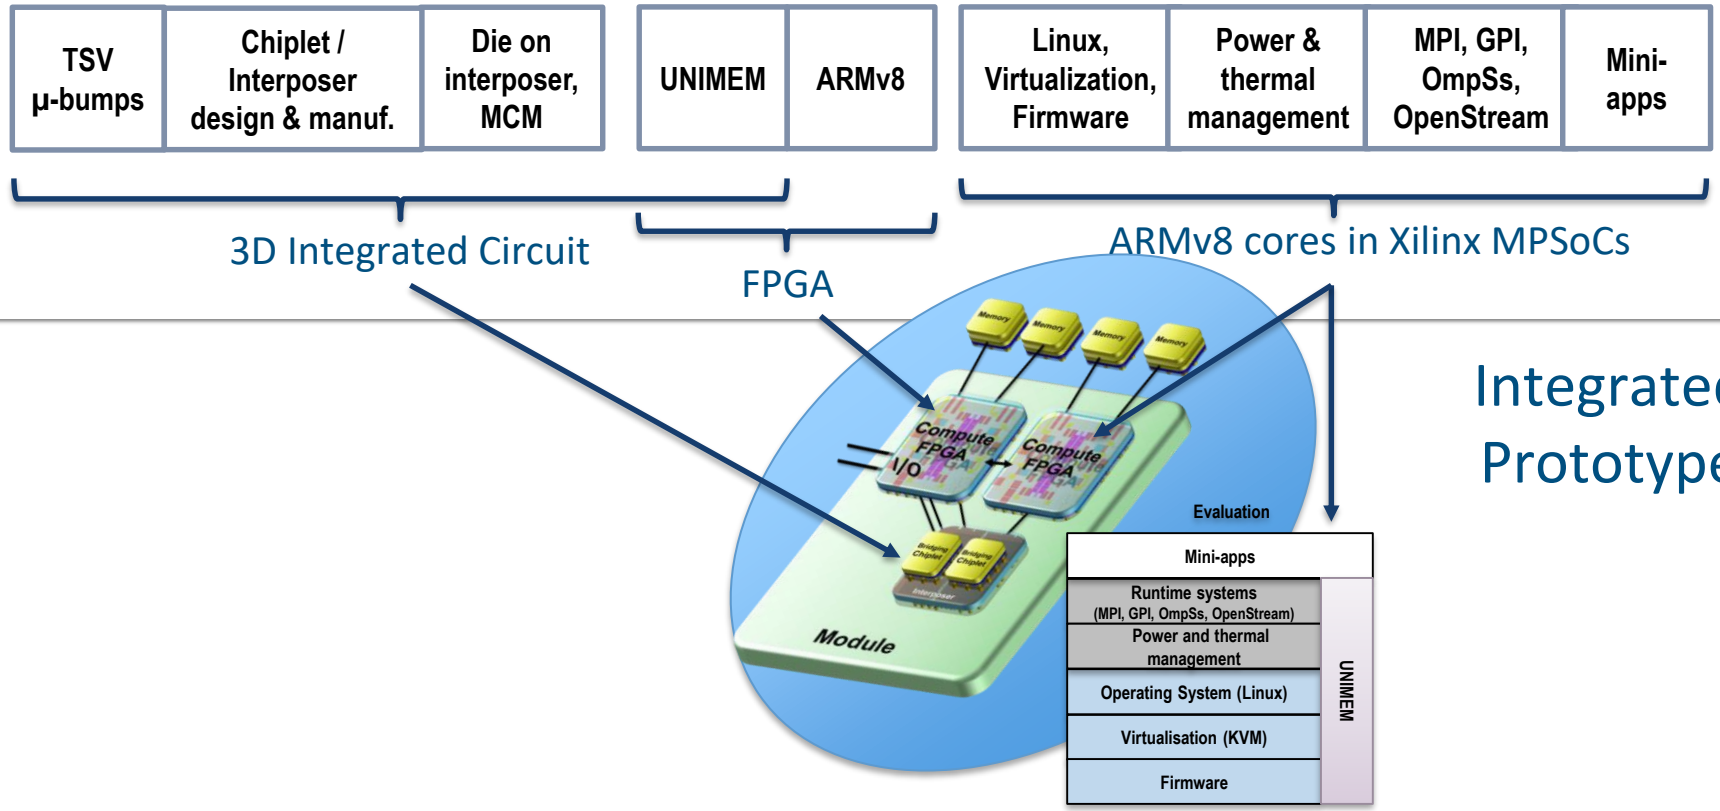

Integrated Prototype

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

ExaNoDe Hardware Prototype

TSV µ-bumps | Chiplet / Interposer design & manuf. | Die on interposer, MCM | UNIMEM | ARMv8 | Linux, Virtualization, Firmware | Power & thermal management | MPI, GPI, OmpSs, OpenStream | Mini-apps

3D Integrated Circuit

FPGA

ARMv8 cores in Xilinx MPSoCs

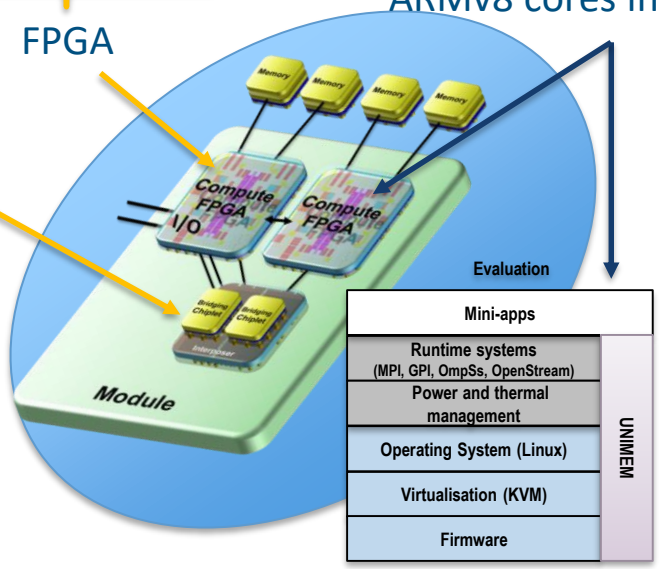

Integrated Prototype

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

ExaNoDe Objective: Deliver a Multi Chip Module

MCM Step 1: Q4-2017
- Without 3D-IC
- UNIMEM FPGA

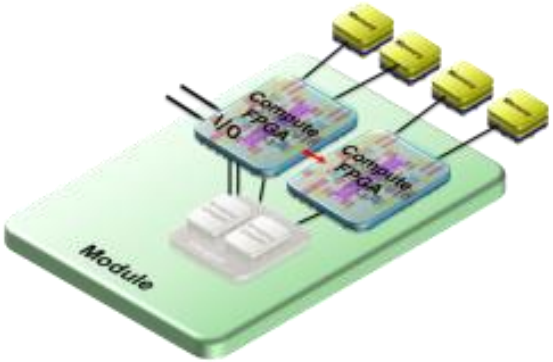

MCM Step 2: Q3-2018
- With 3D-IC
- UNIMEM-capable 3D-IC

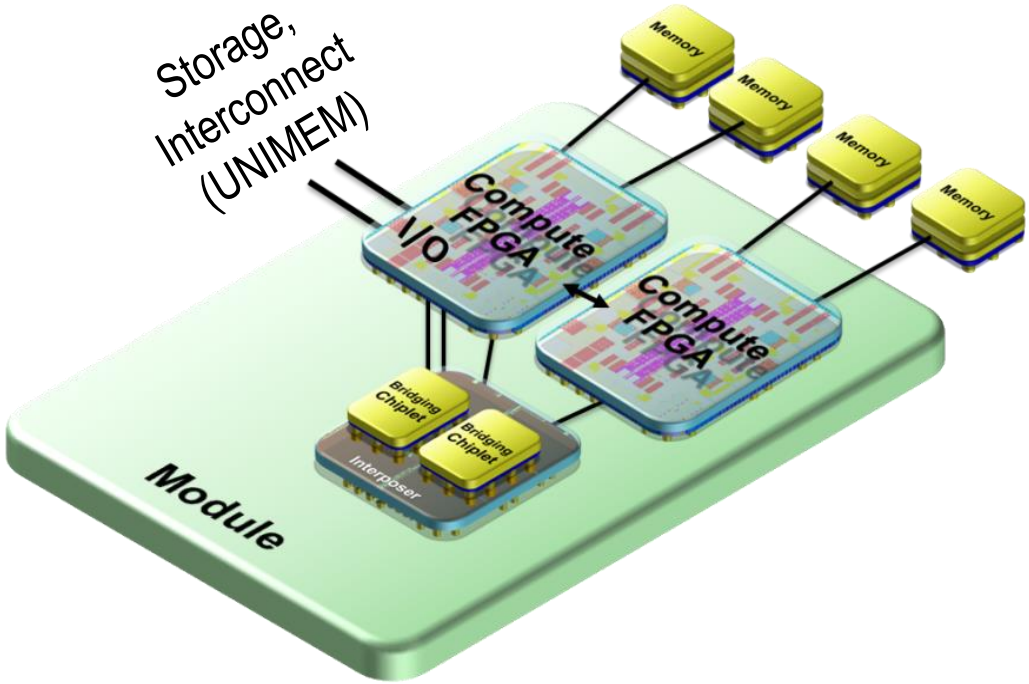

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

Early prototyping platforms

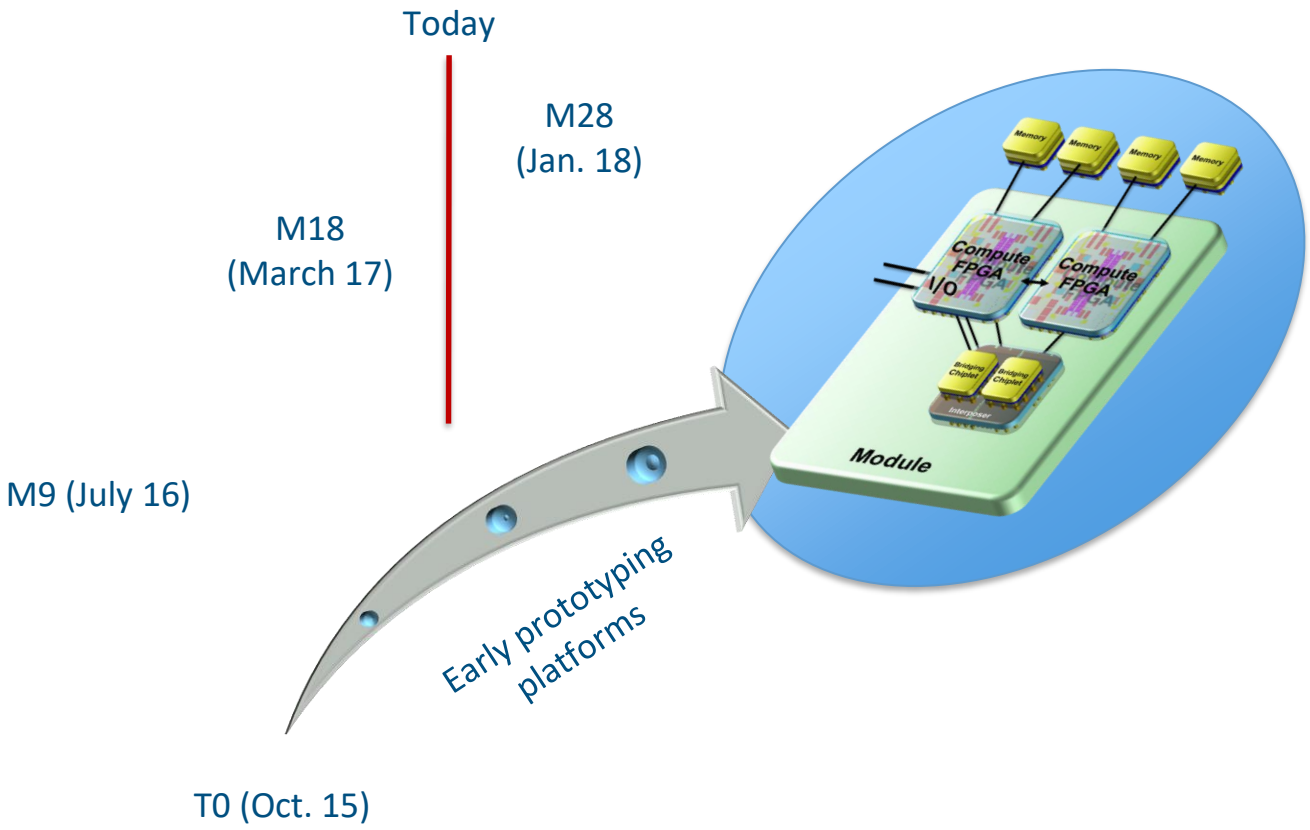

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

Early prototyping platforms

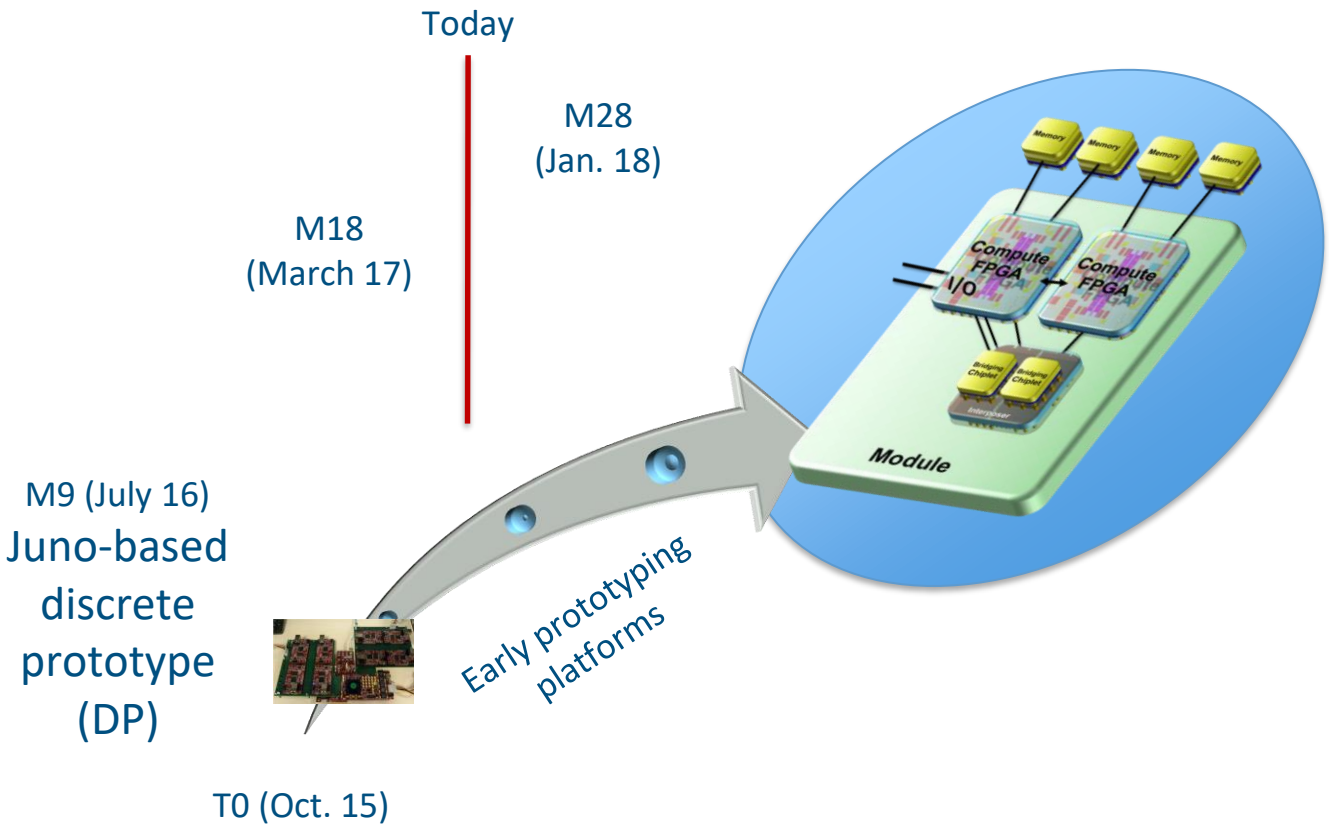

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

High-level Architecture of 64-bit Discrete Prototype (DP)

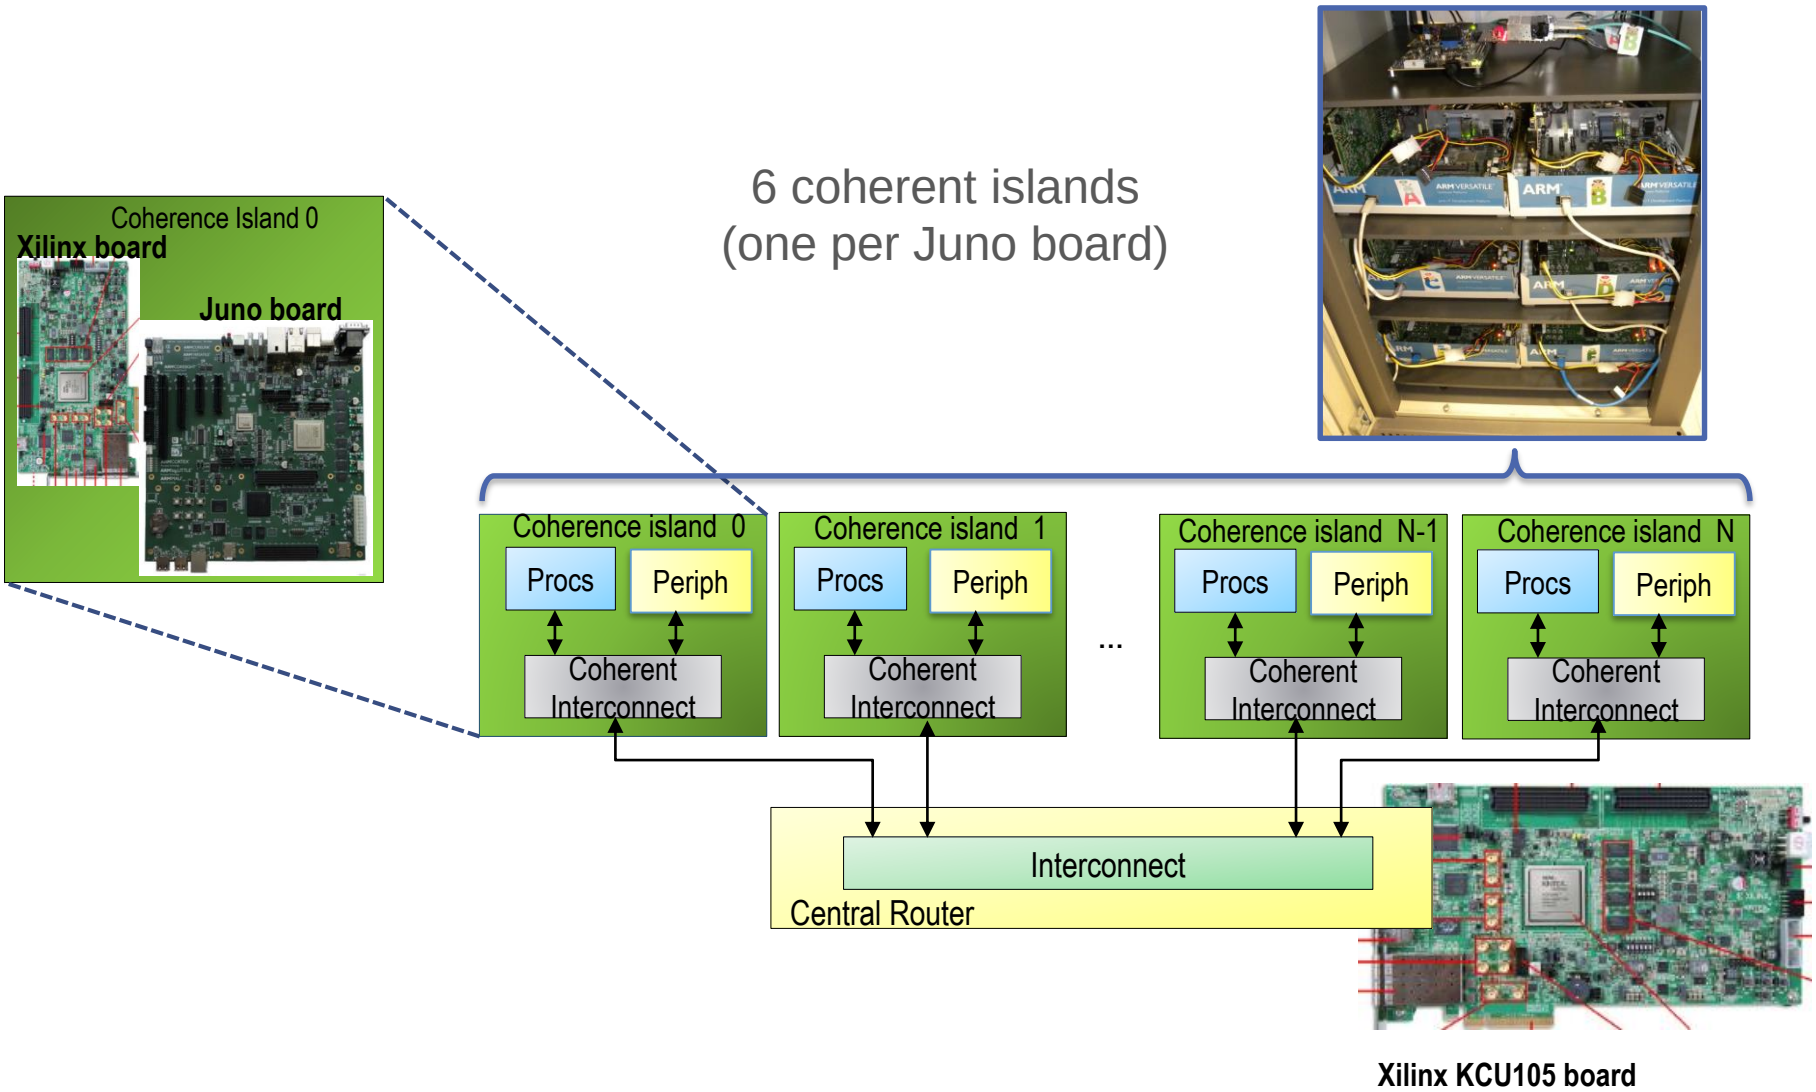

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

Early prototyping platforms

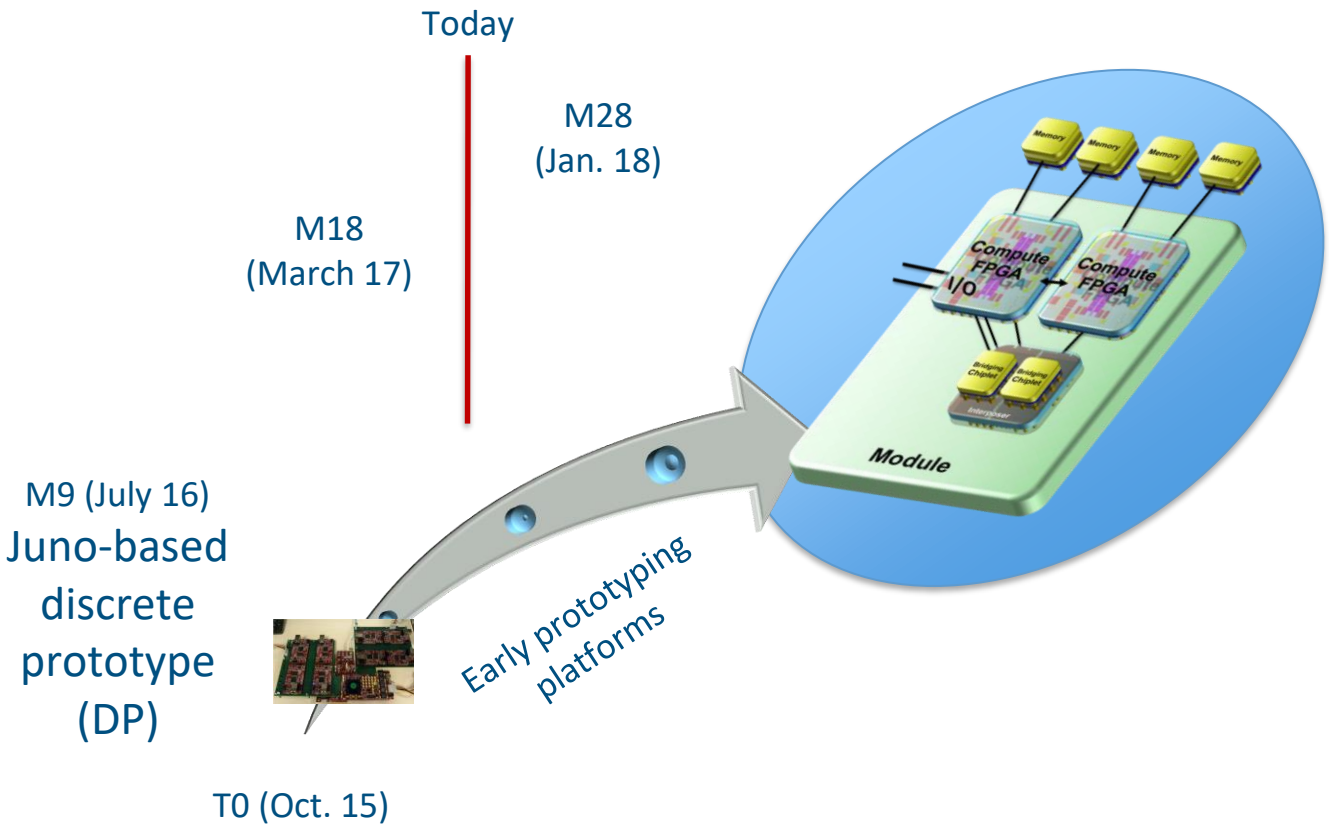

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

Early prototyping platforms

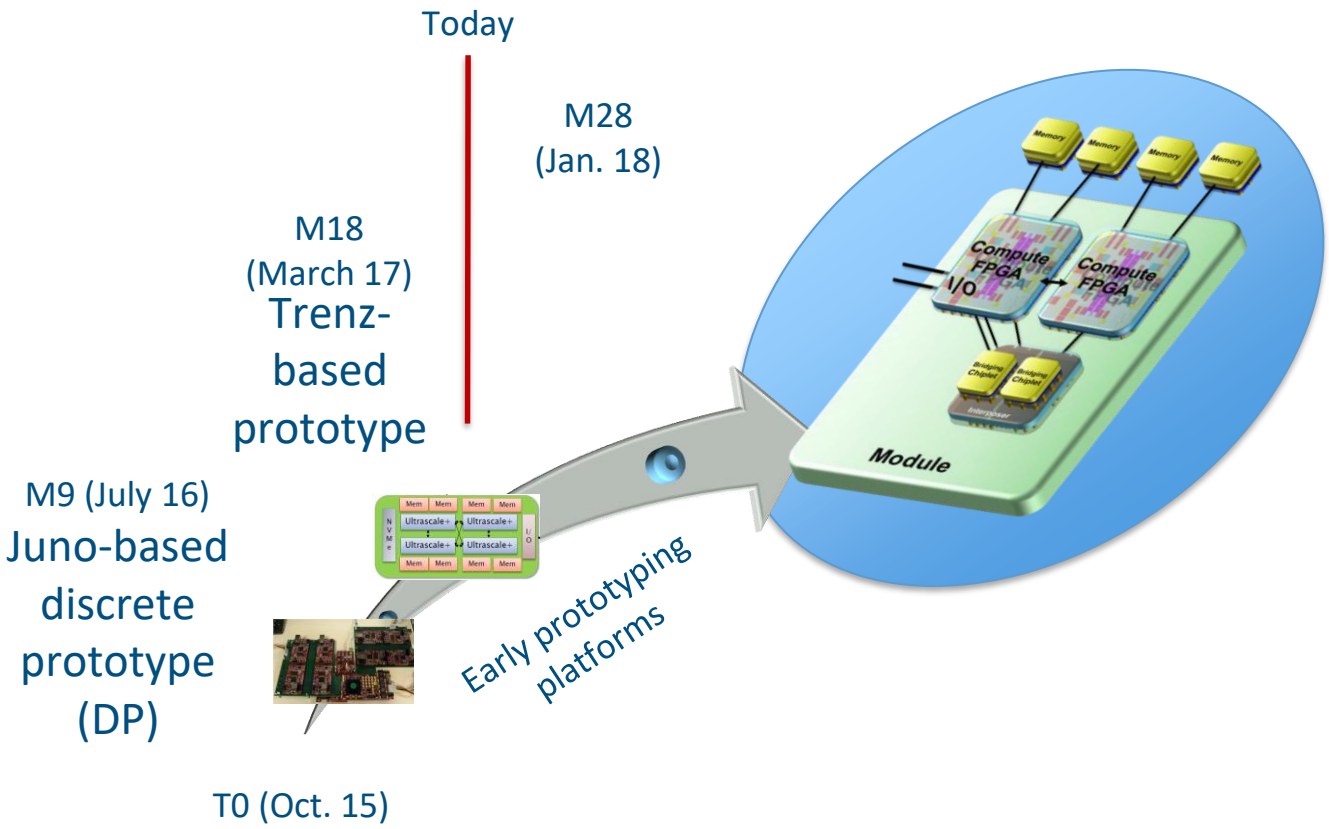

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

Trenz-based Prototype

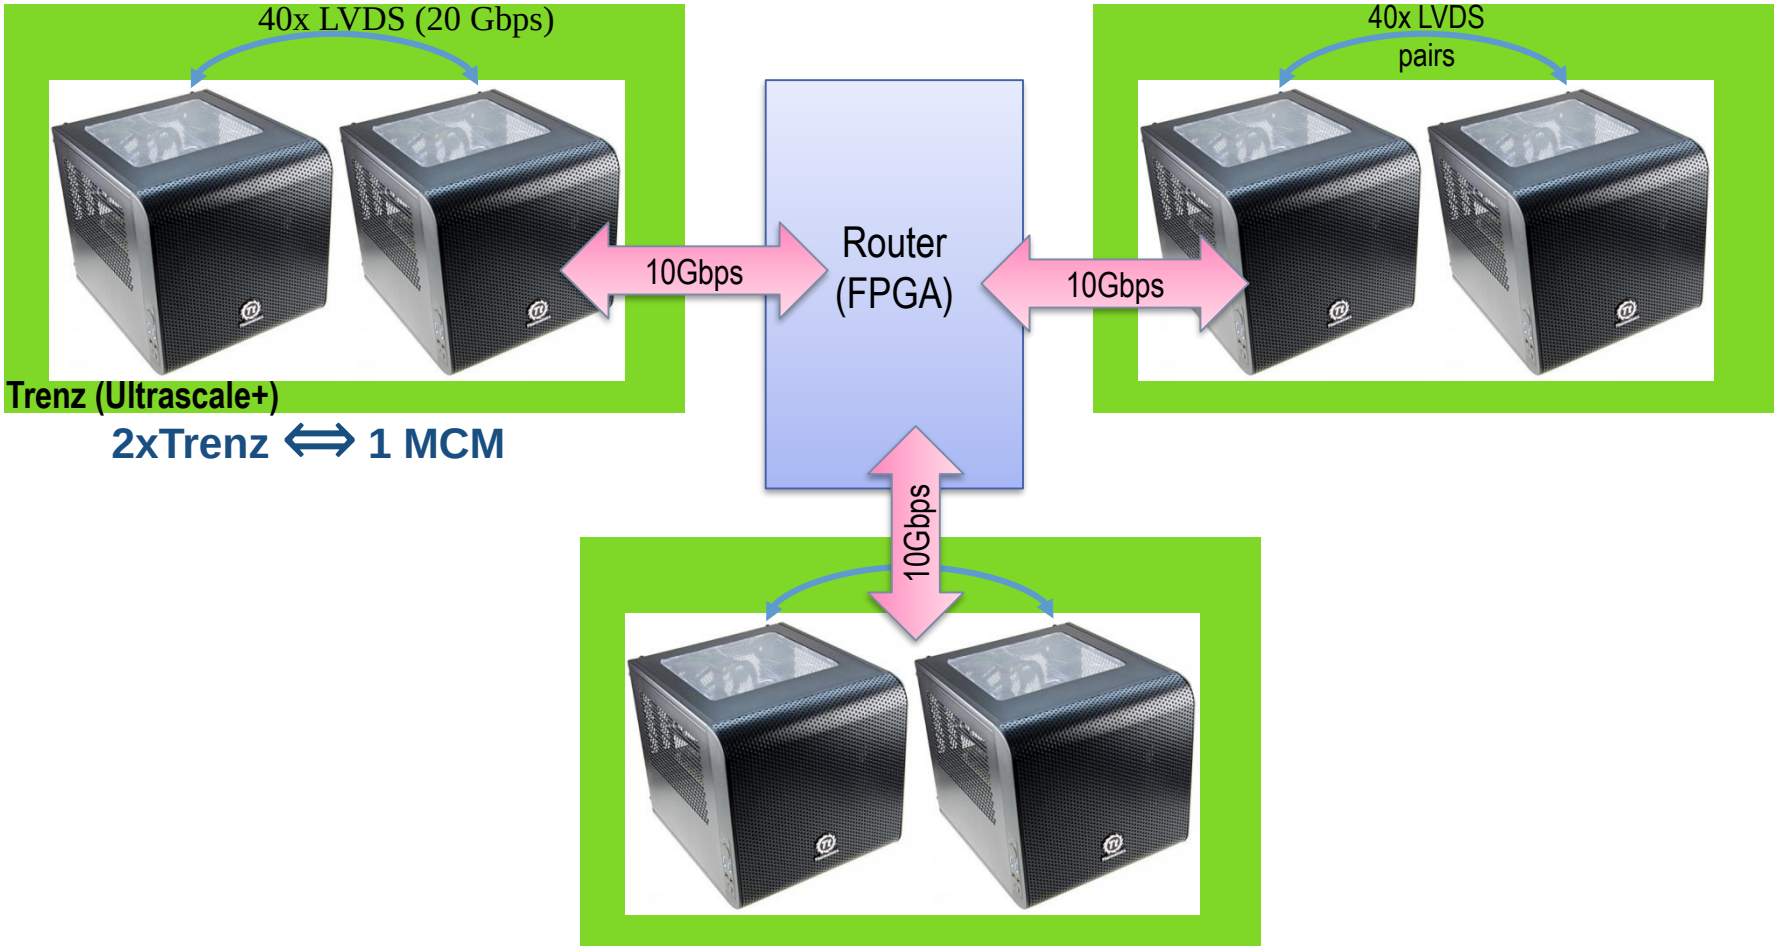

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

Early prototyping platforms

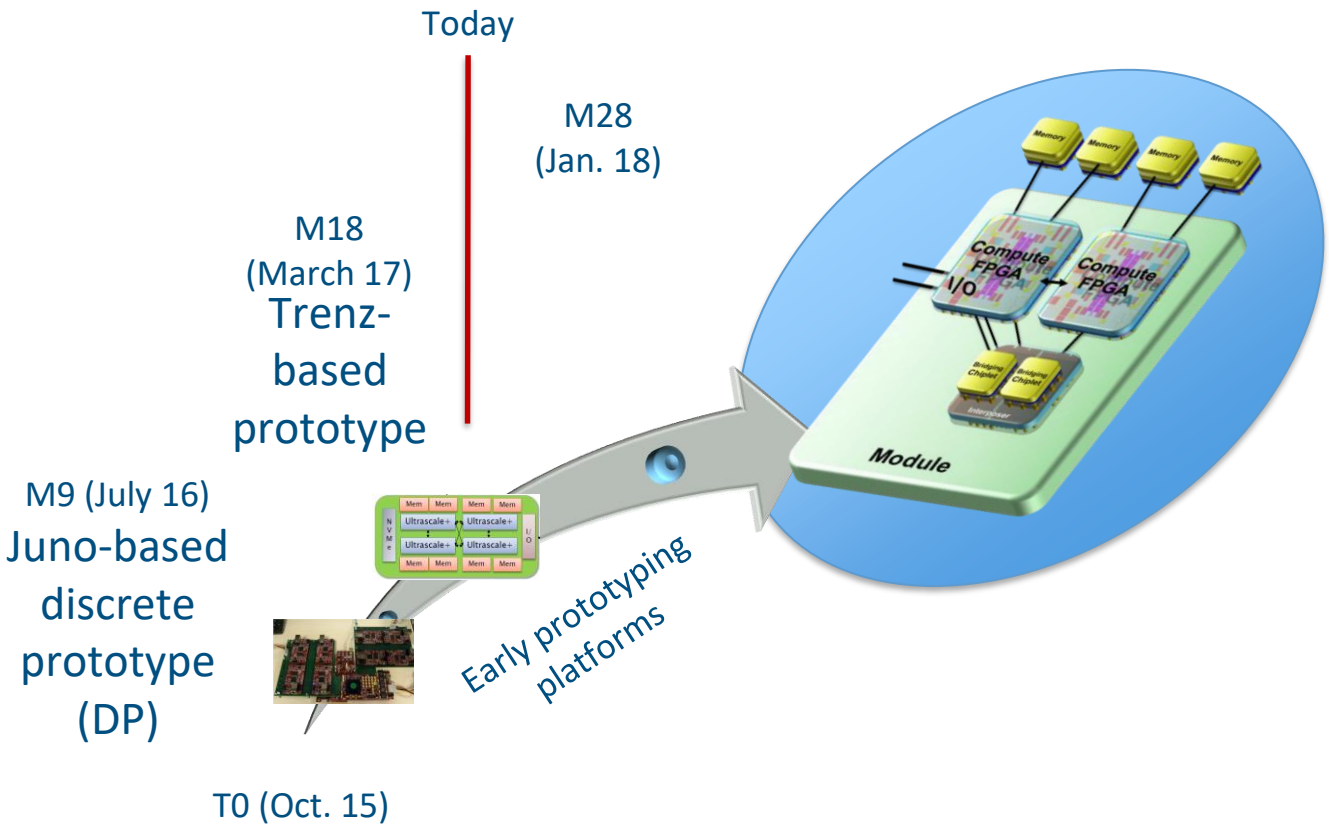

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

Early prototyping platforms

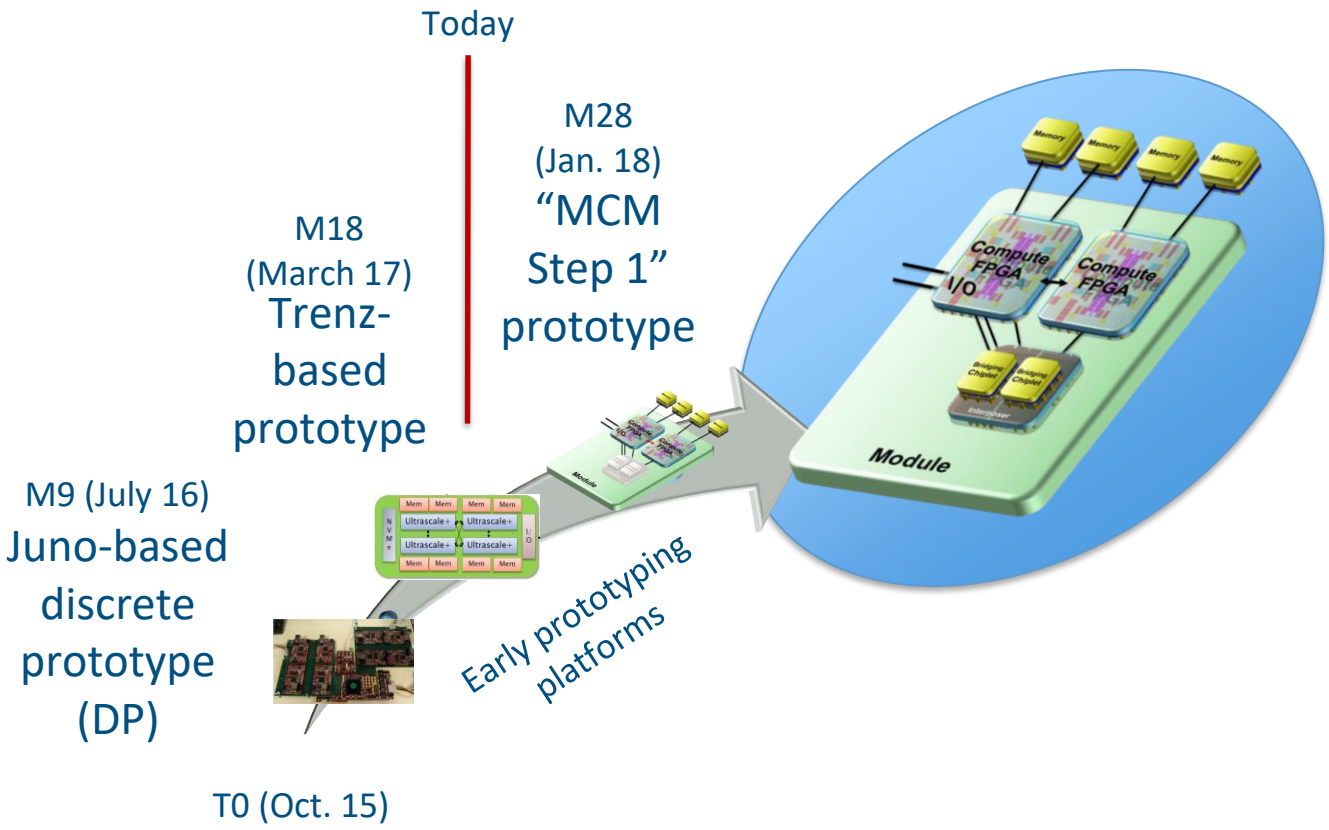

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

Early prototyping platforms

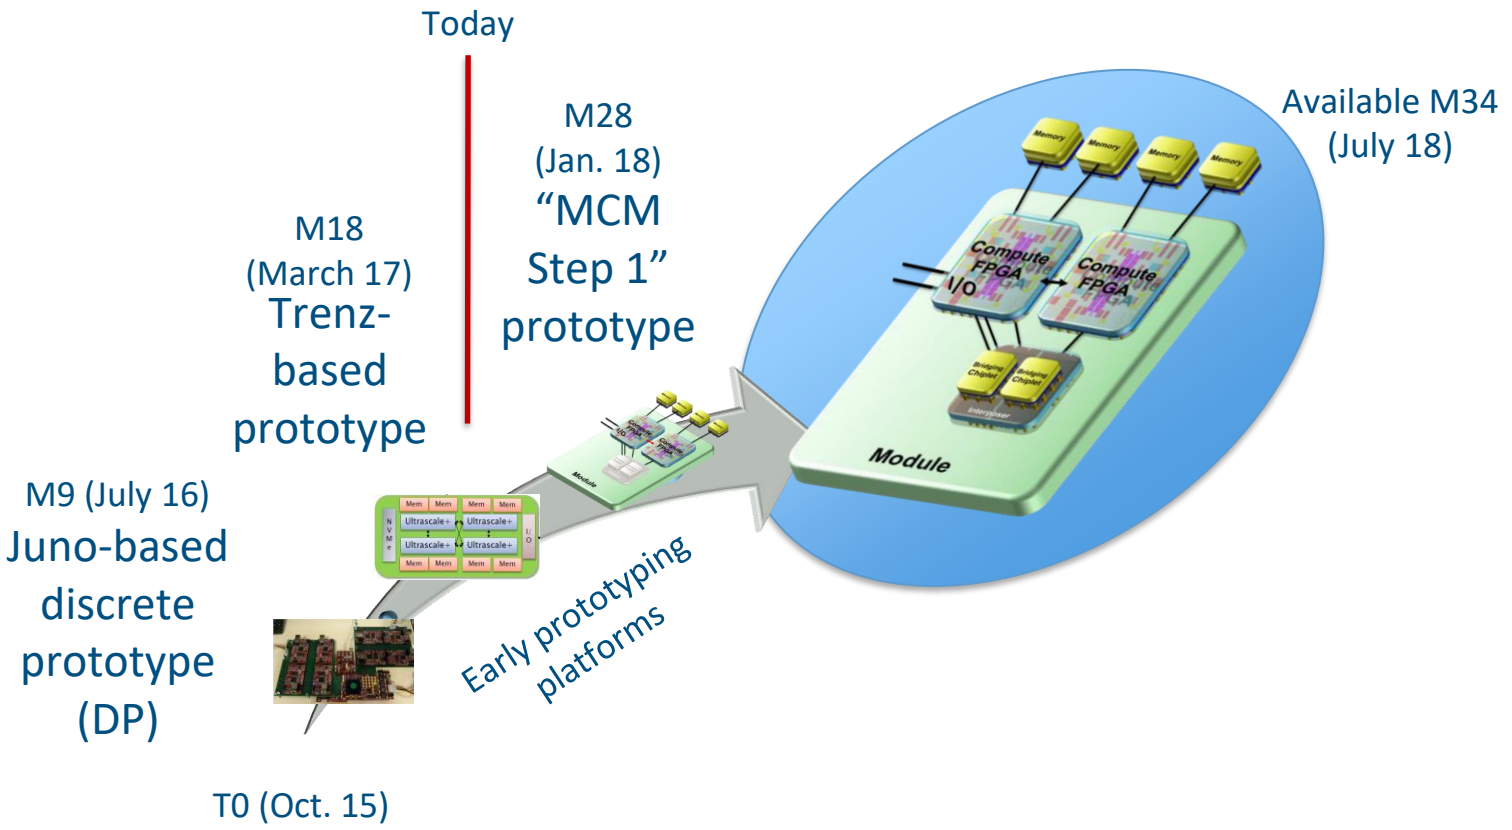

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

From approach to prototype

What has to be demonstrated?

- 3D technology capability for HPC module
- UNIMEM capability
- Scalability:
  - in term of architecture
  - in term of performance

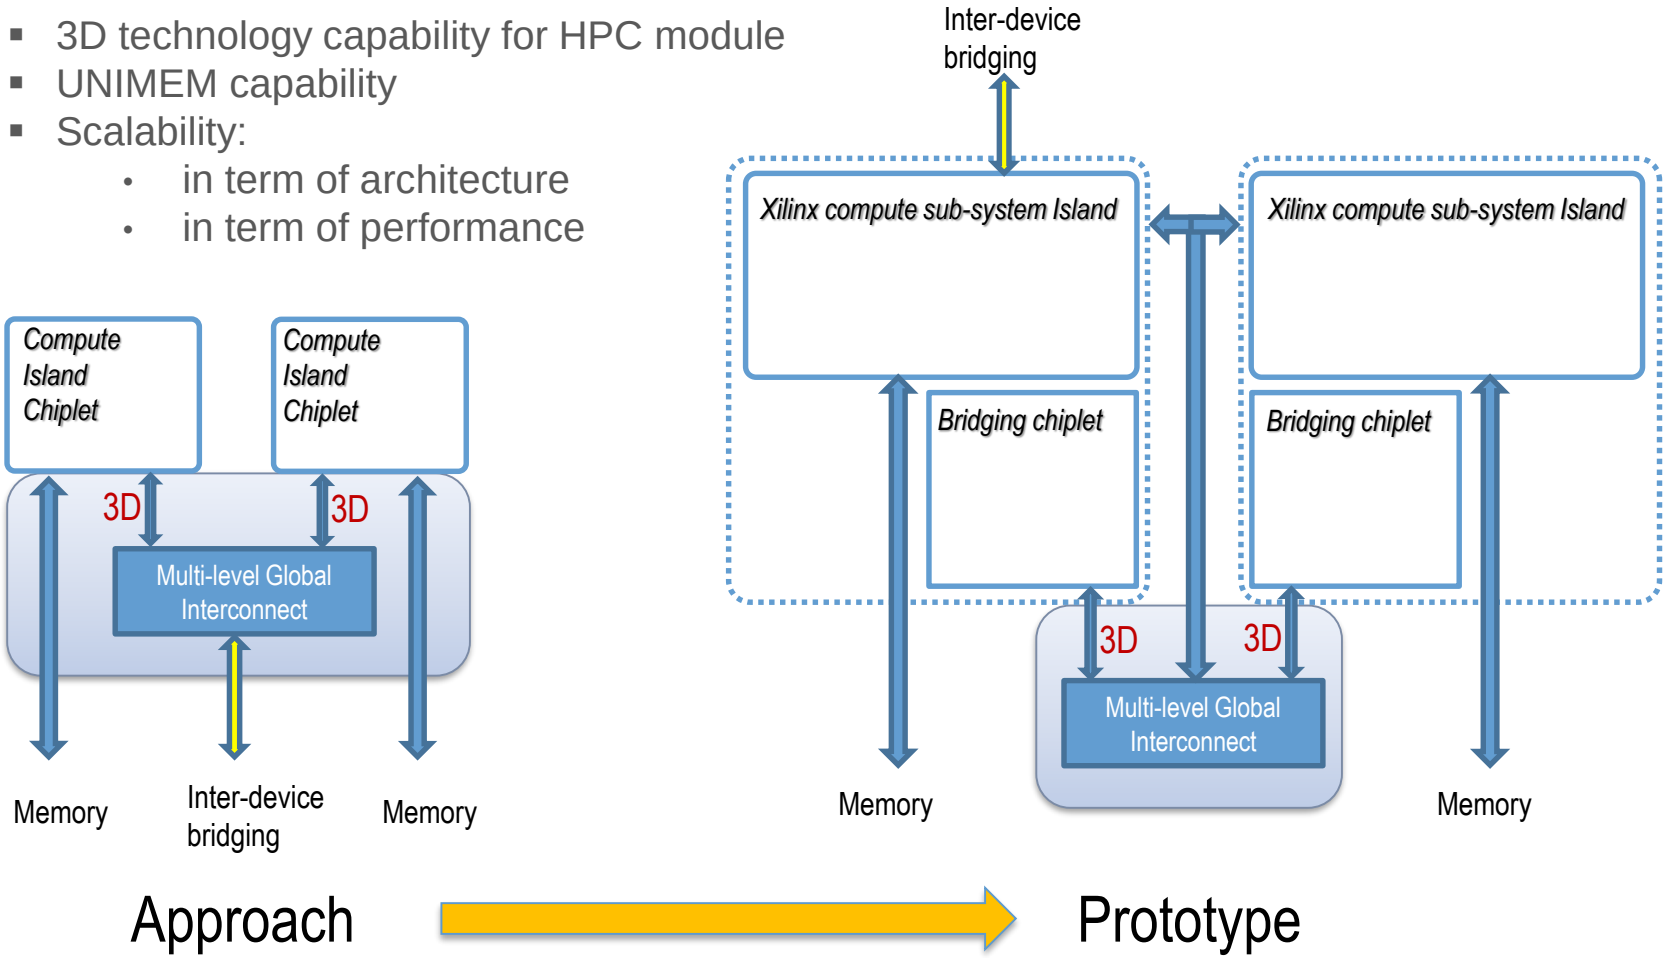

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

4. Application and Software

TSV µ-bumps	Chiplet / Interposer design & manuf.	Die on interposer, MCM	UNIMEM	ARMv8	Linux, Virtualization, Firmware	Power & thermal management	MPI, GPI, OmpSs, OpenStream	Mini-apps

ARMv8 cores in Xilinx MPSoCs

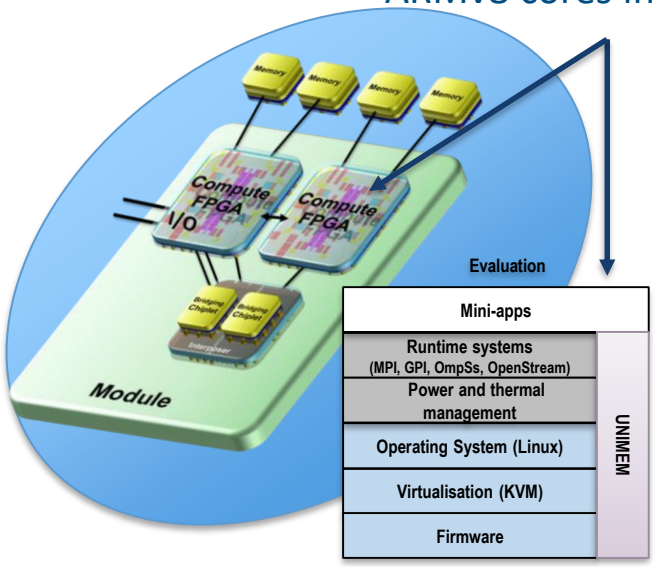

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

ExaNoDe SW objectives

- **Mini-apps for co-design process**
  - Select of HPC applications to co-design the ExaNoDe architecture.
- **Software infrastructure**
  - Deploy a software ecosystem for the ARM-based compute node …
  - in conjunction with the UNIMEM system architecture.

ExaNoDe Software Architecture

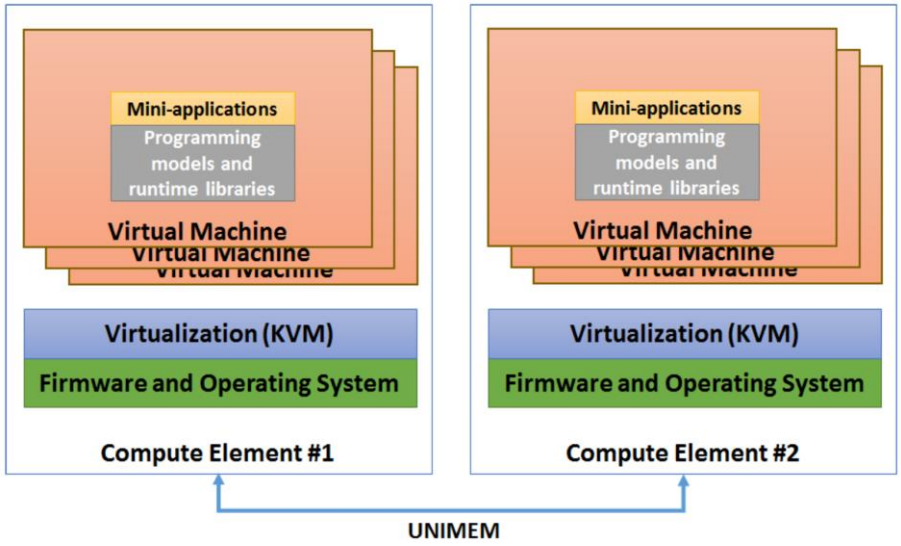

- **Evaluation**
  - Analyse and compare ExaNoDe architecture.

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

Application Portfolio

- **BQCD**
  - Simulation of Quantum Chromodynamics on a lattice (LQCD)

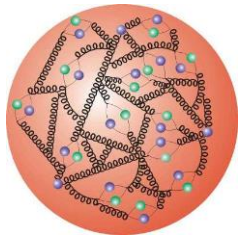

- **HydroC**
  - Simplified version of the astrophysical code RAMSES

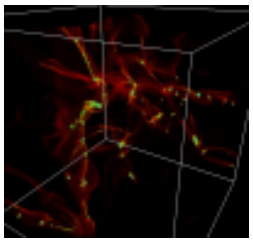

- **KKRnano**
  - Materials science application based on the Density Functional Theory (DFT) method

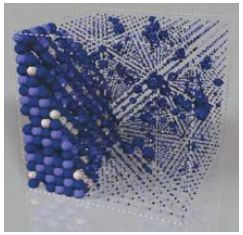

- **MiniFE**
  - Mini-application that mimics the finite element generation, assembly and solution for an unstructured grid problem

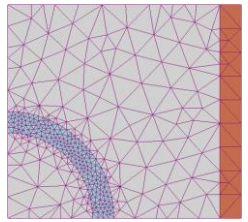

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

SW stack: deployment support

- **Deliver firmware and Numa-aware OS**
  - For multi-board and ExaNoDe prototypes

  (Unimem data movement, memory protection and integration with peripheral devices; OS interface for light RDMA operations)

- **Provide virtualization layer**

  (next slide)

- **Support programming models:**
  - Enable portable exploitation of UNIMEM (MPI, GPI, OmpSs, OpenStream)

- **Evaluate UNIMEM architecture**
  - Latency, bandwidth, memory footprint, CPU usage, …

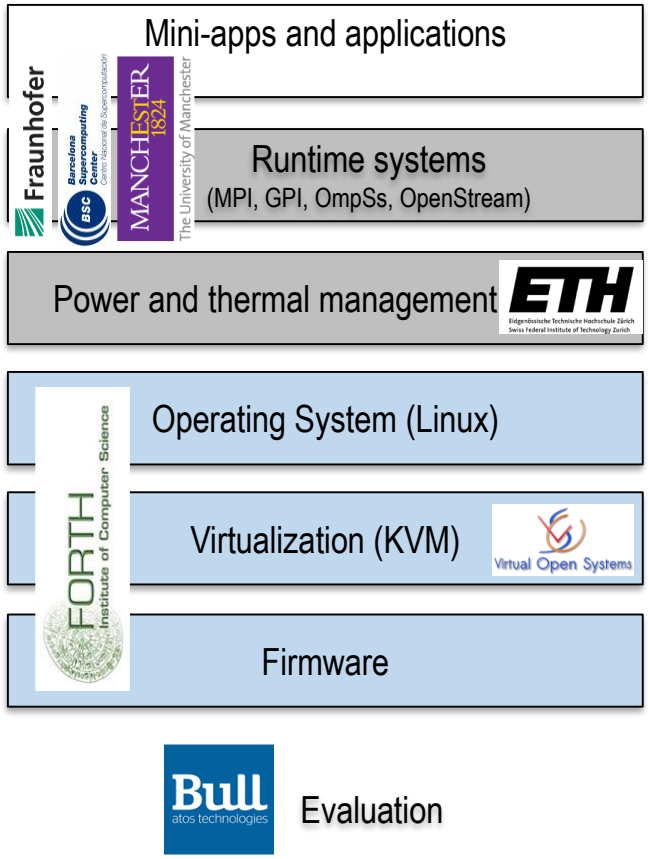

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

SW stack: Virtualization Layer

Virtual Open Systems

- For deployment ease,
  compatibility,
  efficiency of resource usage

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

SW stack: Virtualization Layer

Virtual Open Systems

- **For deployment ease,**
  compatibility,
  efficiency of resource usage

- **For an increased flexibility and reliability**
  - **Snapshot:** save now, restore later, maybe elswhere
  - **Checkpoint**: periodic and incremental, to cope with HW or SW failures

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

SW stack: Virtualization Layer

Virtual Open Systems

- **For deployment ease,**
  compatibility,
  efficiency of resource usage

- **For an increased flexibility and reliability**
  - **Snapshot:** save now, restore later, maybe elsewhere
  - **Checkpoint**: periodic and incremental, to cope with HW or SW failures

- **For performance KVM + paravirtualization**
  - **API Remoting:** use accelerator APIs inside VM
    UNIMEM atomics and RDMA, MPI, OpenCL …

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

SW stack: Virtualization Layer

Virtual Open Systems

- **For deploymer**
  compatibil
  efficiency
- **For an increas**
  - **Snapshot:** save
  - **Checkpoint**: peri r SW failures

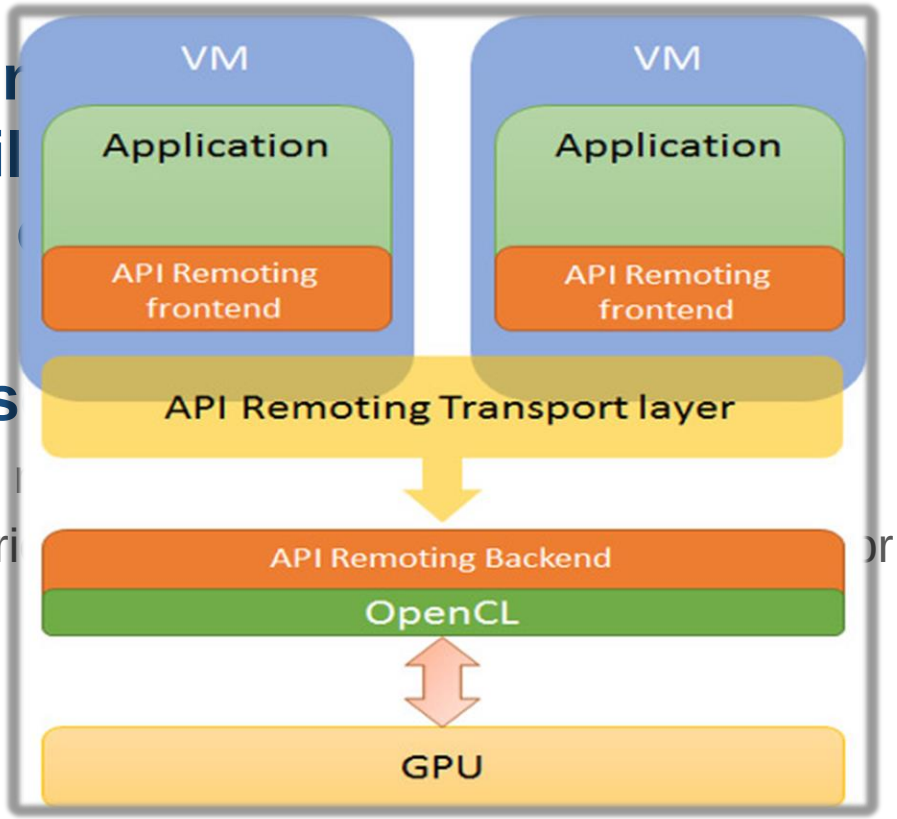

- **For performance KVM + paravirtualization**
  - **API Remoting:** use accelerator APIs inside VM
    UNIMEM atomics and RDMA, MPI, OpenCL …

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

Platform comparisons

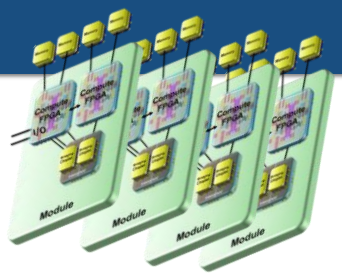

- **ExaNoDe prototype**
- **EUROSERVER prototype based on Juno** ARM
- **Mainstream servers based on Intel Xeon**
  - Nehalem, Sandybridge, Haswell

- **Emerging server architectures**
  - Intel KNL
  - IBM POWER8
  - Cavium ThunderX, APM X-Gene

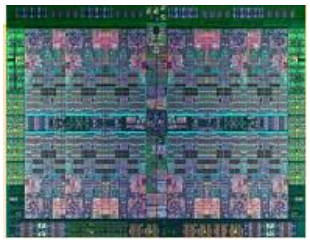

- **Low-end platforms**
  - Intel Atom
  - Raspberry Pi

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

Conclusion

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

ExaNoDe: Conclusion

- **ExaNoDe concept:**
  - **Architecture**: Many simple cores instead of few complex ones
  - **Integration**: Many simple heterogeneous chiplets instead of few complex System-on-Chip
  - **Compute node interconnect**: Active silicon interposer

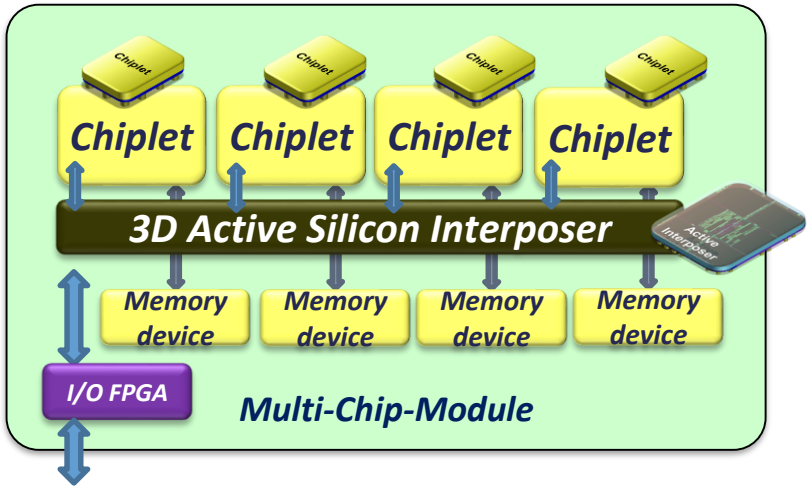

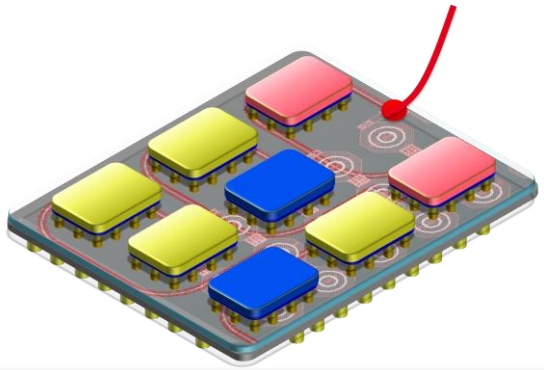

Copyright © 2017 Members of the ExaNoDe Consortium

contact: denis.dutoit@cea.fr, project coordinator

Thank you!

European Exascale Processor & Memory Node Design

Contact: denis.dutoit@cea.fr (Project coordinator)
k.pouget@virtualopensystems.com (Presenter)