Technologies for the future HPC systems

Jean-Pierre Panziera

Teratec 2011

3 petaflop systems : TERA 100, CURIE & IFERC

1.25 PetaFlops

140 000+ Xeon cores
SMP nodes

- **256 TB** memory
- **30 PB** disk storage
- **500 GB/s** IO throughput
- **580 m²** footprint

>1.6 PetaFlops

90 000+ Xeon cores
+SMP + GPU

- **360 TB** memory
- **10 PB** disk storage
- **250 GB/s** IO throughput
- **200 m²** footprint

>1.4 PetaFlops

70 000+ Xeon cores

- **280 TB** memory
- **15 PB** disk storage
- **120 GB/s** IO throughput
- **200 m²** footprint

~.5 PB/s memory total BW

More than PetaFlops !

Tera, Peta, Exa, Zetta, … Flops / Scale

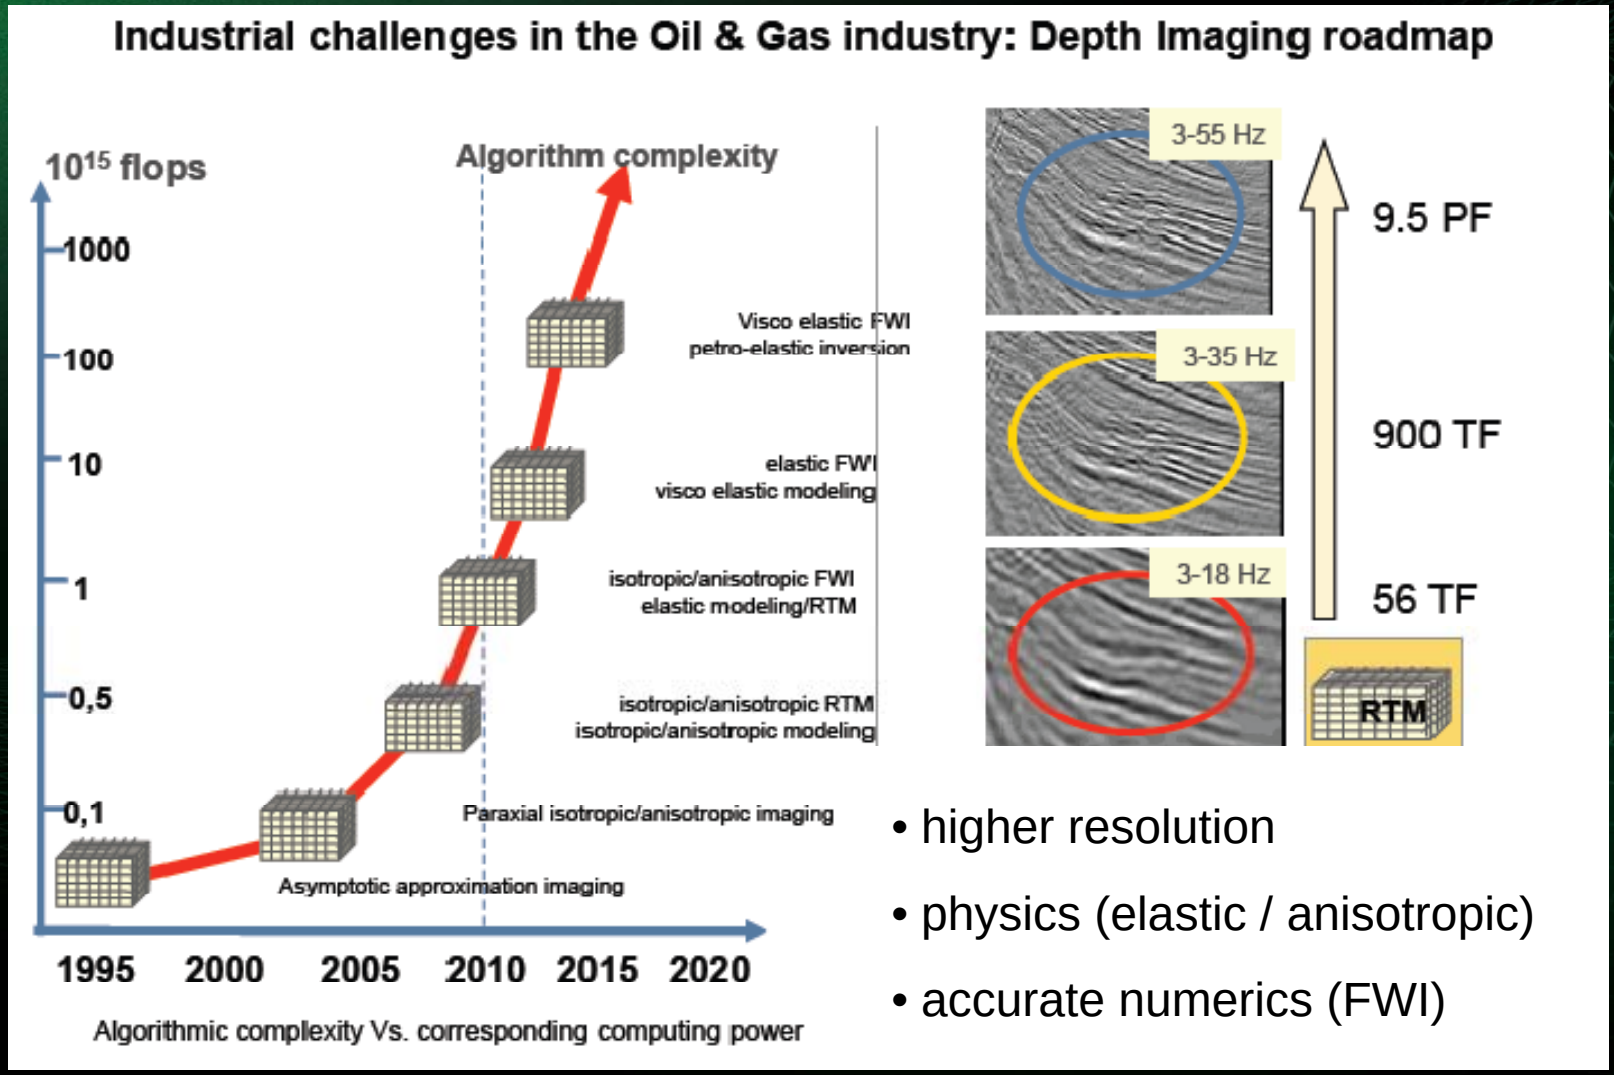

- higher resolution
- physics (elastic / anisotropic)
- accurate numerics (FWI)

source: exascale.org

HPC Applications

Electro-Magnetics

Computational Chemistry
Quantum Mechanics

Computational Chemistry
Molecular Dynamics

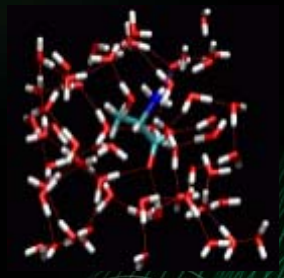

Computational Biology

Structural Mechanics
Implicit

Structural Mechanics
Explicit

A wide range of Domains and Applications

Seismic Processing

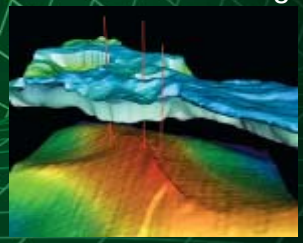

Computational Fluid
Dynamics

Reservoir Simulation

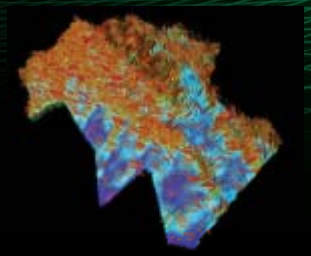

Rendering / Ray Tracing

Climate / Weather
Ocean Simulation

Data Analytics

More than PetaFlops !

Exascale Technology Challenges

- Processing Element : architecture and frequency
  - Multi/Many-cores, Accelerators, …
- Memory Capacity & BW → MCM, 3D Packaging ?
  - Feeding enough Bytes to the FP engines, fast enough
- Network bandwidth, latency, topology and routing
  - Optical connections/cables , fewer hops, compact packaging
- I/O scalability and flexibility
  - XXXLarge datasets + faster computations → data explosion
- System-level resiliency and reliability
  - Month(s) long jobs getting through HW failures
- Power and Cooling
  - Fewer & less consuming components, improved PUE
- Price ?

HPC requirements, constraints and technology evolution

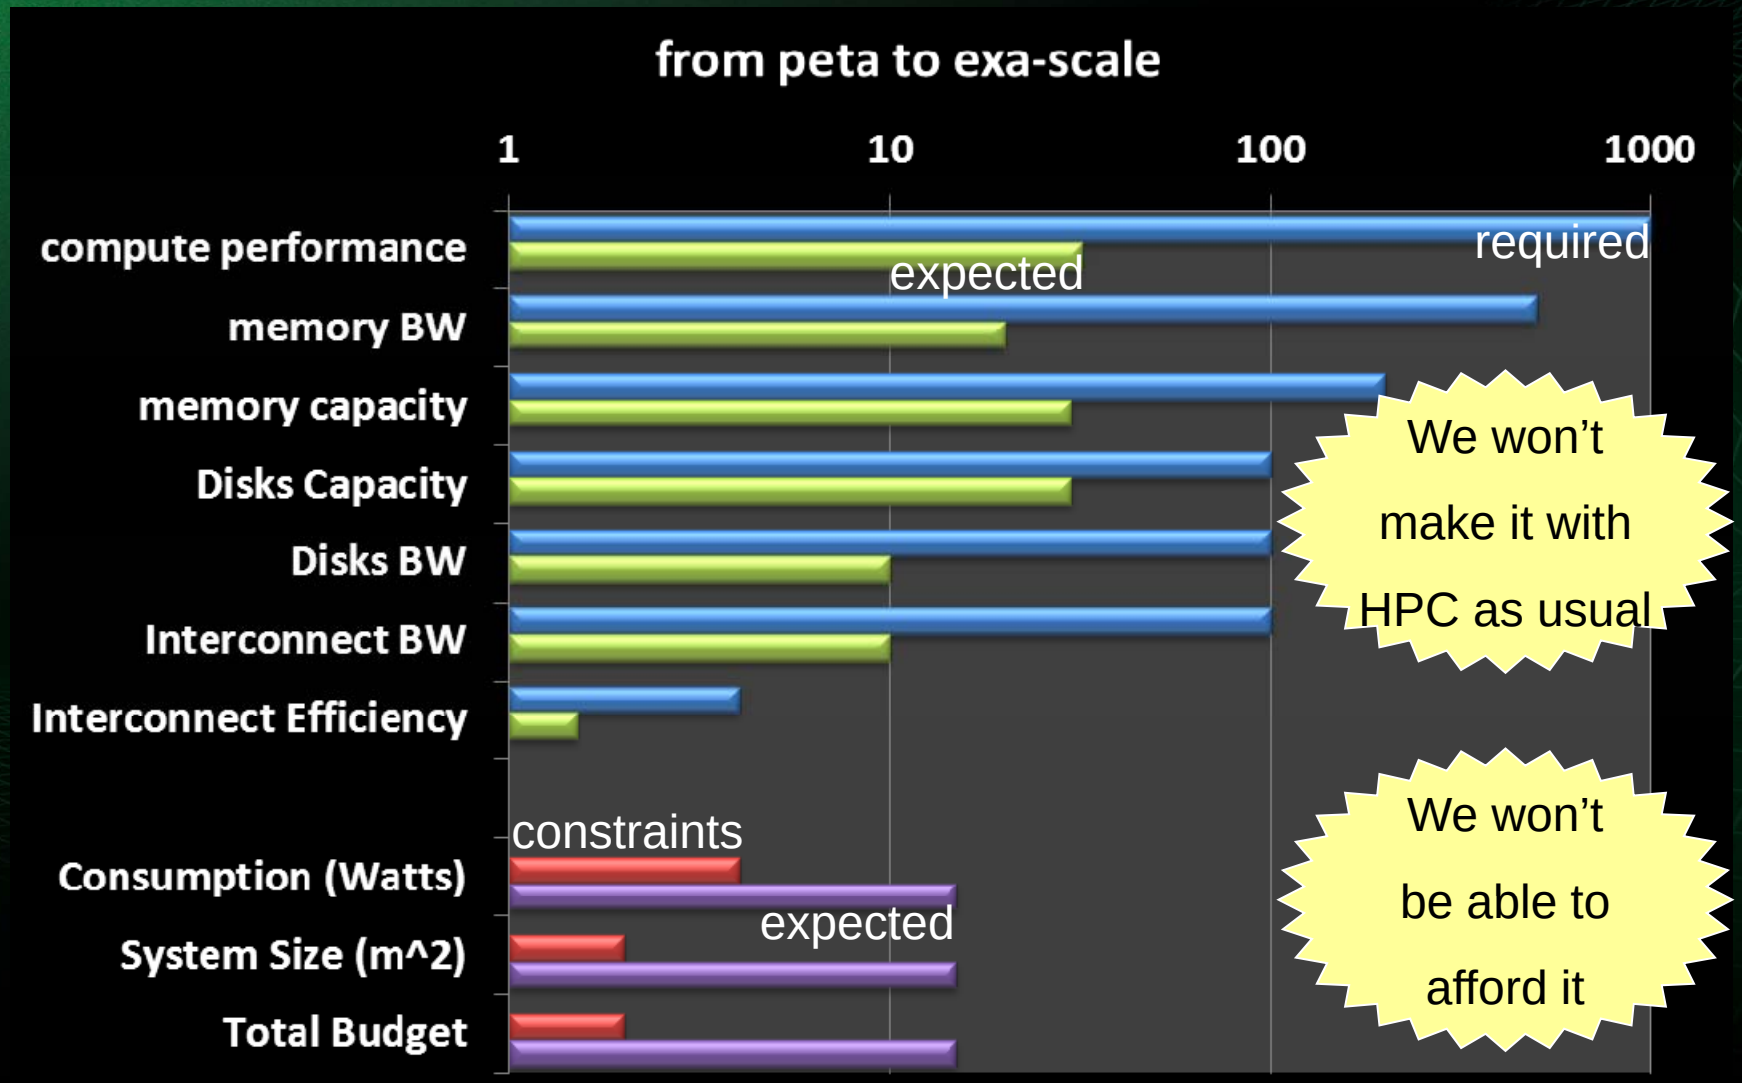

More than PetaFlops !

2011's CPU-GPU Hybrid Architecture

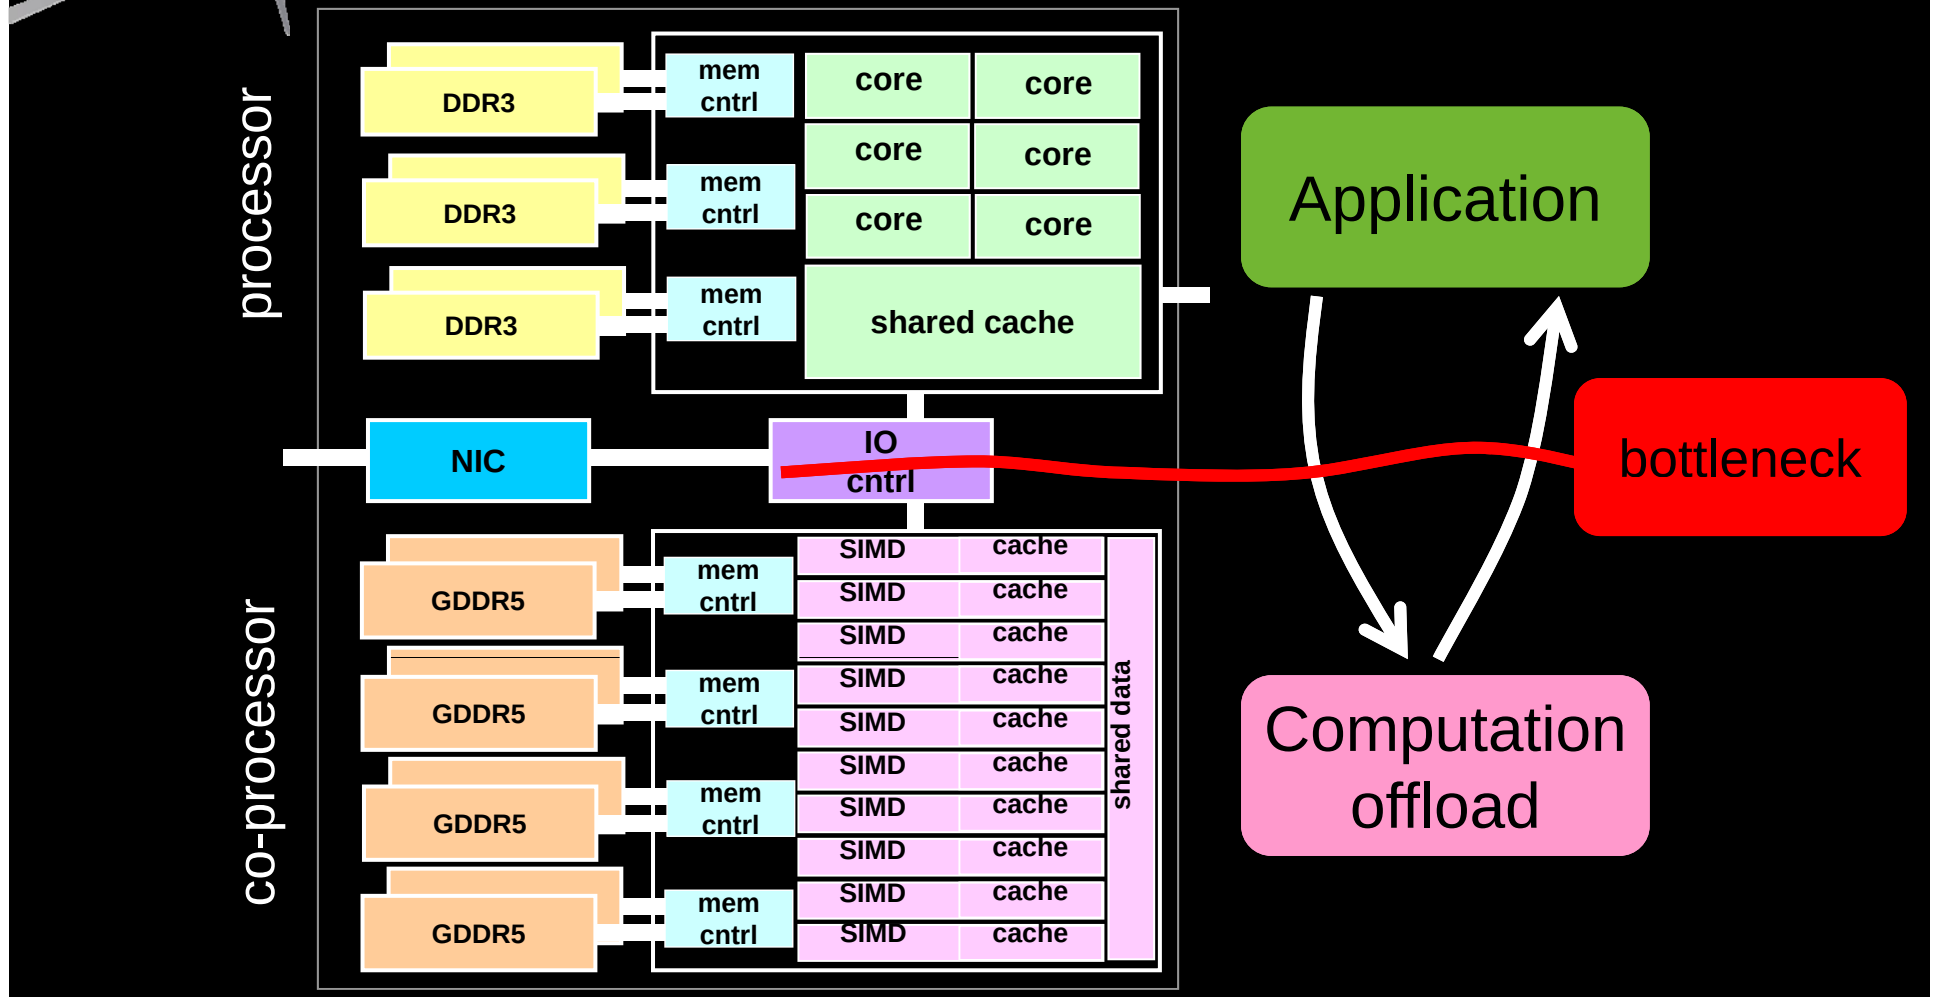

GPU Accélérators and Applications

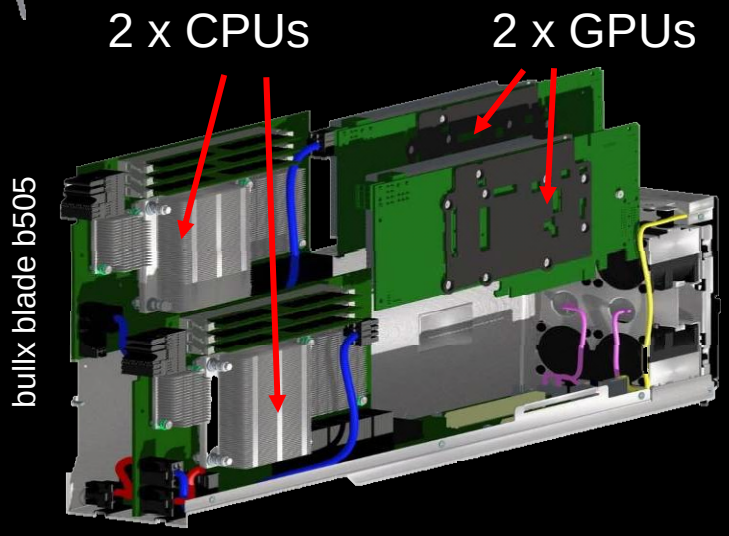

	GPU / CPU ratio
GFlops (DP)	**7**
Memory BW	**4.5**
consumption	**2**
Memory Size	**1 / 8**

Applications suited for GPUs :

- **Small source code (kernels)**
- **Limited datasets or very good locality**
- **(Single Precision)**
- **Little communications**
- **Cuda or OpenCL or HMPP**

- ***Graphics Rendering***
- ***Seismic Imaging***
- ***Molecular Dynamics, Astrophysics***
- ***Financial Simulations***
- ***Structure Analysis , Electromagnetism***
- ***Genomics***
- ***Weather / Climate / Oceanography***
- ***… and more…***

Hybrid CPU-GPU architecture evolutionS

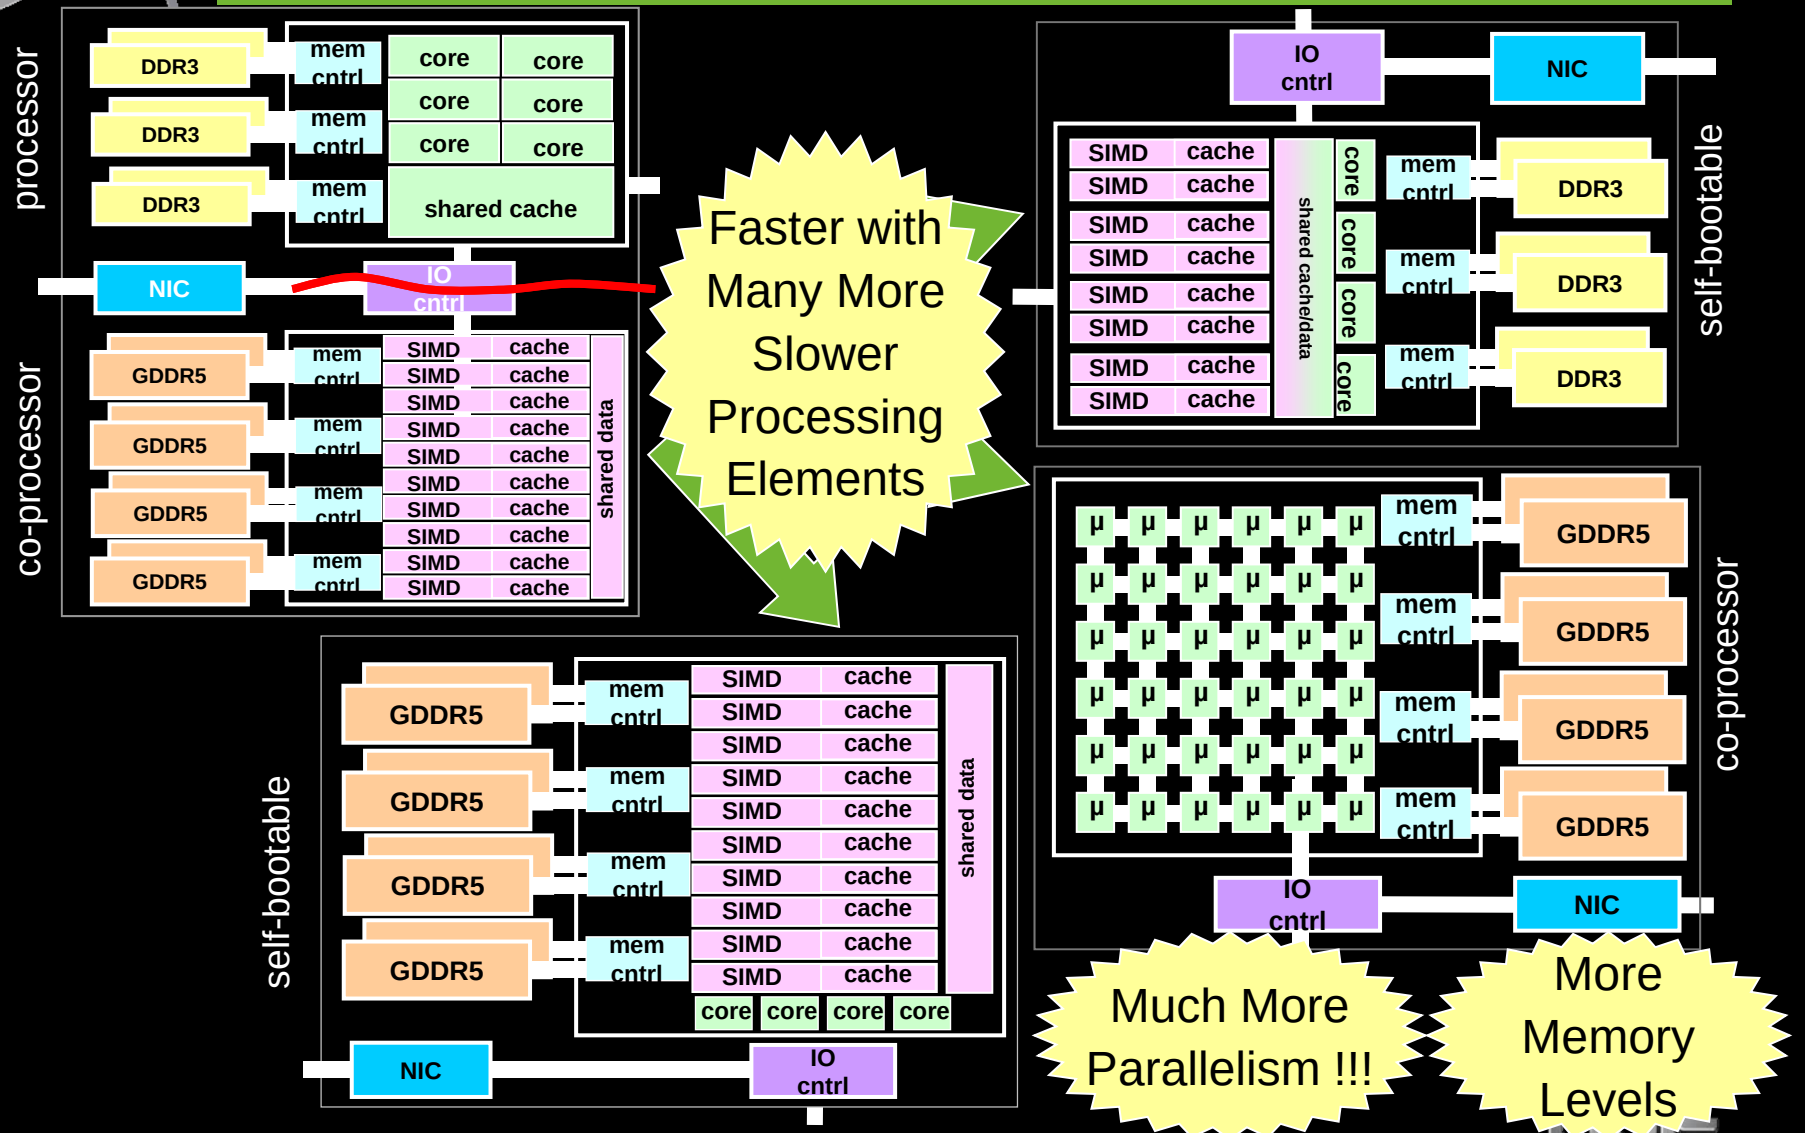

InfiniBand for Communications and IO

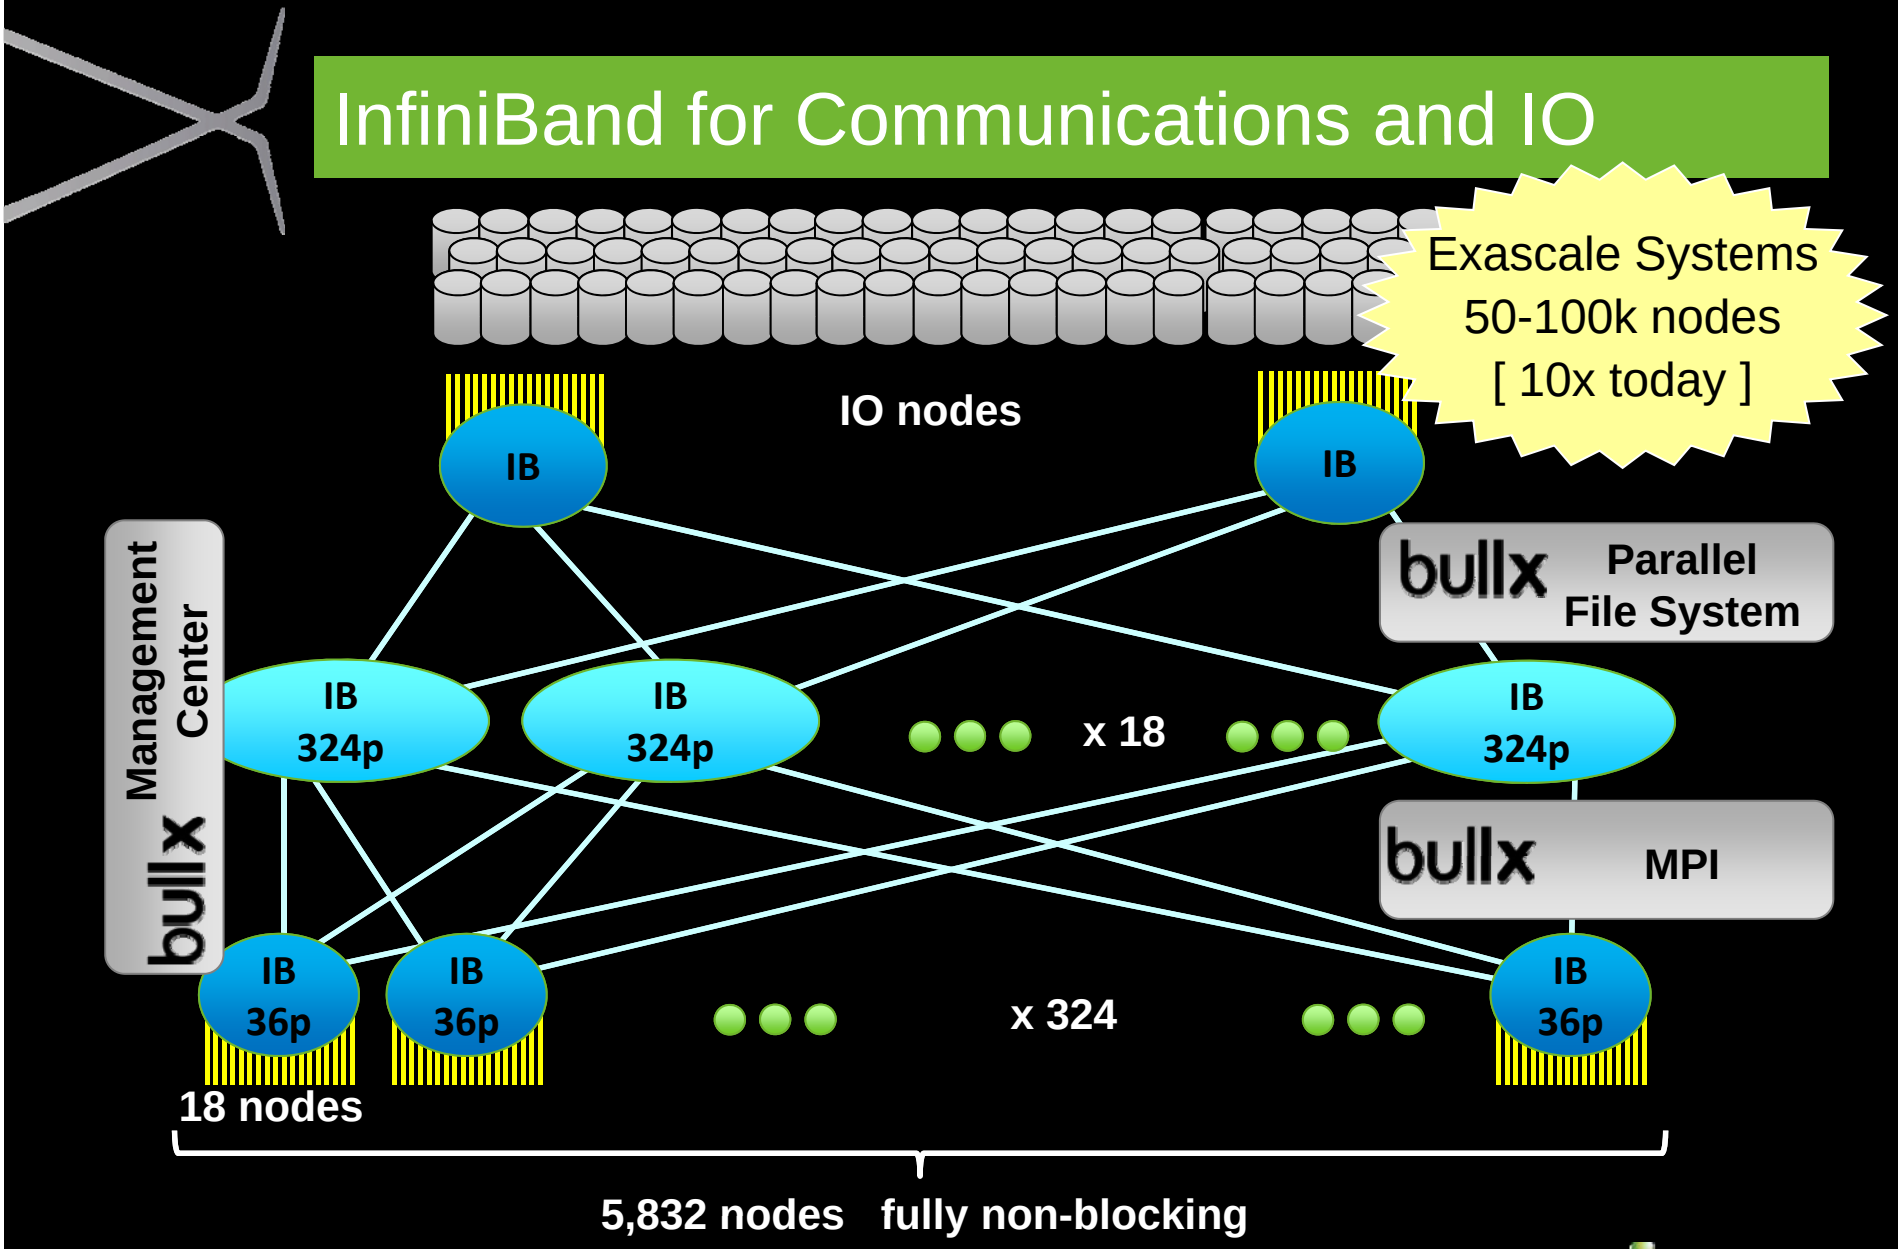

System reliability and Applications runtime

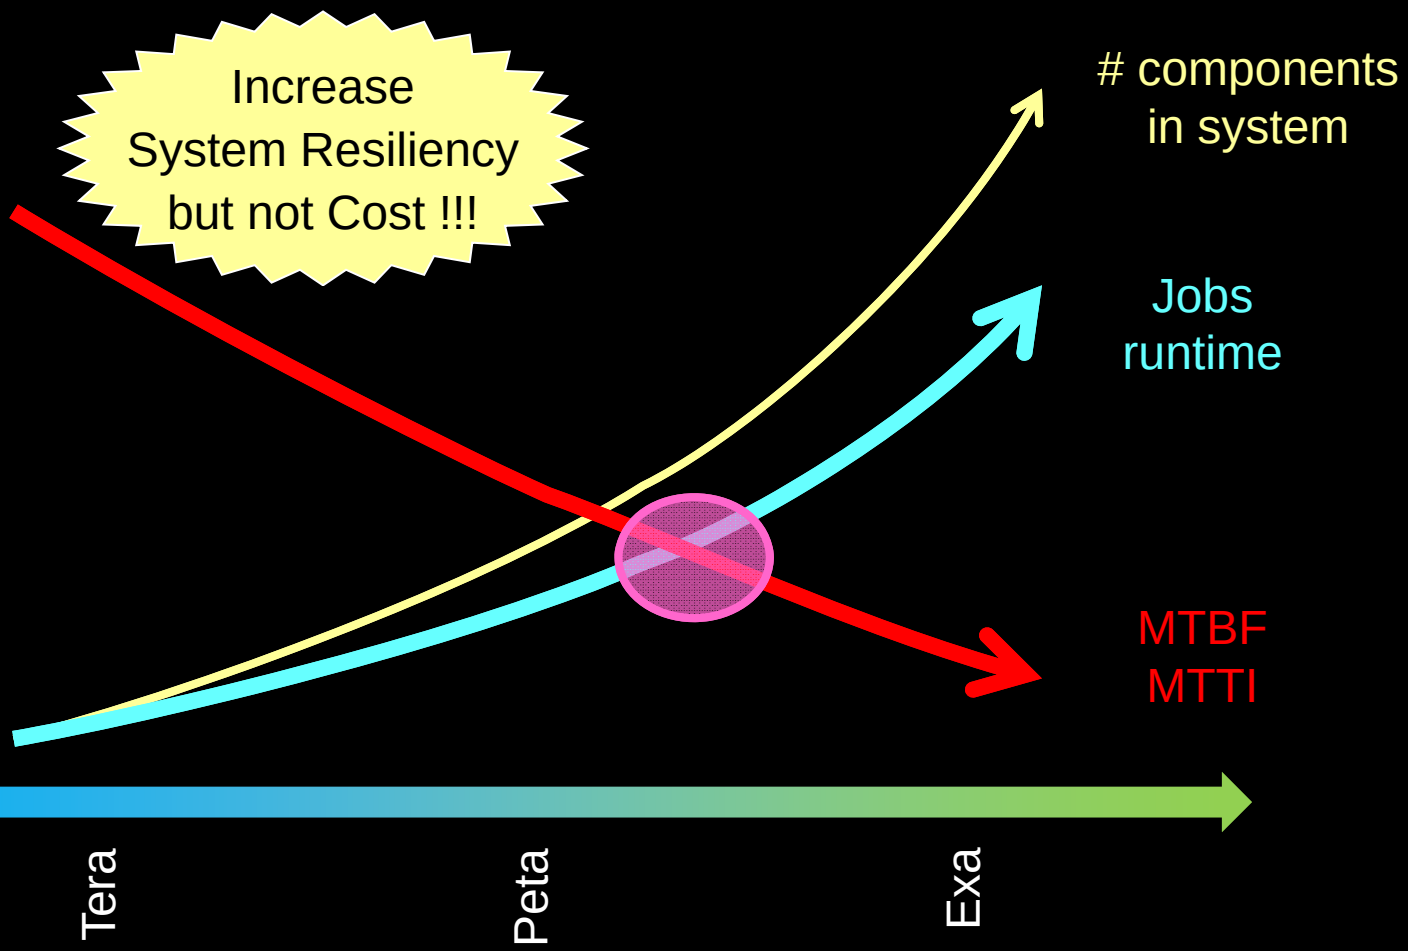

Cooling & Power Usage Effectiveness (PUE)

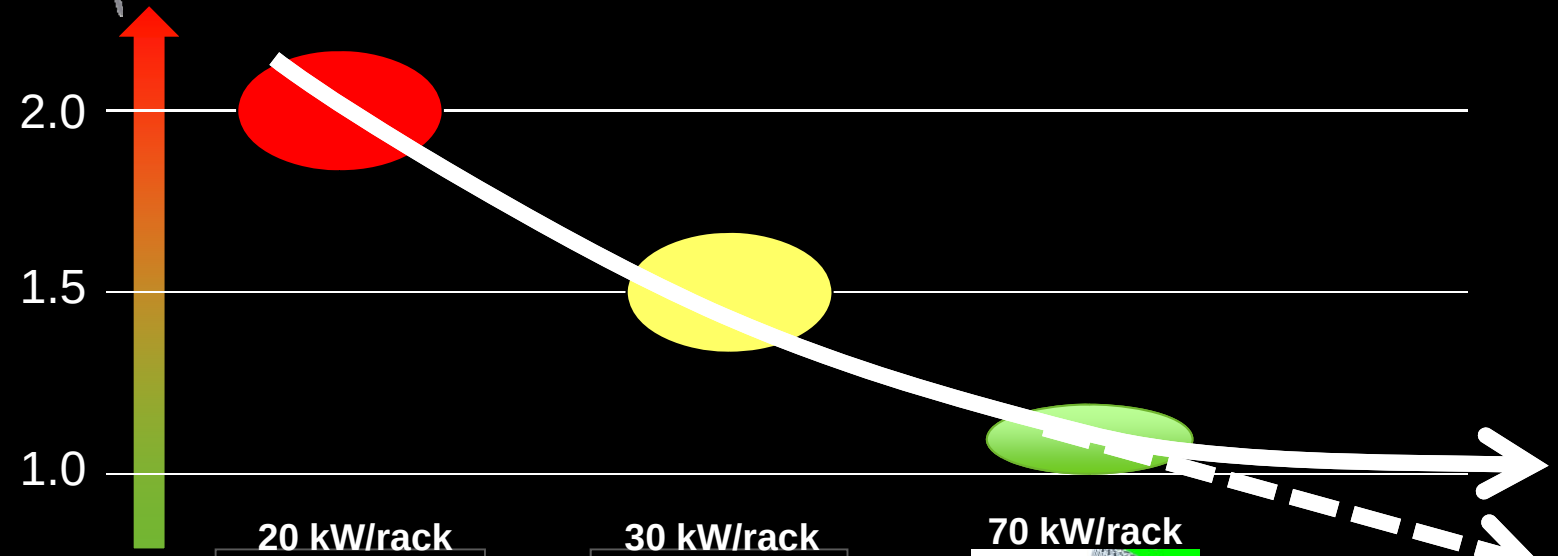

20 kW/rack

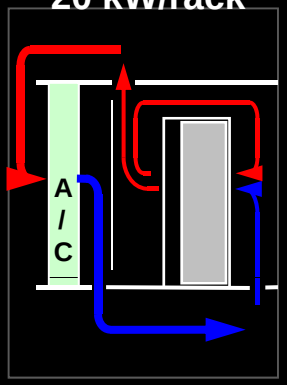

Air-cooled

30 kW/rack

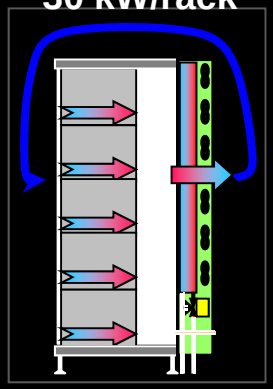

Water-cooled door

70 kW/rack

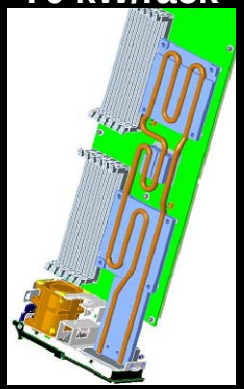

Direct-Liquid-cooling

Co-generation

HPC systems and future technologies:

- HPC Applications stress all aspects of a system (seldom flops)
- Processing Units (PUs) cross-over between CPUs & GPUs
  - CPUs ease of use + GPUs performance and power efficiency
- Balanced Design Required
  - Memory BW & latency, Interconnect, local IO
- Optimal data centers cooling
- Improved nodes energy efficiency
  - Accelerators, better integration (NIC in PU)
- Resilient but Affordable
- Software, Software, Software, …
  - Administration, Monitoring, Development, Runtime

More than Flops !

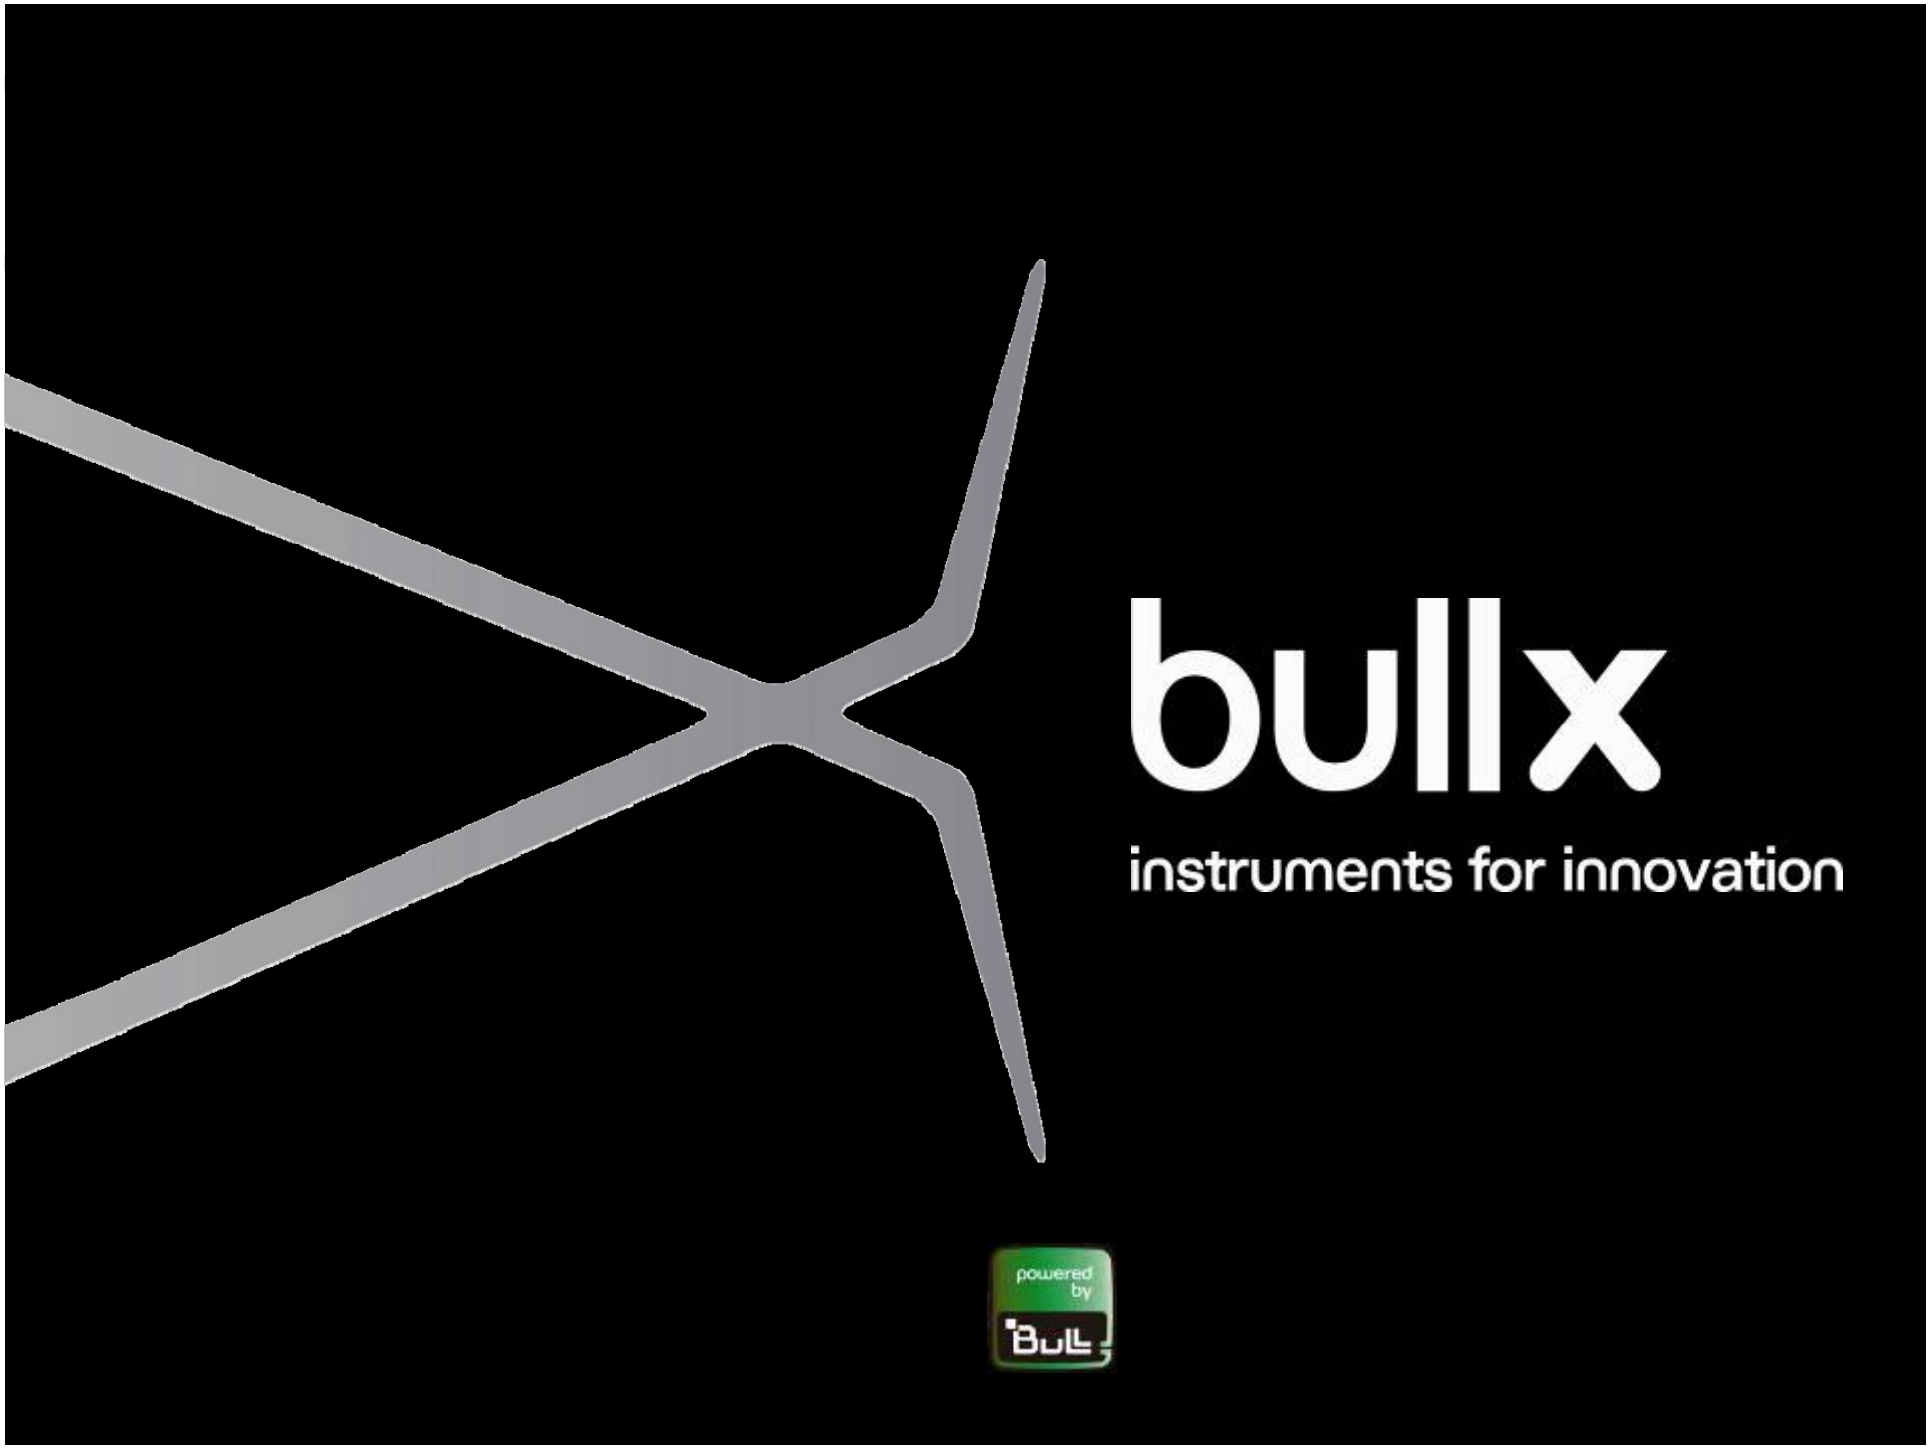
bullx
instruments for innovation
powered by
Bull