




# E4

COMPUTER  
ENGINEERING

WHEN  
PERFORMANCE  
MATTERS

# E4 COMPUTER ENGINEERING: TECHNOLOGICAL LEADERSHIP FOR



Fabrizio Magugliani  
Strategic Planning and Business Development Unit  
[fabrizio.magugliani@e4company.com](mailto:fabrizio.magugliani@e4company.com)

The 12th International Conference on Internet and Distributed Computing Systems  
NAPLES, OCT. 10-12, 2019



# THE COMPANY

Since 2002, E4 Computer Engineering has been innovating and actively encouraging the adoption of new computing and storage technologies. Because new ideas are so important, we invest heavily in research and hence in our future. Thanks to our comprehensive range of hardware, software and services, we are able to offer our customers complete solutions for their most demanding workloads in: HPC, Big-Data, AI, Deep Learning, Data Analytics, Cognitive Computing and for any challenging Storage and Computing requirements.

**E4. When Performance Matters.**

# OUR MEMBERSHIPS



Silver Level

Cosimo Gianfreda IBM **CHAMPION** 



Member of CERN openlab



Member of the Steering Board

<http://www.etp4hpc.eu>



Member of the OEHI (Open Edge and HPC Initiative)



Member of the Consortium

<http://european-processor-initiative.com>



Member of the MaX Center of Excellence

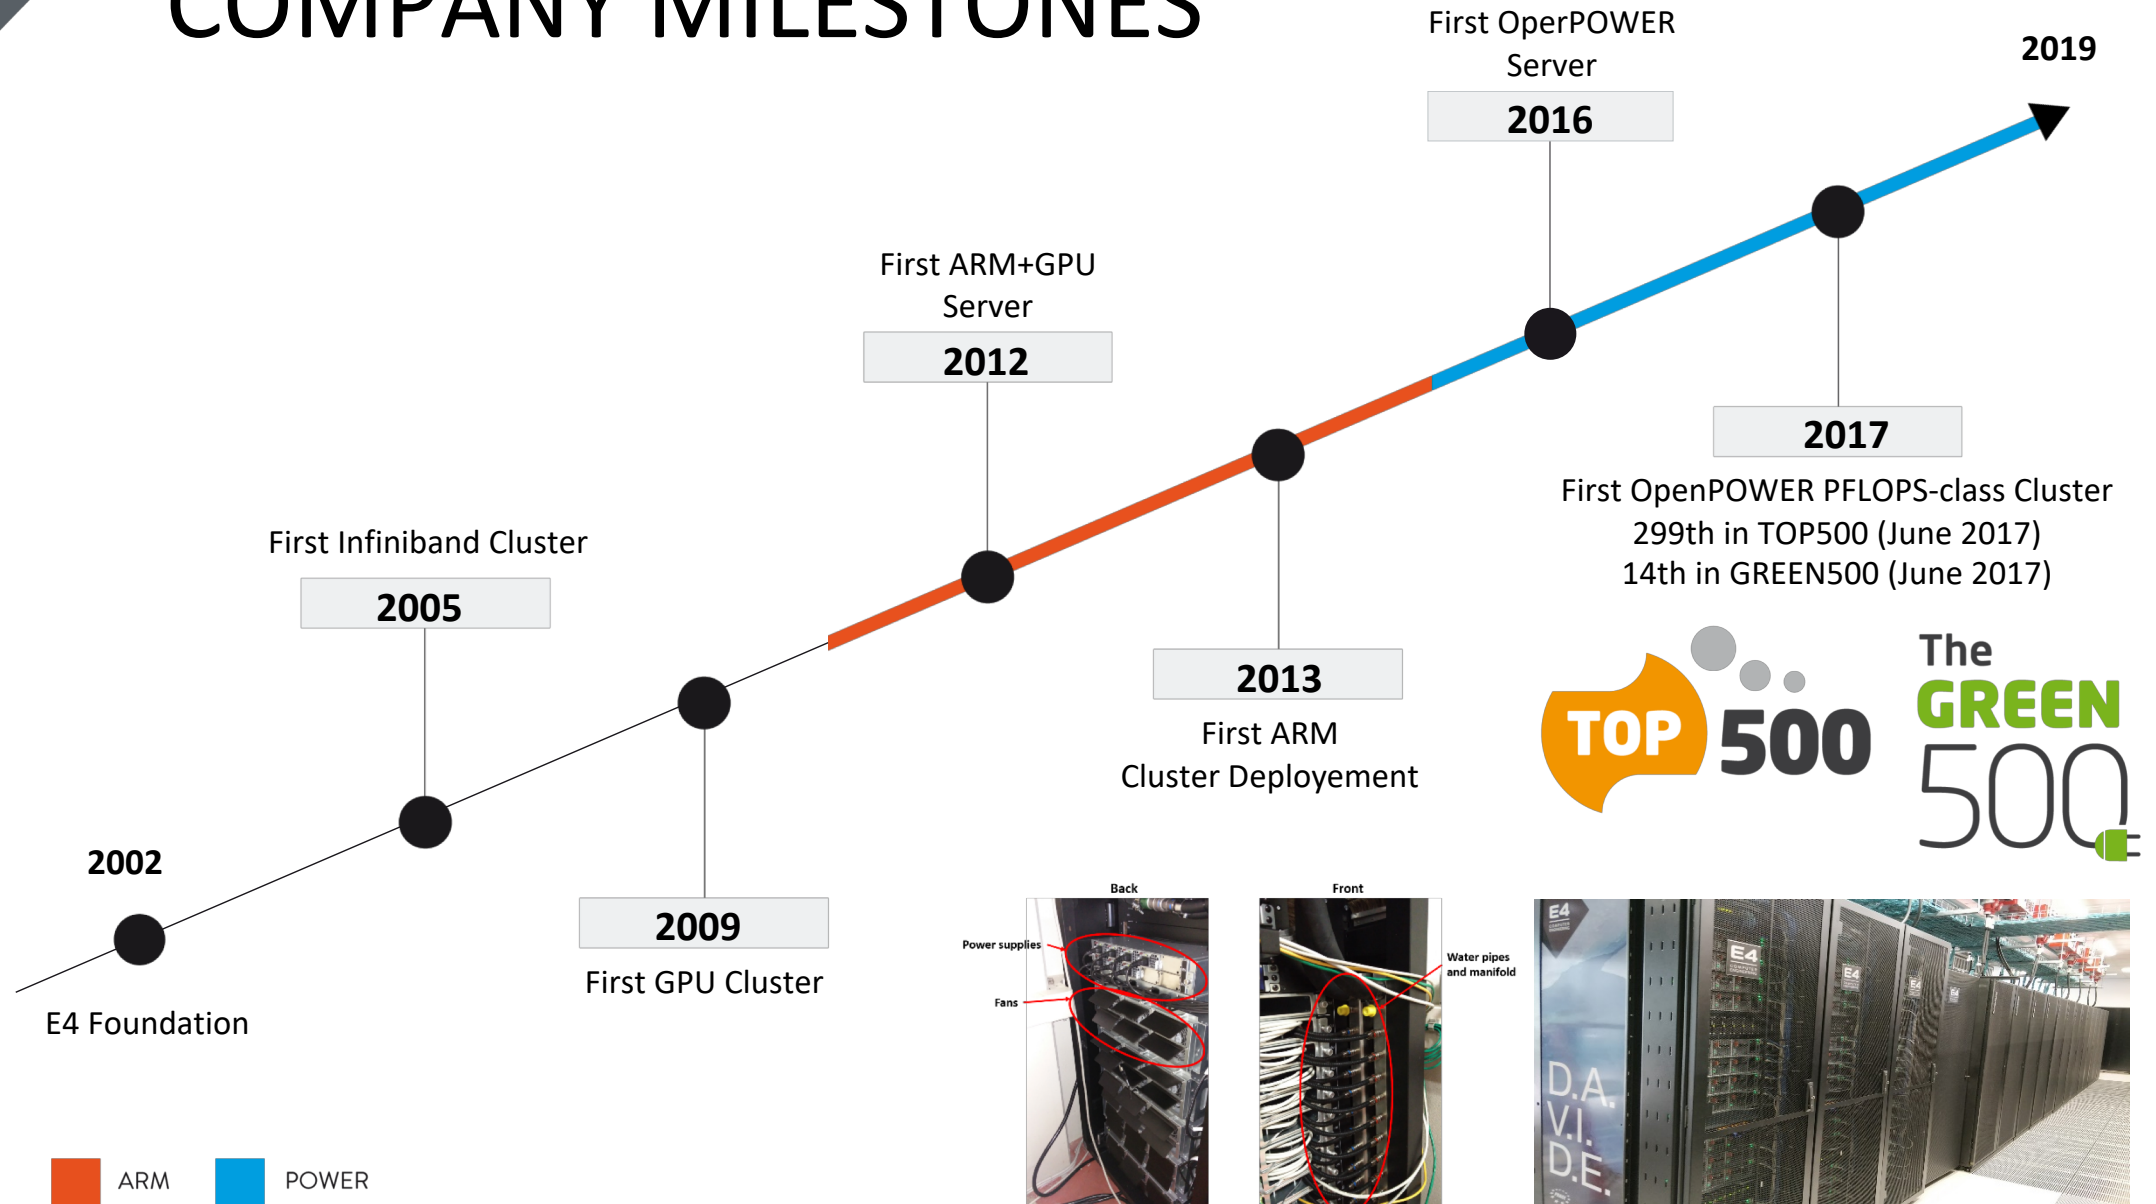
Member of the  
**Open  FOAM® Partnership Programme**



Member of HiPEAC



# COMPANY MILESTONES



# A BIT OF HISTORY....

Once upon a time



was looking for a compute cluster...



# ... AND BLACKJEANS WAS BORN!



## Front-end node

Intel(R) Xeon(R) CPU E5606 @ 2.13GHz

## Nodes 1/12

Intel(R) Xeon(R) CPU X5650 @ 2.67GHz

Tesla M2050

Infiniband Mellanox MT26428

## Nodes 13/18

Intel(R) Xeon(R) CPU E5-2680 v2 @ 2.80GHz

Infiniband Intel TrueScale QLE 7340

## WebServ Node

Intel(R) Xeon(R) CPU E5-2650 0 @ 2.00GHz







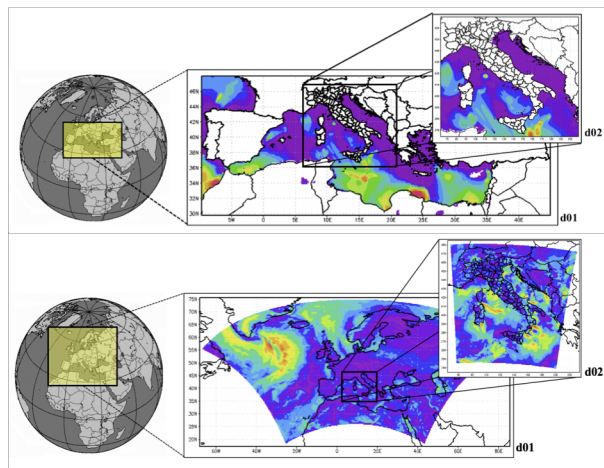
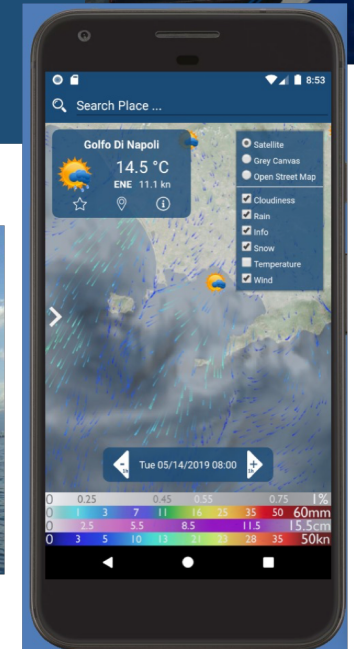
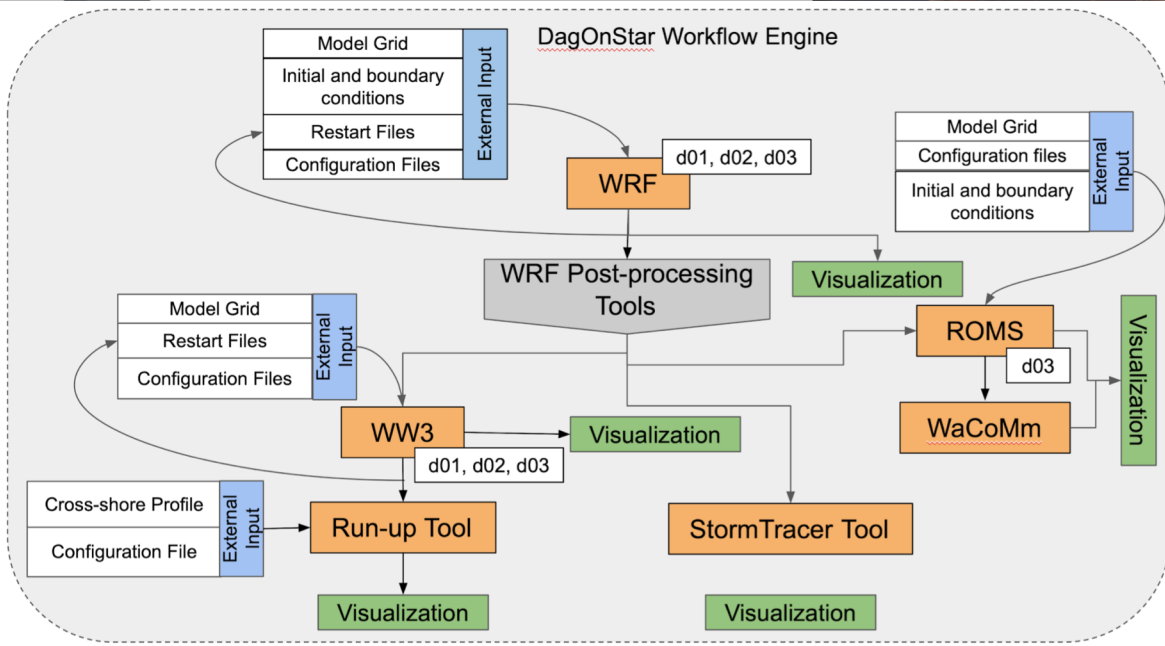

Switch Infiniband QLogic 12200

## Storage

Dot Hill D2732


E4 HPC OpenSuite 2.1 (cluster management suite)







# AND GREW STRONGER...



## **2019 EXPANSION**

### **Nodes 19/22**

**Intel(R) Xeon(R) Xeon 20-Core 6230 @ 2,1Ghz 27.5MB**

**Infiniband Mellanox CX4 VPI SinglePort FDR IB 56Gb/s x8**

Front-end node

Intel(R) Xeon(R) CPU E5606 @ 2.13GHz

Nodes 1/12

Intel(R) Xeon(R) CPU X5650 @ 2.67GHz

Tesla M2050

Infiniband Mellanox MT26428

Nodes 13/18

Intel(R) Xeon(R) CPU E5-2680 v2 @ 2.80GHz

Infiniband Intel TrueScale QLE 7340

WebServ Node

Intel(R) Xeon(R) CPU E5-2650 0 @ 2.00GHz

Switch Infiniband QLogic 12200

Storage

Dot Hill D2732

E4 HPC OpenSuite 2.1 (cluster management suite)

THEN THE REQUIREMENTS SKYROCKETED...

AND I HEARD ABOUT PURPLEJEANS...




# PURPLE WHAT????


Purple rain, [Prince](#) and [The Revolution](#), 1984




Deep Purple in Rock, [Deep Purple](#), 1970



*The Color Purple*, [Alice Walker](#), 1982



The Original Purple Bed. *The ultimate No-Pressure® mattress.*



# REALLY, *PurpleJeans!*

This one??????

No, this one....  
I mean, this one....



Frontend : SuperServer 7049GP-TRT  
Intel(R) Xeon 10-Core 5215 2,5Ghz 13.75MB  
Infiniband Mellanox CX4 VPI SinglePort FDR IB 56Gb/s x16  
PNY Quadro RTX6000

Gnodes 1-4: SuperServer 1029GQ-TVRT  
Intel(R) Xeon(R) Xeon 16-Core 5218 2,3Ghz 22MB  
NVIDIA Tesla V100 32GB SXM2  
Infiniband Mellanox CX4 VPI SinglePort FDR IB 56Gb/s x16  
Nvidia Deep Learning Framework, Nvidia.Docker

Wnodes 1-4 SuperServer 6019P-WTR  
Intel(R) Xeon(R) Xeon 16-Core 5218 2,3Ghz 22MB  
Infiniband Mellanox CX4 VPI SinglePort FDR IB 56Gb/s x16

Switch Mellanox InfiniBand SX6012F 12 Port QSFP FDR

SuperServer 6029P-WTRT  
Stor1  
Intel(R) Xeon(R) Xeon 10-Core 4210 2,2Ghz 13.75MB  
Infiniband Mellanox CX4 VPI SinglePort FDR IB 56Gb/s x16

E4 HPC OpenSuite 2.1 (cluster management suite )

THE REASONS?

OUR PEOPLE

OUR SKILLS

OUR RELENTLESS PURSUE FOR ACHIEVING THE  
COMPLETE SATISFACTION OF THE CUSTOMER'S  
NEEDS

OUR “MANIACAL” FOCUS TOWARDS THE PEOPLE  
WE SERVE



AN(OTHER...) EXAMPLE:

PRACE-3IP PCP: Whole-System Design  
for Energy Efficient HPC





**D.A.V.I.D.E.  
SUPERCOMPUTER**  
(Development of an  
Added  
Value  
Infrastructure  
Designed in  
Europe)





# D.A.V.I.D.E. SUPERCOMPUTER

(Development of an Added Value Infrastructure Designed in Europe)

OCP form-factor compute node  
based on IBM Minsky

4x  **nvidia**. Tesla P100 HSMX2

2 x  **POWER8** with NVLink

2xIB EDR

BusBar

E4/Bologna University  
POWER MANAGEMENT  
COMPONENTS



LIQUID COOLING

# PCP PHASE III – D.A.V.I.D.E. SUPERCOMPUTER

(Development of an Added Value Infrastructure Designed in Europe)


Power- and Energy-monitoring & management infrastructure (in collaboration with University of Bologna, ETHZ)

- Off-the-shelf components
- High speed and accurate per-node power sensing synchronized among the nodes
- Data accessible out-of-band and without processor intervention
- Out-of-Band and synchronized fine grain performance sensing
- Dedicated data-collection subsystem running on management nodes
- Predictive Power Aware job scheduler and power manager

# KEY ENABLING TECHNOLOGIES



ALMA MATER STUDIORUM  
UNIVERSITÀ DI BOLOGNA





E4 PPBB (Power and Performance «Black Box»)


- Out-of-band => Zero overhead
- Fine grain => down to ms scale
- IoT communication technology => scalable
- Time synchronous (NTP, PTP)

1 Node -20 min

E4 PPBB Vi





45 Nodes -4s



45 Nodes -1s



# KEY ENABLING TECHNOLOGIES



Job info Performance Energy

Job ID: [REDACTED] | User ID: [REDACTED] | Account name: [REDACTED]

## Times

Duration: 9 mins 59 secs

Queue time: 11:07:55 21/08/17

Start time: 11:07:55 21/08/17

End time: 11:17:54 21/08/17

Status  
Finished



Average Power  
684.4 W

Sys Utilization  
77.4 %

CPUs Utilization  
92.34 %

Average Temperature  
46.24 °C

## Cores' Load



Overview 3D View

OPEN TOP RESET

Node: node522

Power Consumption : 112.50



# E4: LOOKING FORWARD

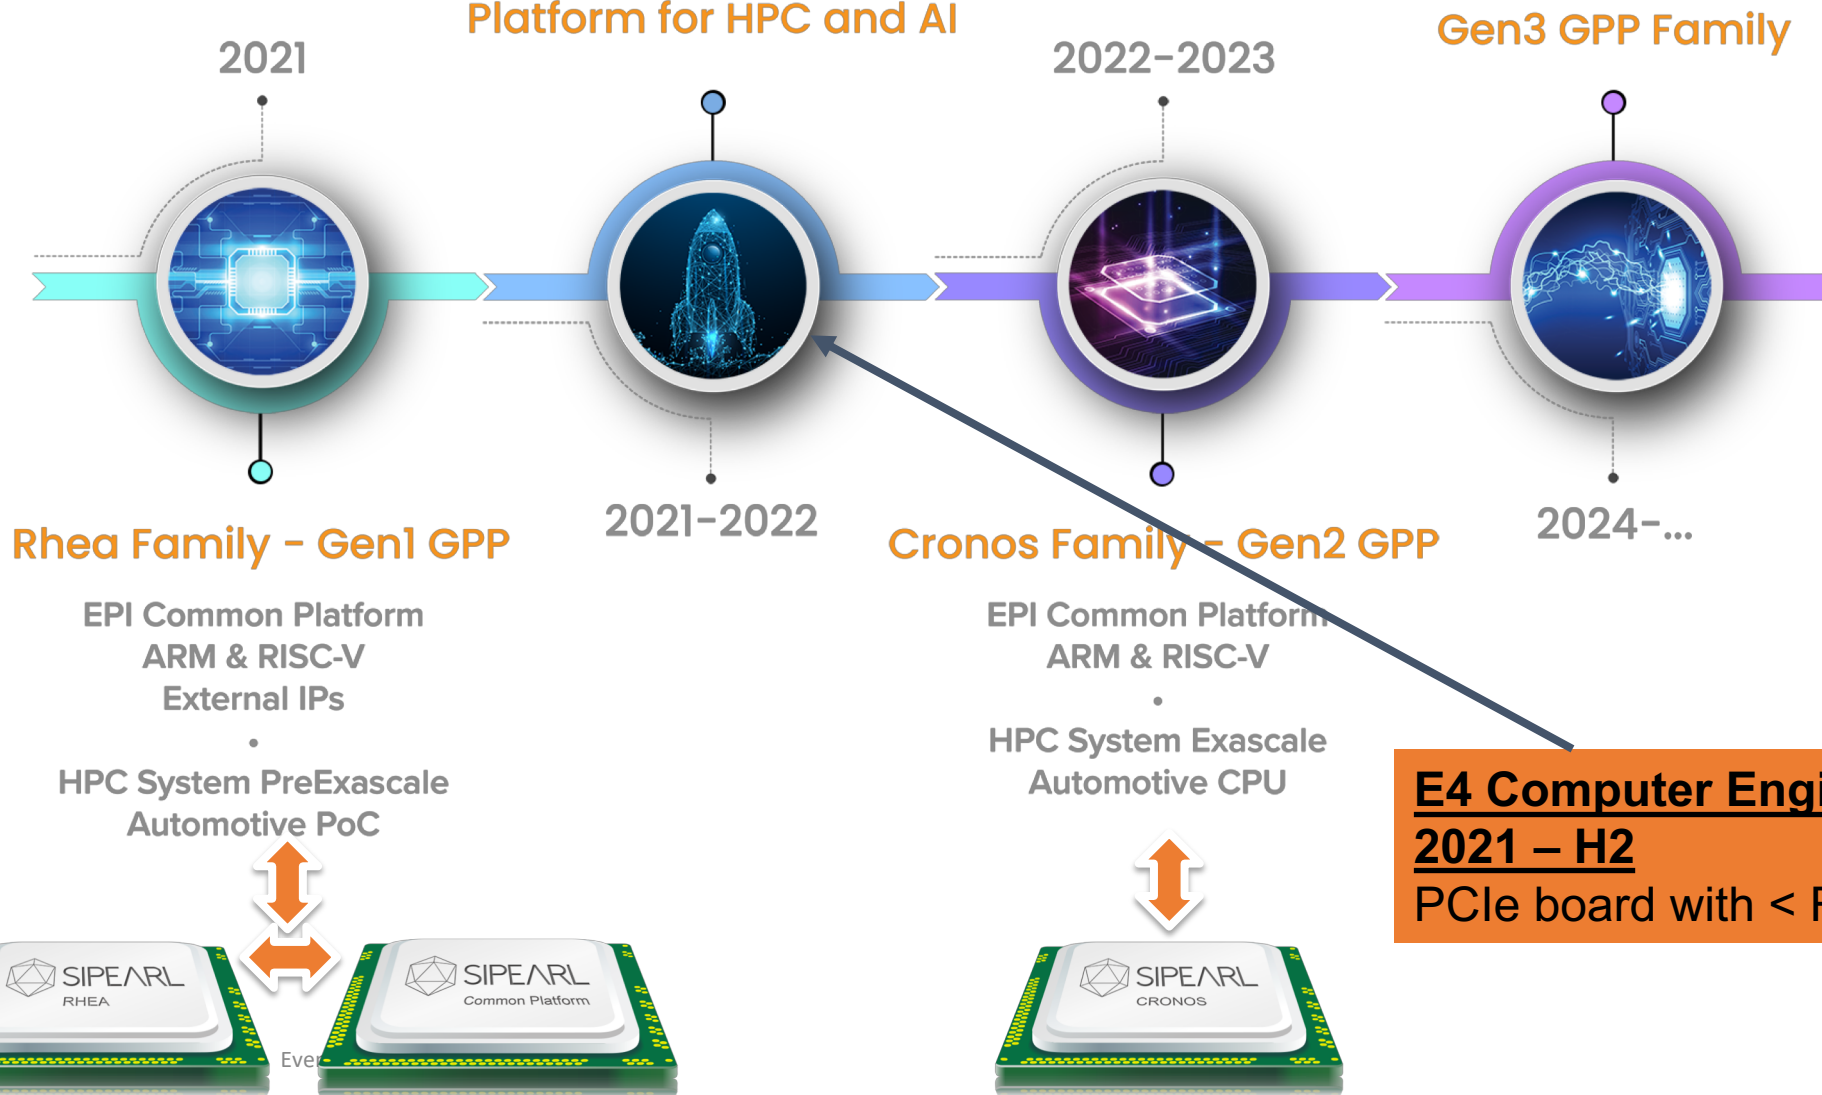
## FRAMEWORK PARTNERSHIP AGREEMENT IN EUROPEAN LOW-POWER MICROPROCESSOR TECHNOLOGIES

THIS PROJECT HAS RECEIVED FUNDING FROM THE EUROPEAN UNION'S HORIZON 2020 RESEARCH AND INNOVATION PROGRAMME UNDER GRANT AGREEMENT NO 826647



# E4: LOOKING FORWARD

EPI IP's Launch Pad  
&  
Pan European Research  
Platform for HPC and AI





# E4'S FACILITIES



# E4 SYSTEM INTEGRATION FACILITY

Our production takes place in 2 Burn-In rooms through automated tools for installation, configuration and quality control.

In 2016 two major CERN orders required the production of ~ 2300 servers and 200 JBOD, in total 30PB and 46756 cores.

We adopt Open Source Software: Cobbler – Puppet – Nagios – Graphite.  
We internally support the “E4tools” application development for quality control.

High quality systems: failure rate < 1%



Cobbler Web Interface

100.100.6.181/cobbler\_web/system/list

Logged in: cobbler [Logout](#)

**E4 COBBLER**

**Configuration**

- Distros
- Profiles
- Systems
- Repos
- Images
- Kickstart Templates
- Snippets
- Management Classes
- Settings

**Resources**

- Packages
- Files

**Systems**

[Create New System](#) [Batch Actions](#) [Go](#) [Items/page: 50](#) [Page 1](#)

Name ↓	Profile	Status	Netboot_Enabled	Actions
<input type="checkbox"/> 617-1	SLC67-x86_64	production	<input type="checkbox"/>	<a href="#">Edit</a> <a href="#">Copy</a> <a href="#">Rename</a> <a href="#">Delete</a> <a href="#">View Kickstart</a>
<input type="checkbox"/> 617-2	SLC67-x86_64	production	<input type="checkbox"/>	<a href="#">Edit</a> <a href="#">Copy</a> <a href="#">Rename</a> <a href="#">Delete</a> <a href="#">View Kickstart</a>
<input type="checkbox"/> 617-3	SLC67-x86_64	production	<input type="checkbox"/>	<a href="#">Edit</a> <a href="#">Copy</a> <a href="#">Rename</a> <a href="#">Delete</a> <a href="#">View Kickstart</a>
<input type="checkbox"/> 617-4	SLC67-x86_64	production	<input type="checkbox"/>	<a href="#">Edit</a> <a href="#">Copy</a> <a href="#">Rename</a> <a href="#">Delete</a> <a href="#">View Kickstart</a>
<input type="checkbox"/> 618-1	Destroy	production	<input type="checkbox"/>	<a href="#">Edit</a> <a href="#">Copy</a> <a href="#">Rename</a> <a href="#">Delete</a> <a href="#">View Kickstart</a>
<input type="checkbox"/> 618-2	Destroy	production	<input type="checkbox"/>	<a href="#">Edit</a> <a href="#">Copy</a> <a href="#">Rename</a> <a href="#">Delete</a> <a href="#">View Kickstart</a>
<input type="checkbox"/> 618-3	Destroy	production	<input type="checkbox"/>	<a href="#">Edit</a> <a href="#">Copy</a> <a href="#">Rename</a> <a href="#">Delete</a> <a href="#">View Kickstart</a>



# E4 R&D FACILITY

## TESTING (HARDWARE)

- New Components
  - Systems
- Integration



## TESTING (SOFTWARE)

- Inter-compatibility
  - Inter-operability
- SDDC





# E4 MANUFACTURING FACILITY

OUR PRODUCTION FACILITIES INCLUDE R&D, DATA CENTER, BURN-IN ROOMS AND REGULAR CHECKS TO ENSURE ONLY THE HIGHEST QUALITY SYSTEMS ARE ALLOWED TO LEAVE OUR PREMISES


- 5.382 ft<sup>2</sup> of committed production area
- 2 burn-in rooms kept at test temperature 27/30°c
- 1 Datacenter
- 1 R&D Facility
- Each system individually tested
- Regular stock rotation according to day of arrival
- Environment 100% compliant to health & safety regulations
- Stock split by vendor and typology
- Periodic stock check and cleaning
- Thorough and frequent in/out checks to minimize human




## SETTINGS



## COMPONENTS UPDATE



## STRESS TEST



## PERFORMANCE MEASUREMENT



## SYSTEM TESTS

That's what we mean with  
**TECHNOLOGICAL LEADERSHIP!**



# AND NOW...

## THE «mandatory» TIMELAPSE

BUILDING PurpleJeans!