

# *HPC Interconnect Power Measurement*

Presented to: EEHPCWG Workshop @ SC14

November 17, 2014

Robin Goldstone  
HPC Solutions Architect, LLNL

 Lawrence Livermore  
National Laboratory



LLNL-PRES-664358

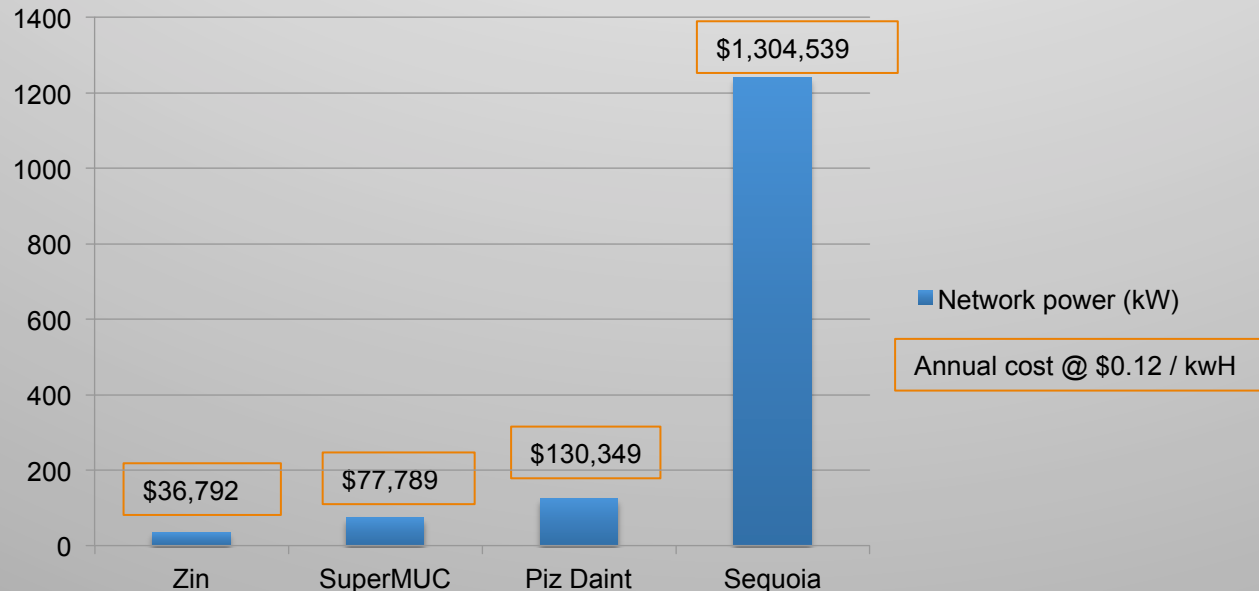
This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

# HPC Interconnect Power

- Measurement of network power was not required by original Green500 rules.
- Network power is not insignificant

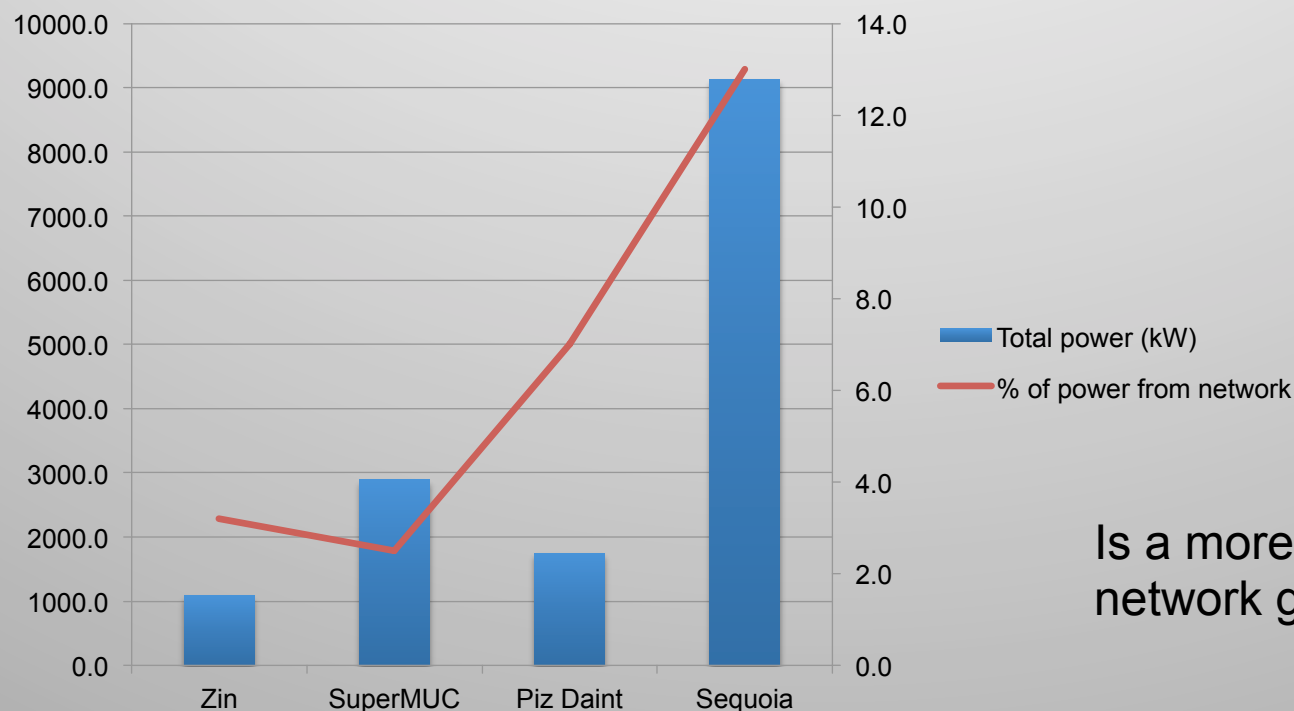
...especially if you have to pay for it!

## LINPACK Network power



# Network Power vs. Total Power

Now that we can measure network power, what can we glean from the data?



Is a more power-hungry network good or bad?

# Food for Thought

- What is the right metric to judge the energy efficiency of the interconnect?
  - Bandwidth: GB/s/watt
  - Message injection rate: messages/sec/watt
  - Other ?
- What benchmark(s) capture the desired metric?
- How much does network power increase under a communication-heavy workload?
  - Optics use power whether transmitting data or not
  - Some interconnects offer power-saving modes (Mellanox IB ~8% switch power reduction possible)
  - If network power is mostly static, do we even need a workload to measure it?

