


Big Data/High Performance Computing Perspectives for Unmanned Systems

Dan Andresen, Ph.D.

Director, Institute for Computational Research in
Engineering and Science

dan@k-state.edu






= \$\$\$\$

\$

=



**Bottlenecks are temporary.
Be opportunistic.**






**You will never know everything.
Embrace inconsistency.**



~1TB, 70g, \$600
Use 2.

Keep your data.



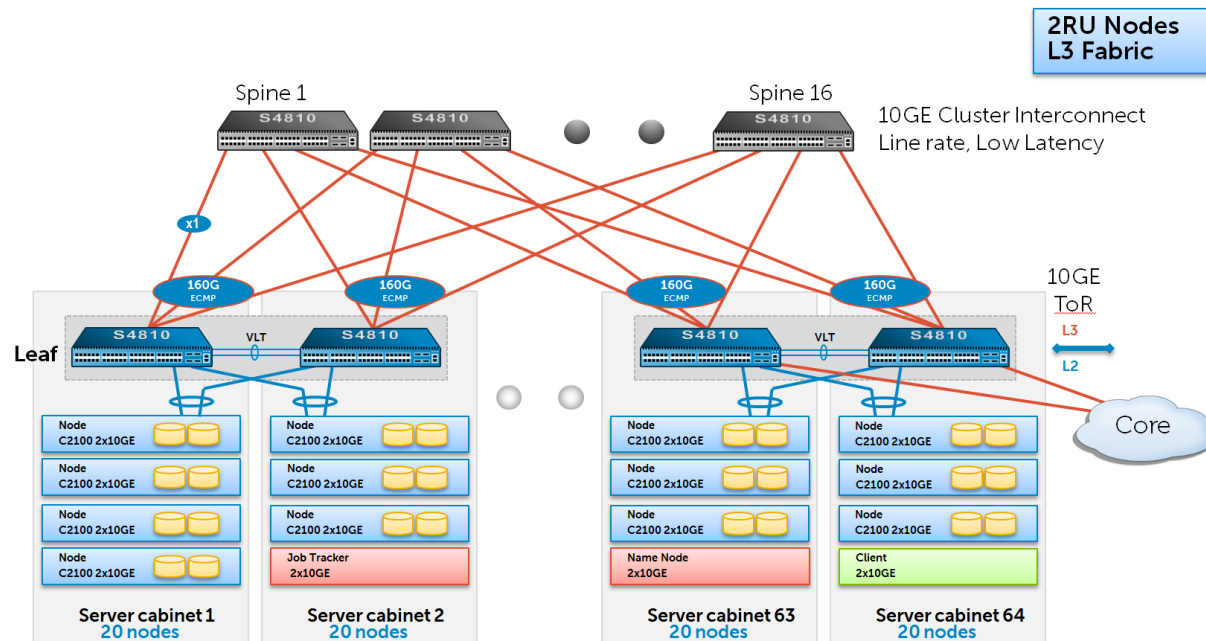
Pipeline your workflow.



<http://www.marketingoops.com/news/tech-update/ipad-lego/>

“[G]et rid of monolithic software and build bite-size software.”

10GE Hadoop Cluster – Scaled to 1280 nodes

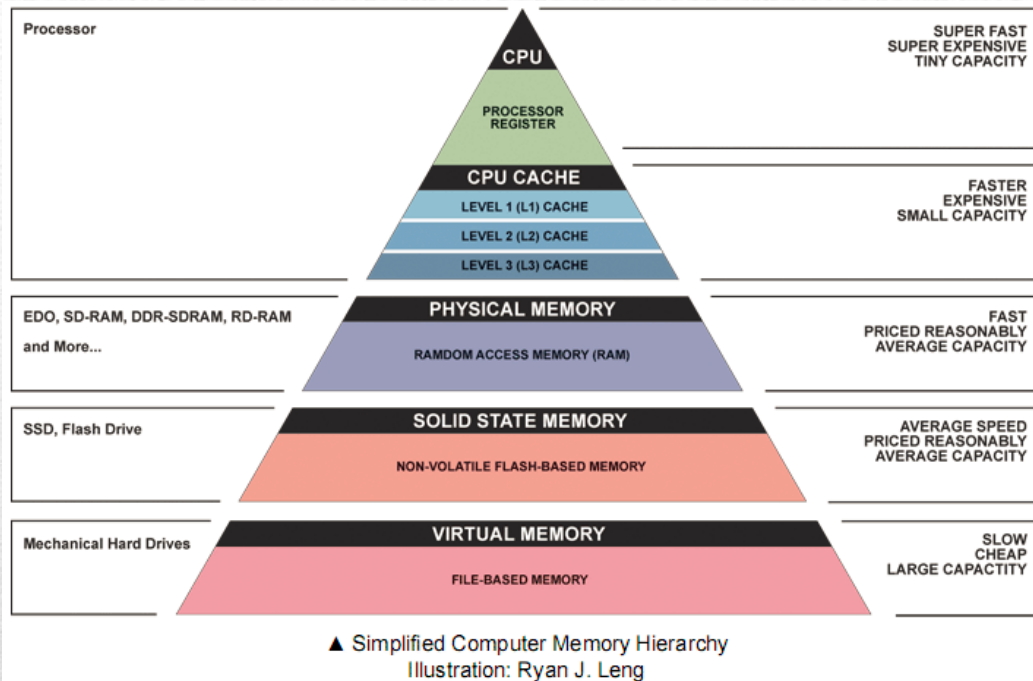


- 2RU Nodes 2 x 10GE (C2100)
- 20 Nodes per rack
- 64 racks, **1280 nodes**
- 2.5:1 oversubscription @ ToR
- Expand as needed by adding Spine switches
- Leaf QSFP+ optics, Spine SFP+ optics
- Leaf QSFP optical breakout cables
- 150m Leaf to Spine

BRAD HEDLUND.com



Scalability beats performance.

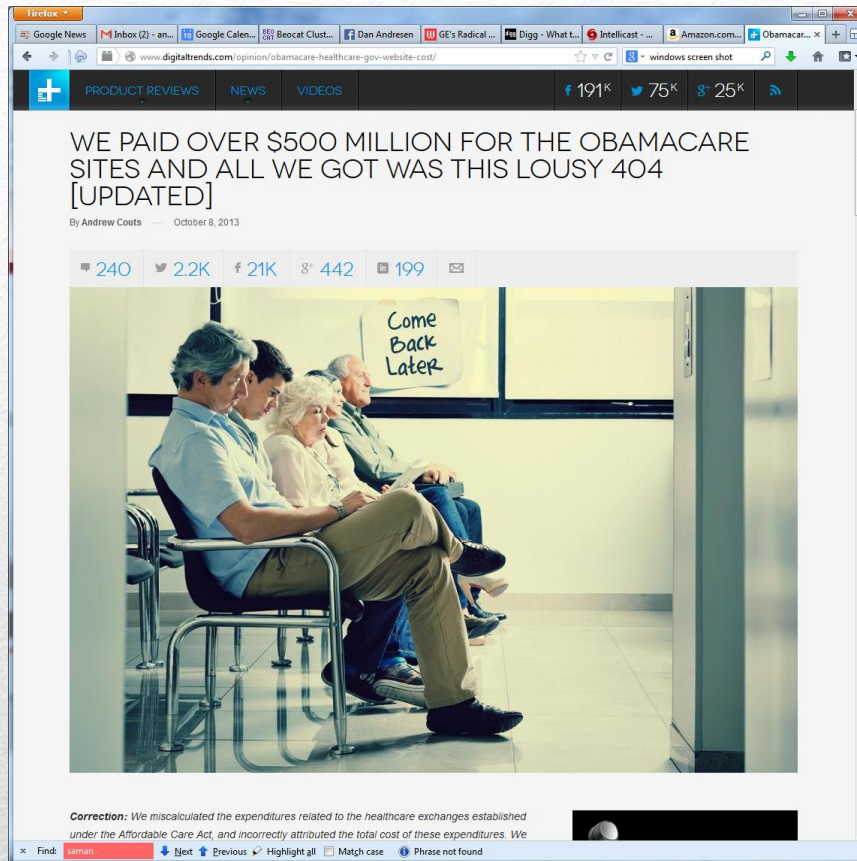


<http://cs230.wikispaces.com/Architecture+Part+2+Notes>

Big Idea: The memory hierarchy creates a large pool of storage that costs as much as the cheap storage near the bottom, but that serves data to programs at the rate of the fast storage near the top.

Randal E. Bryant and David R. O'Hallaron, CMU

The Memory Hierarchy is there for a reason. Play nice.



**When in doubt, bring in a
Computer Scientist.**