Accelerating Research
An off the shelf 10U Mini-Supercomputer

ANDREAS HAUSER
SCIENTIFIC COMPUTING EXPERT
JUNE 2013

GENE CENTER MUNICH AT LMU

- Established 1984 by Ernst Ludwig Winnacker
  - Now 15 Labs, over 300 people from over 20 countries
  - Huge growth under current director Patrick Cramer
- Research focus areas
  - Structural Biochemistry, Genome Maintenance, Regulation of Gene Expression
- High-profile publications
  - Science, Nature, Cell
- Research Techniques
  - Cryo-Electron Microscopy, X-Ray Crystallography, Next Generation Sequencing, Computational Biology

DIVERSE COMPUTATIONAL NEEDS

- High End Instruments
  - Cryo-EM: FEI Titan Krios (4M€)
  - Sequencing: Illumina GA II
  - High Throughput Robot: Beckman Coulter
  - Mass Spec: Thermo Scientific Orbitrap Elite
- 3D Reconstruction from Cryo-EM
  - 10k Raw Micrograph: 4096x4096x16 = 32MB
  - Up to 1M ~256x256x32 2D particle images
  - Resolution of 3D mass density down to 4Å
- Software Packages
  - SPIDER supports MPI (Frank, 1978)
  - Relion needs MPI (Scheres, 2012)

3D Ribosome Model

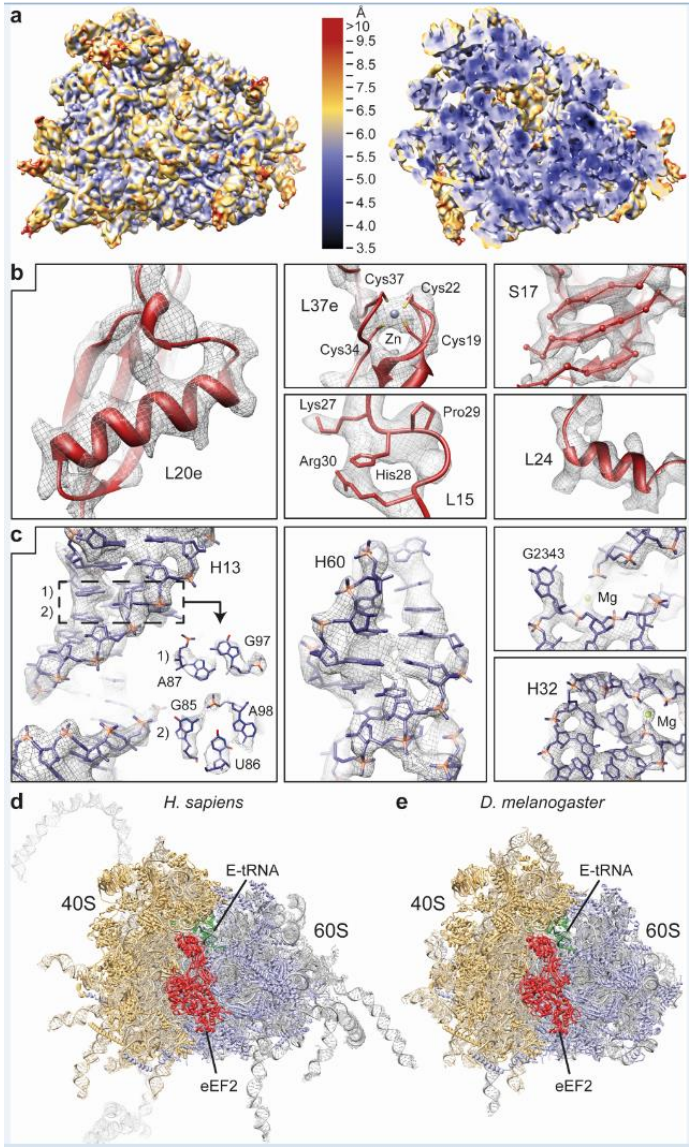

"Structures of the human and Drosophila 80S ribosome."
Anger AM, Armache JP et al., Nature, 2013

INTEGRATED SUPERCOMPUTERS ARE GREAT BUT …

… can we have that in half a Rack within budget?

“OFF THE SHELF” INTEGRATED MINI-SUPERCOMPUTER: AMD’S SEAMICRO

- Standard requirements
  - 4-6 normal power cables
  - No special cooling
  - 128 nodes productive within 48h
- High computational density at low power
  - 3.5KW per 10U
  - 3 SMs + fileservers in a Rack < 12KW
  - 2048 cores per Rack possible
- Near linear MPI scaling for our needs
  - 8 Gbit/s per low power CPU
  - Amazing bandwidth per Instruction
- Significant reduction in strain on budget
  - Total Cost of Ownership
  - Entry price
  - Personnel
  - FLOPs per Watt
  - Interconnect

AMD SeaMicro SM-15000

"Given that the OCP disaggregated server concept is not even in the market yet, AMD deserves respect for their forward thinking.“

Forbes, Apr 2013

A SEAMICRO NODE IS BASICALLY JUST A PCI-E CARD

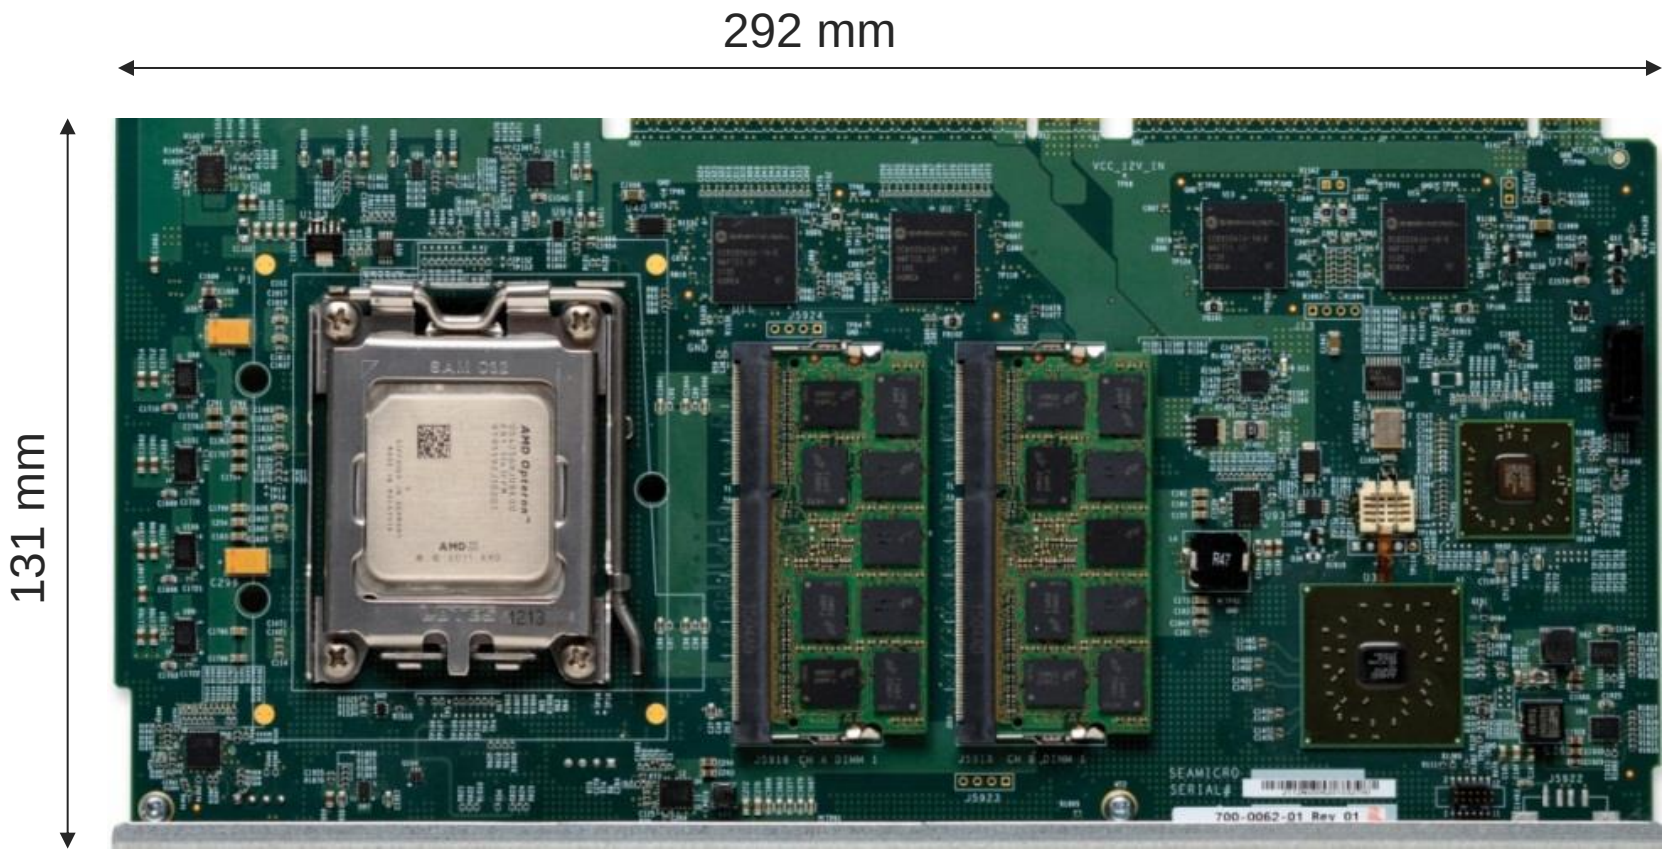

COMPARISON TO A BLADE SERVER

Option 1

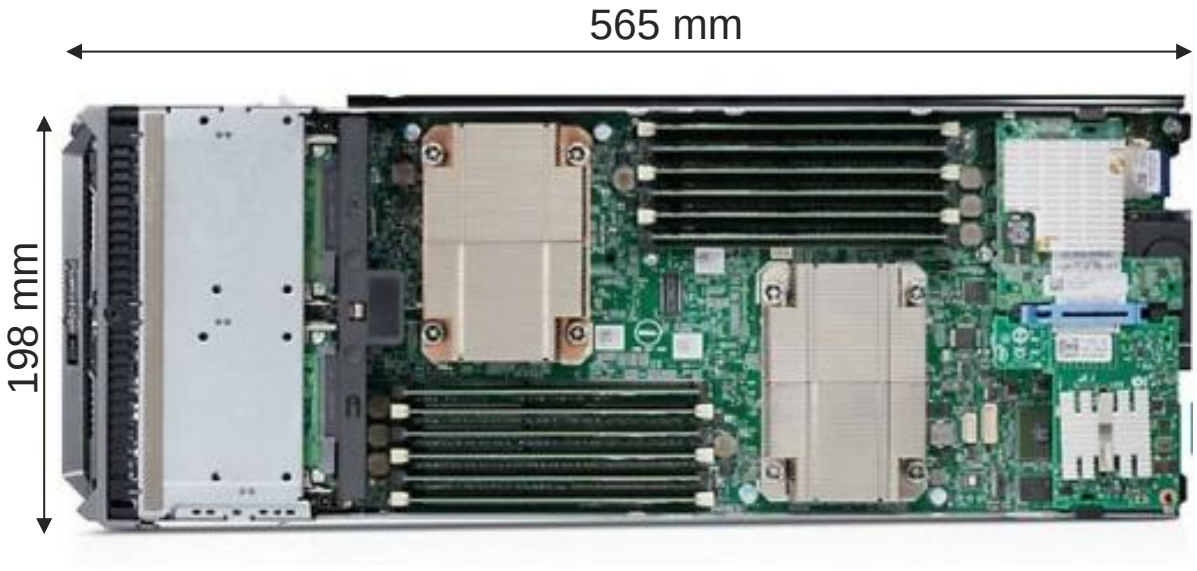

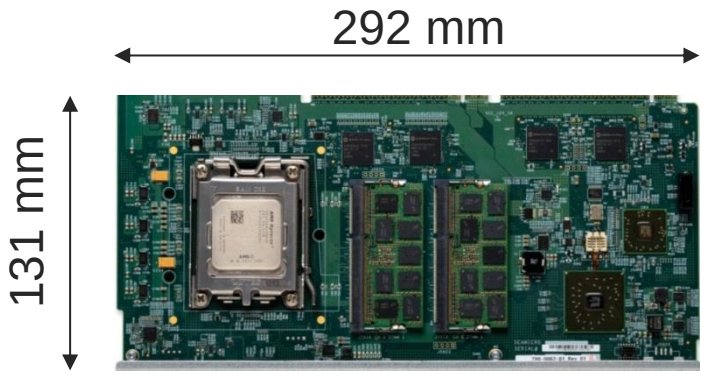

COMPARISON TO A BLADE SERVER

Option 2

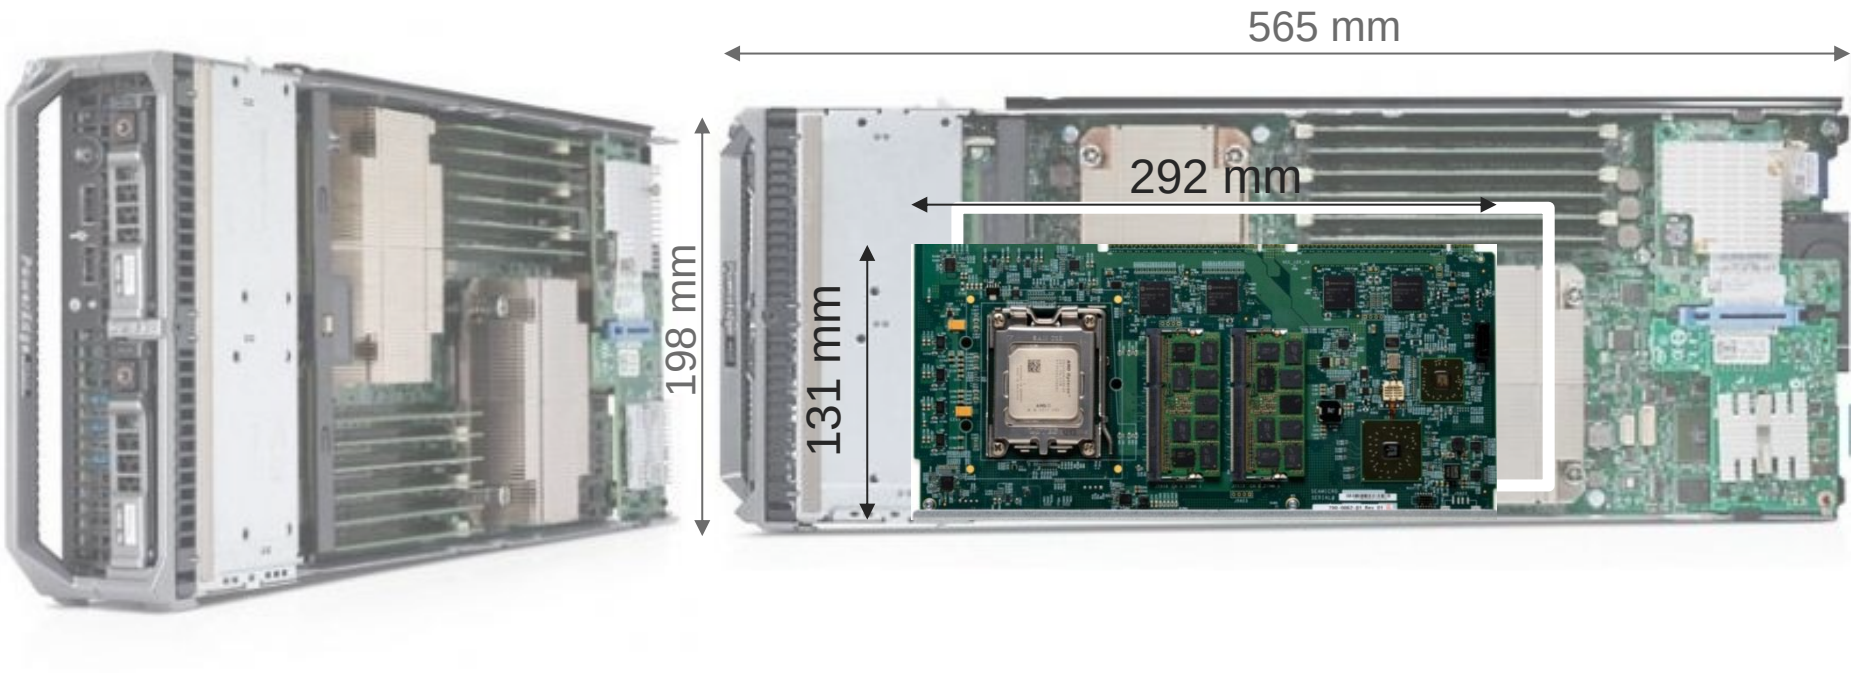

COMPARISON OF SM15000 TO A BLADE CENTER

	AMD's SeaMicro SM15000-Opteron	AMD's SeaMicro SM15000-Xeon	Blade Center
			

Comparison	SM15000-OP AMD Opteron 4365 EE (Measured)	SM15000-XN Intel Xeon E3-1265Lv2 (Measured)	Blade Center Intel Xeon E5-2470 (estimated)
Nodes	64	64	16
Number of Cores	512	256	256
GHz	2000	2500	2300
Network GBit/s per core	1.0	2.0	0.6
HPL Performance (TFLOP/s) Efficiency (estimated %)	2867 70%	3584 70%	3300 70%
Power HPL (kW)	5.82	4.54	6.10
MFLOPS/W	639	785	548
CPU Price per Core	?€	65€	150€

Benchmark measurements and estimates by AMD/Seamicro

COMPARED TO A STANDARD BEOWULF CLUSTER

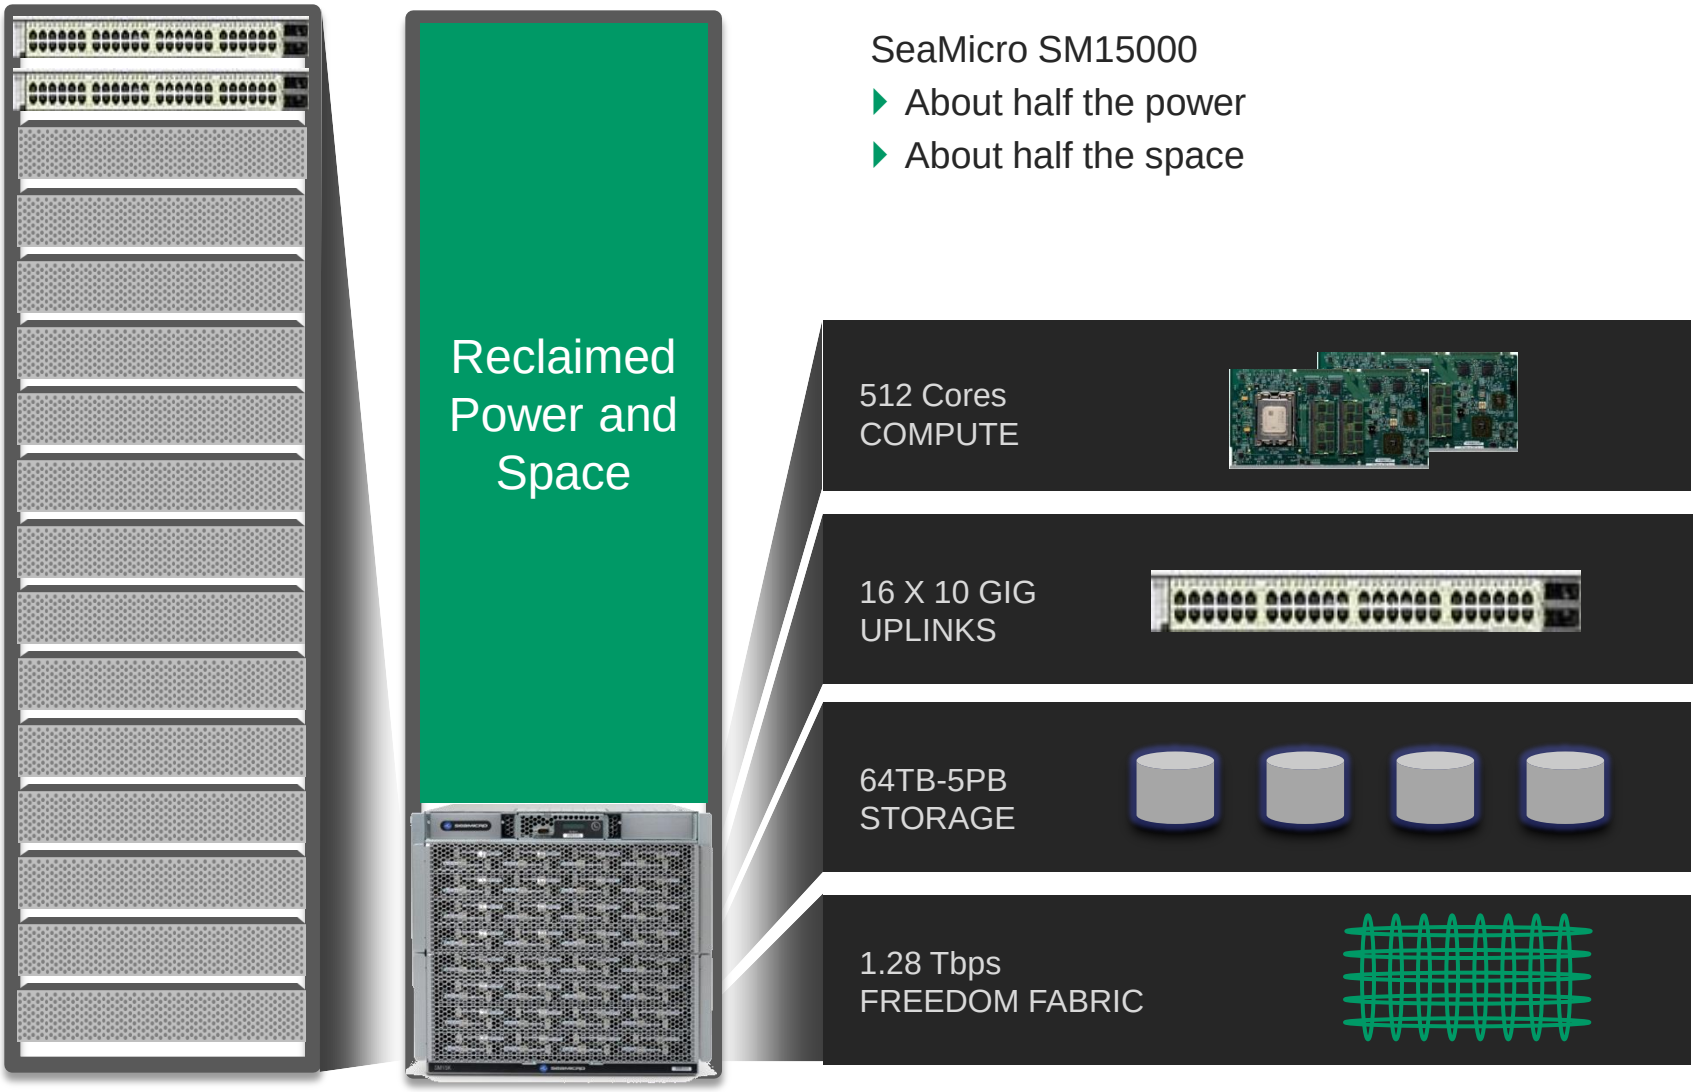

SeaMicro SM15000

- About half the power
- About half the space

FABRIC INTERCONNECT PROVIDES ANY SERVER TO ANY DISK DRIVE CONNECTIVITY AT 10 GBPS

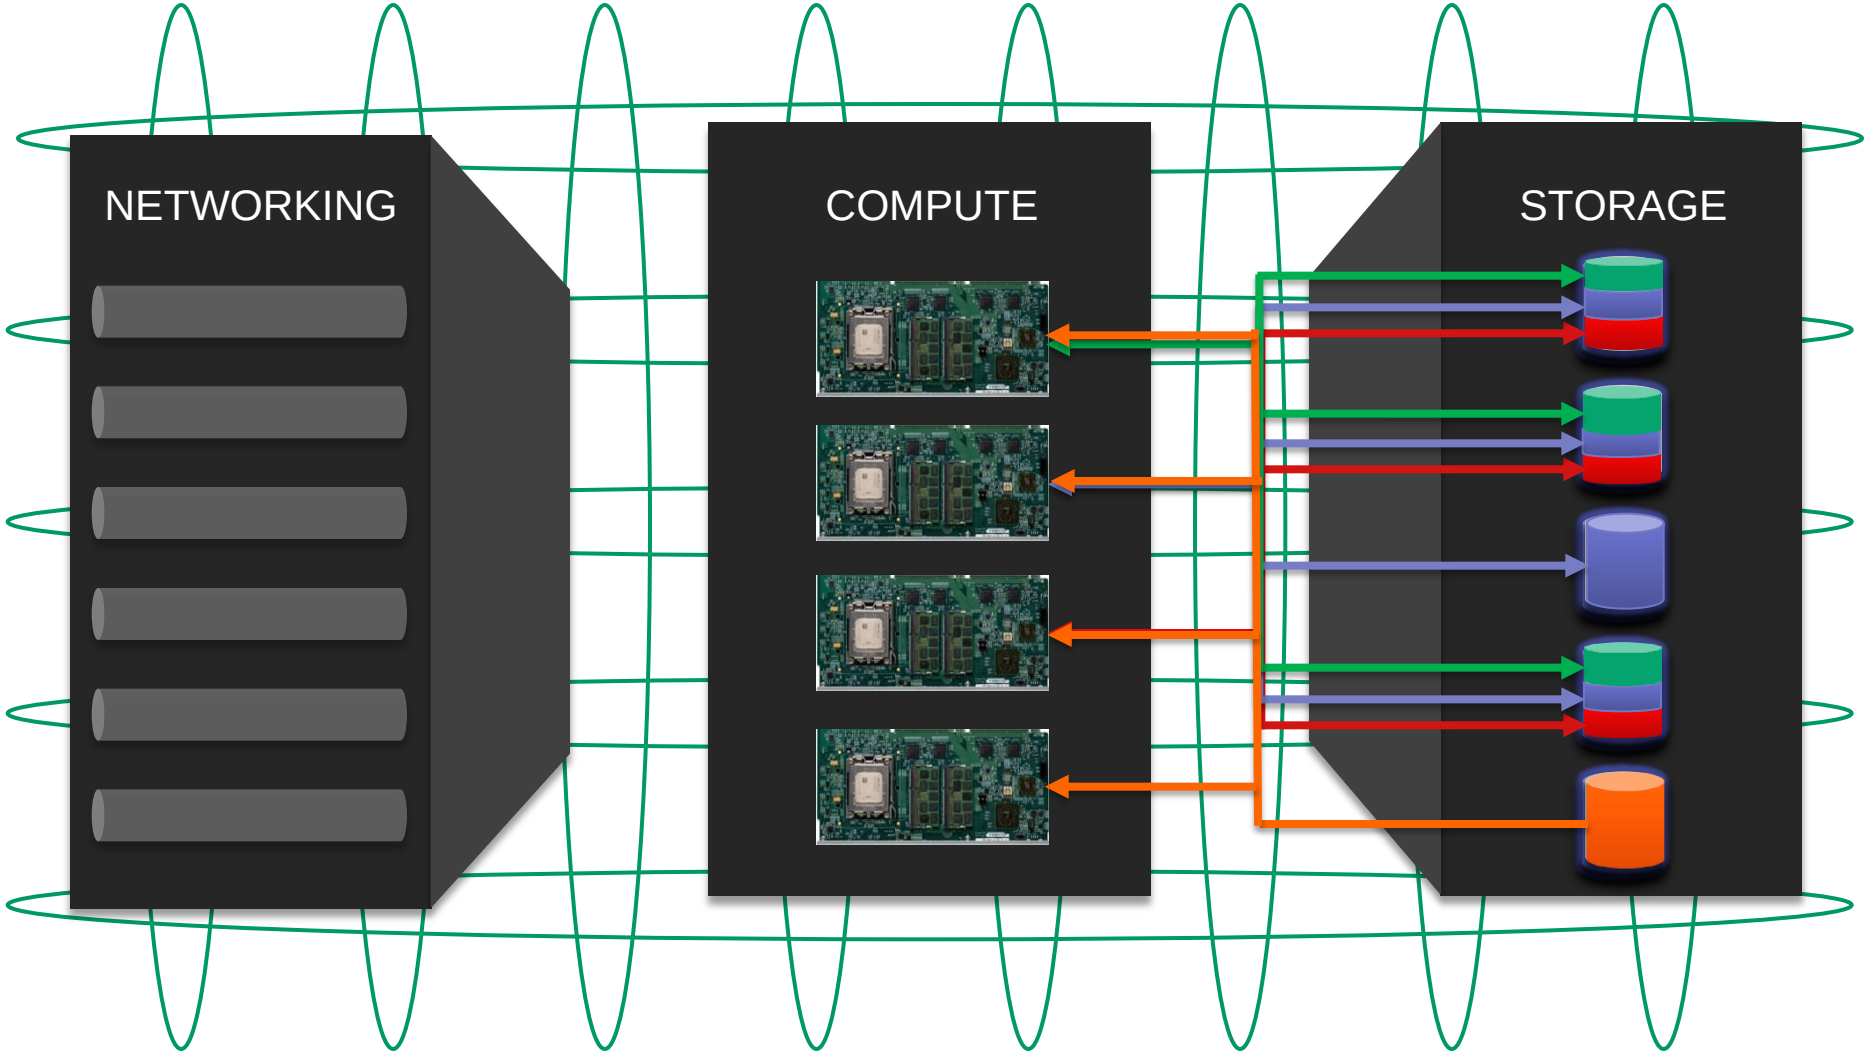

NETWORK I/O IS SHARED TO MINIMIZE NIC COSTS AND AVOID UNDERUTILIZATION OR CONGESTION—NO INFINIBAND

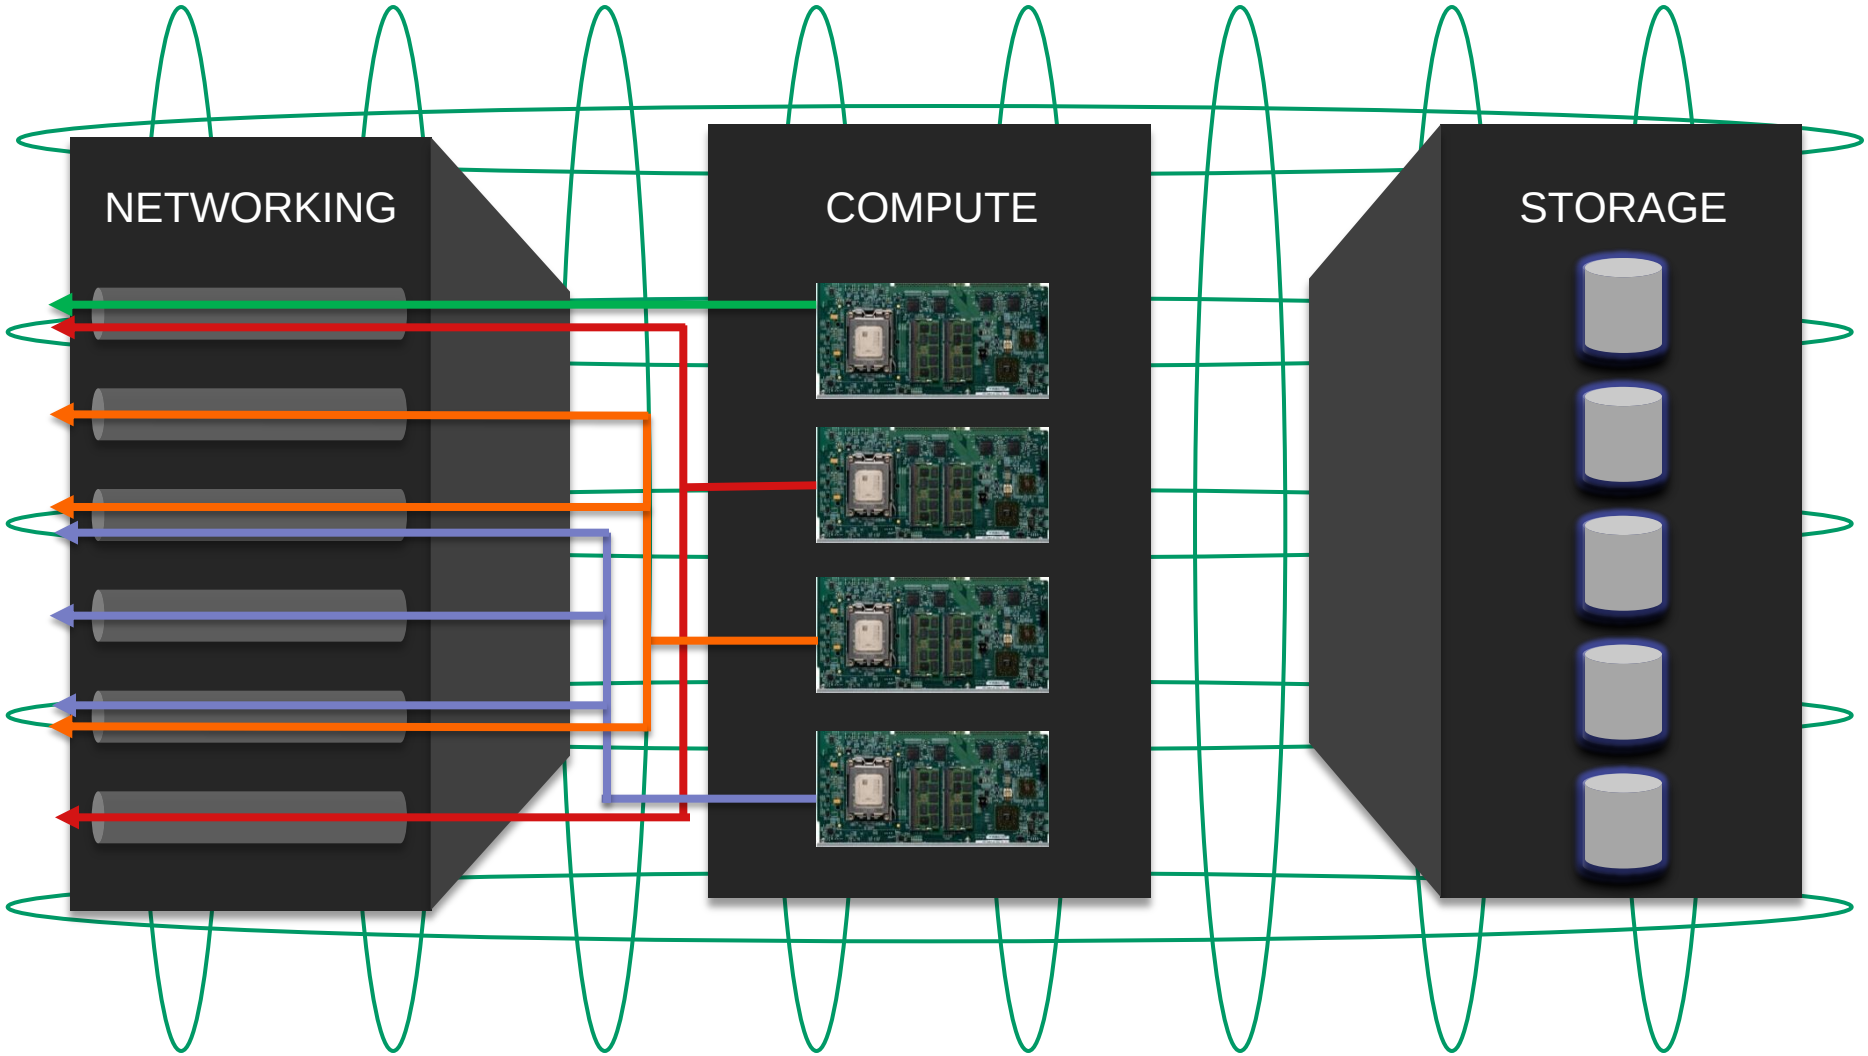

MINI-SUPERCOMPUTER IN A BOX

64 HDDs/SDDs
- Share drives across all servers
- Assign one server to one or more drives as needed
- In service upgrades as needed

64 Industry standard x86 servers
- Opteron, Xeon, Atom
- Energy efficient processor
- 20 Gbps per socket, 16X traditional servers

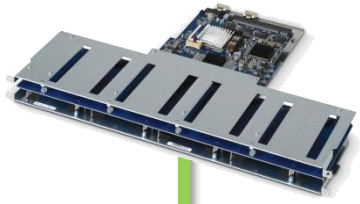

5 Petabytes Fabric Storage
- Extends supercompute fabric to external storage
- 1,344 LFF/SFF drives
- Map to any CPU—same as internal drives

160 Gbps Network I/O
- Share network I/O across all servers
- Eliminate TOR switch
- Minimize cabling
- In service upgrades as needed

Thanks

- Prof. Dr. Roland Beckmann, GCM
- Dr. Andreas Anger, GCM
- Young-Sae Song, AMD/Seamicro
- Joshua Mora, AMD/Seamicro

Questions