

The Bottom Line of Convergence

Martin Grund

Hasso Plattner Institute

State of the Union

HPC

- Asynchronous, job-based “query” execution
- Large scale
- Specialized hardware
- Compute bound (mostly)

Analytics (OLAP)

- Data-based
- I/O dominant (main memory or disk)
- Eventual Consistent

Data Management (OLTP)

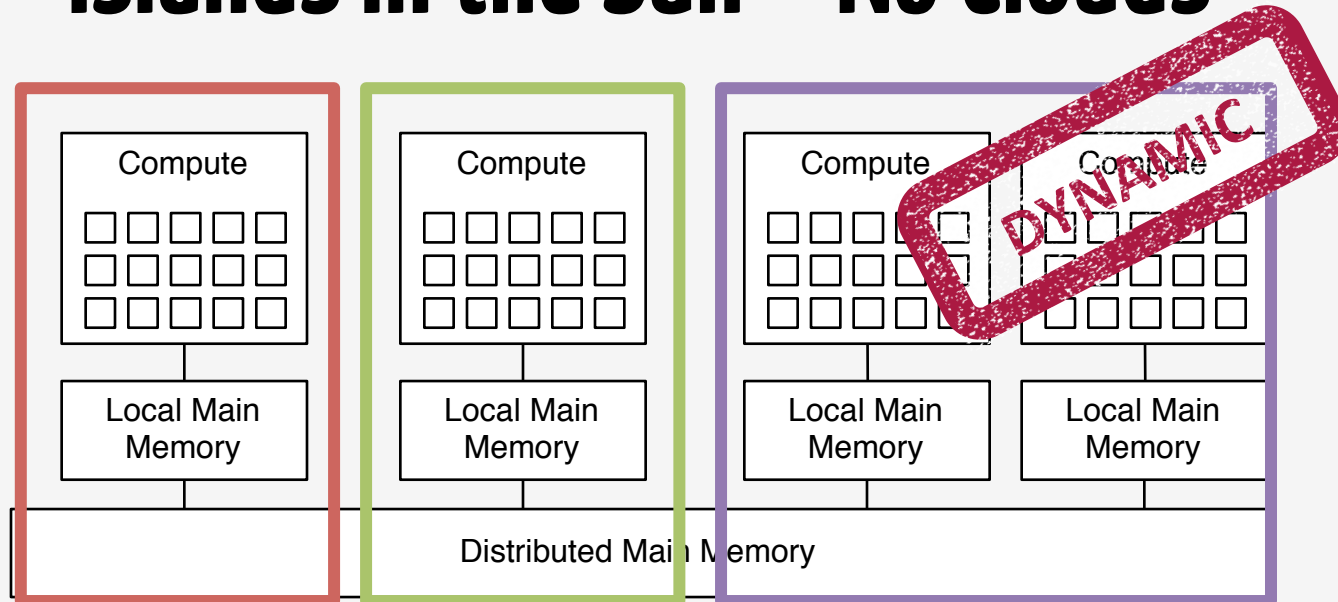
- Monolithic, scalability bound by network latency
- Transactional
- Strongly Consistent

Con...ver...ge..nce ??

State of the Art

- Huge amounts of main memory – 2TB, basically we have more memory than compute resources
- Scalability foe #1 vanishes – Infiniband reduces bandwidth and latency bottleneck to below one order of magnitude
 - Intel Xeon E5 4650 – Bandwidth ~409 Gbits/s, Latency 0.1 μ s
 - Infiniband – Bandwidth ~ 300 Gbits/s, Latency 1-2 μ s
- Many core (even more with GGPUs)
- The Intel way using x86 (why not?)

Islands in the Sun – No Clouds



OLAP ↔ **HPC** ↔ **OLTP**

Research:

- What is locality?
- What is storage?
- What is the API?
- Execution model?