Characterizing Full-system Network Performance and Congestion Management Capabilities with Improved Network Benchmarks

Pete Mendygral, Nathan Wichmann,
Duncan Roweth, Krishna Kandalla, & Kim McMahon

pjm@cray.com

CONTENT

- Purpose
  - HPC workloads generate dynamic and complex communication patterns / interactions that affect
    - Full system performance and throughput
    - User experience
  - The HPC community needs improved methods and benchmarks for
    - Characterizing network performance
    - Measuring impacts from network contention and congestion
- Improved network test infrastructure
- Infrastructure for measuring impacts from congestion
- Results
  - Measurements of isolated performance and performance with congestion using the new infrastructure
- Summary
- Q&A

HPC COMMUNICATION WORKLOADS

- Communication in HPC workloads are dynamic and complex
  - User applications with a range of
    - Communication patterns
    - Intensity
    - Scale
  - System services (e.g., filesystems)

- Interactions within a workload induce contention for network resources (e.g., buffers)
  - Results in
    - Performance variability
    - Reduced system throughput

- Need tools to characterize HPC workload communication performance

NETWORK BENCHMARKS

“Tests like ping pong latency are like trying to understand your commute into NYC by driving the route alone at 4am.” – Steve Scott

Your Commute at 4:00 AM

MPI Ping-Pong Latency

Your Commute at rush hour

Halo Exchange Under Load

NEED FOR NEW BENCHMARK METHODS

- Commonly rely on MPI benchmarks for measuring network performance
  - E.g., OSU and Intel MPI Benchmarks
  - Suitable for measuring performance of MPI library components
  - Not sophisticated enough for characterizing HPC networks with a diverse workload

- HPC community needs standard benchmark methods for
  - Approximating complex communication patterns at scale
  - Measuring performance limiting metrics
  - Impact of contention or congestion

IMPROVED NETWORK TESTS

NETWORK BENCHMARK INFRASTRUCTURE

- Goal was to create a benchmark that
  - Could be a proxy for a range of important communication patterns
    - Latency, bandwidth, and synchronization sensitivity
  - Statically characterizes performance
    - Report average and 99% tail without significant bias from placement and network topology

- In collaboration with NERSC and ANL, developed benchmark with an infrastructure that simulates
  - Nearest neighbor exchanges with natural and random rings
    - Latency, bandwidth, and bandwidth with synchronization
  - FFT transposes with scalable all-to-all
  - Synchronization or small reductions (e.g., convergence test)

- Optionally reports additional metrics
  - E.g., 99.9% tail performance and full histogram data files

- Measures only off-node communication using sub-communicators

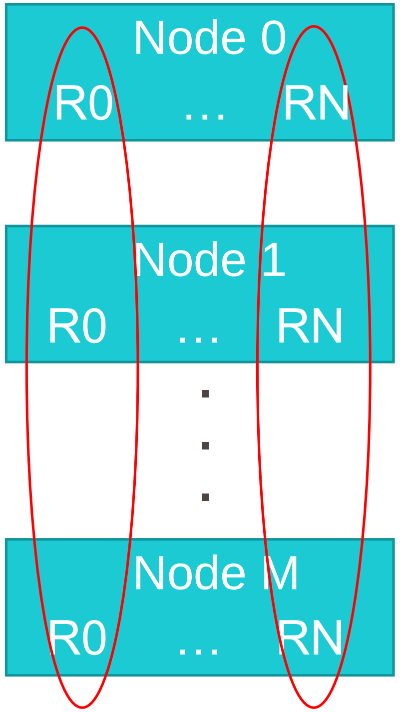

NATURAL RING COMMUNICATION PATTERN

- Proxy for 2D nearest neighbor pattern
  - 2nd dimension is on-node
- Natural ring has low connectivity
  - Each node communicates with only 2 other nodes
  - E.g., Node 0 communicates only with Nodes 1 and 4
- Bandwidth is limited by injection bandwidth
  - Opposed to global bandwidth
- Latency similar to ping-pong

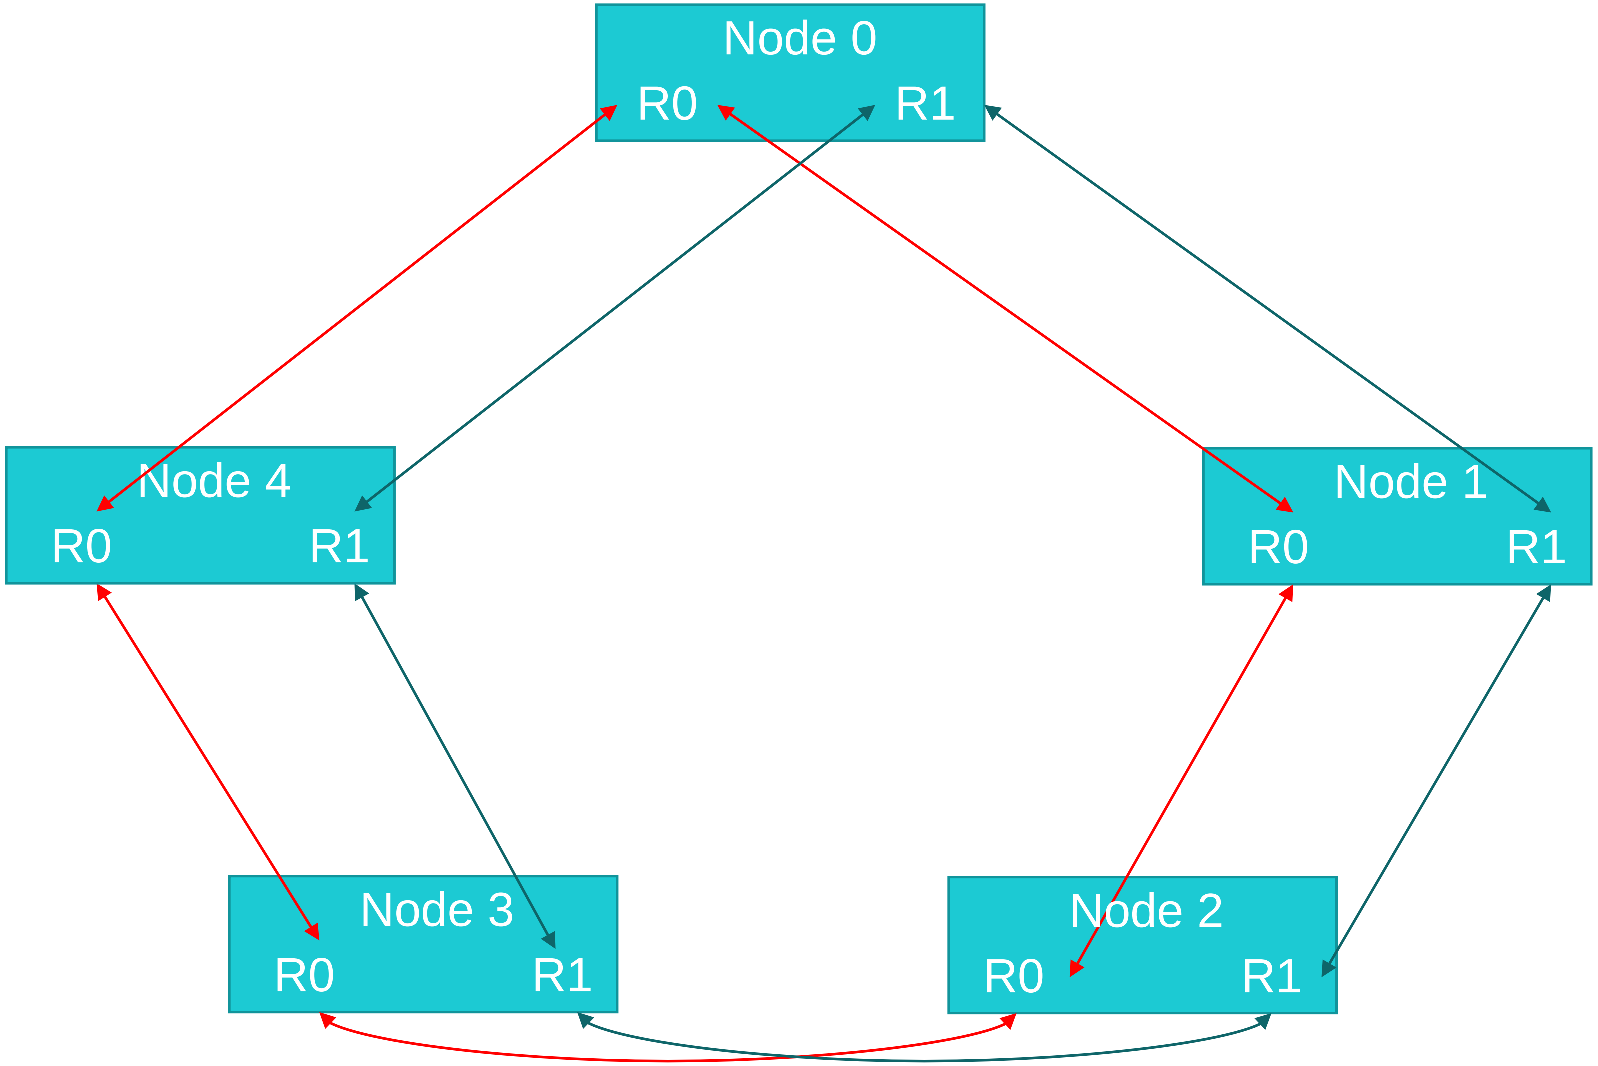

RANDOM RING COMMUNICATION PATTERN

- Proxy for one or multiple applications with
  - High-dimensional nearest neighbor pattern
  - Unstructured mesh
- Random ring has tunable connectivity
  - Each node communicates with up to 2×PPN unique nodes
    - E.g., Node 0 communicates with other 4 nodes
  - Increase PPN to proxy higher dimensional stencils
- Bandwidth sensitive to global bandwidth
- Latency sensitive to
  - Hop count distribution
  - Possibly contention

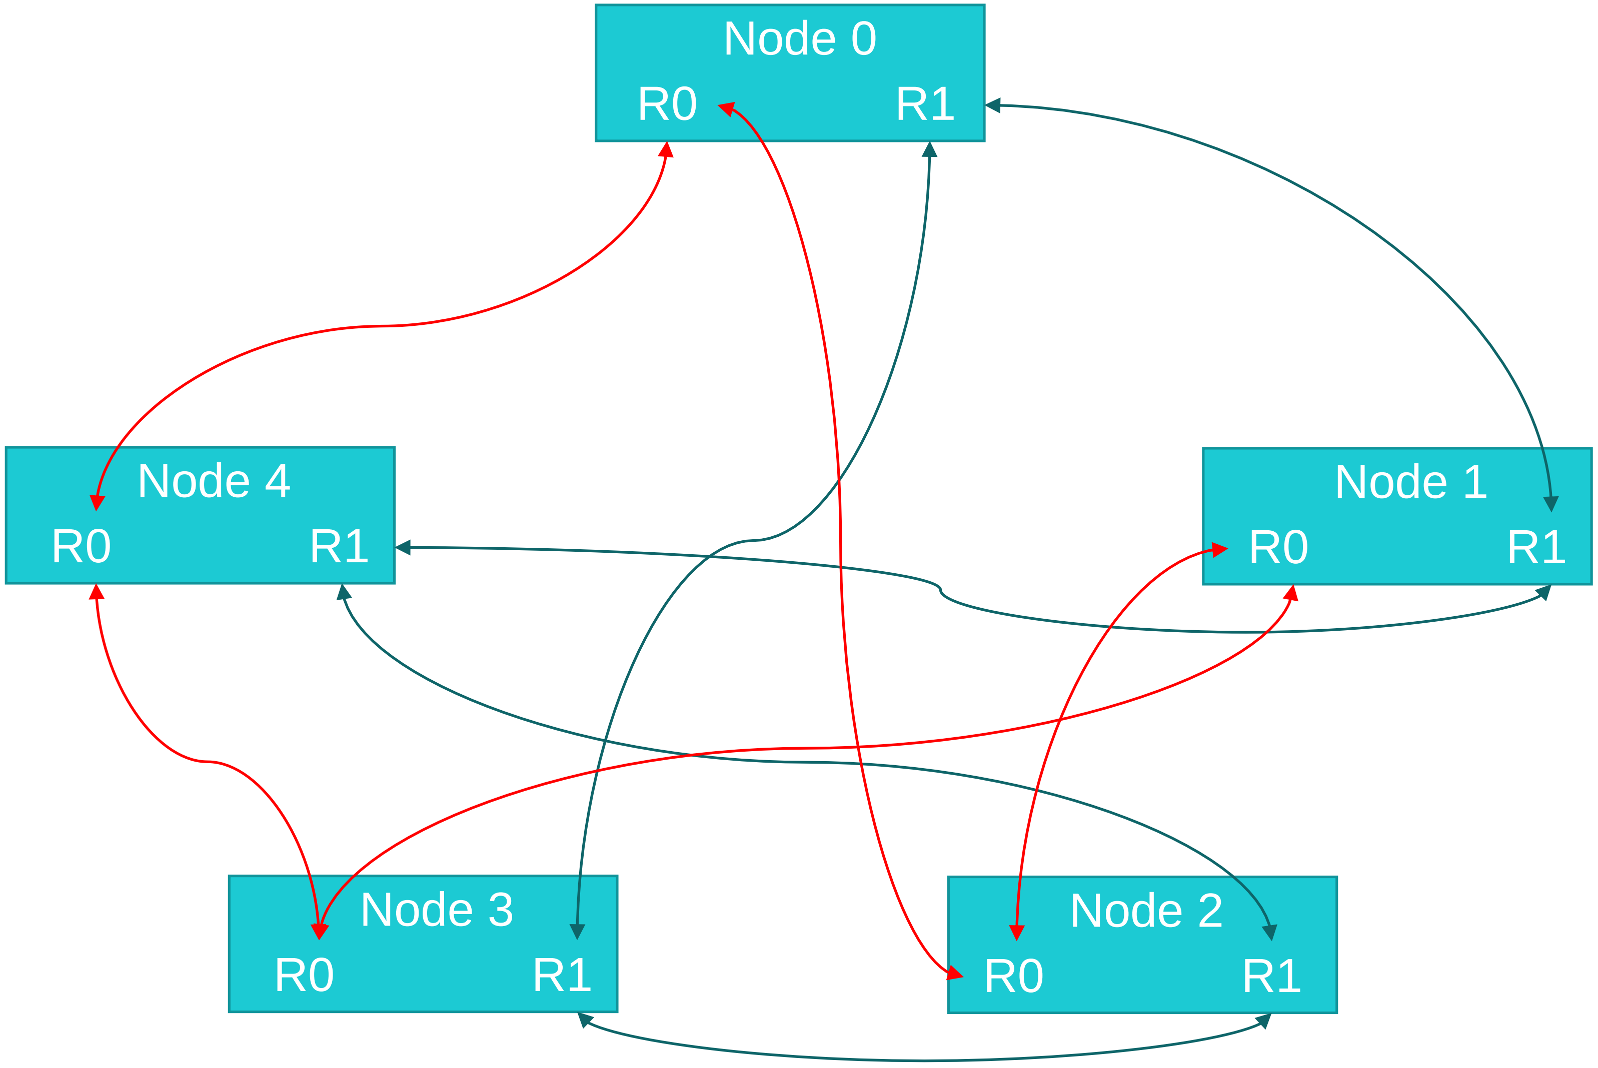

MEASURED LATENCIES – WHAT LIMITS APPLICATION PERFORMANCE?

- OSU ping-pong latency on 2 nodes (1 PPN)
  - ~ 1 usec

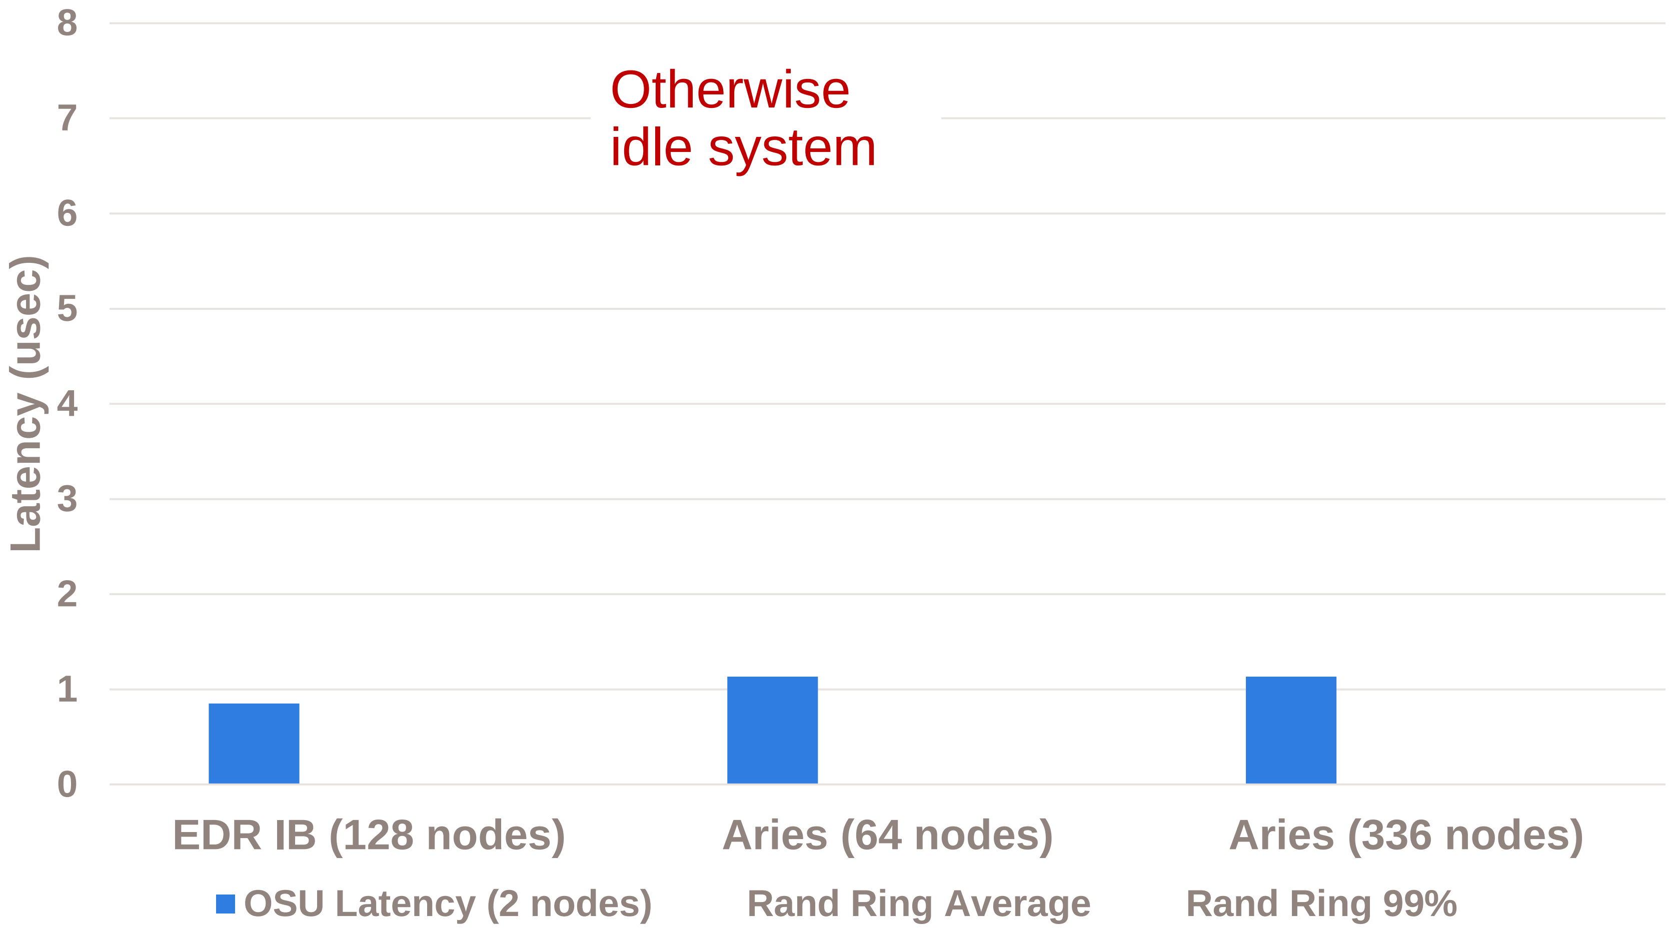

* Aries system: 336 dual-socket BDW nodes in a single dragonfly group

* EDR system: 128 dual-socket BDW nodes in full tree

MEASURED LATENCIES – WHAT LIMITS APPLICATION PERFORMANCE?

CRAY

- OSU ping-pong latency on 2 nodes (1 PPN)
  - ~ 1 usec
- Average Random Ring latency (16 PPN)
  - Depends on hop count distribution
  - ~ 1 – 3 usec

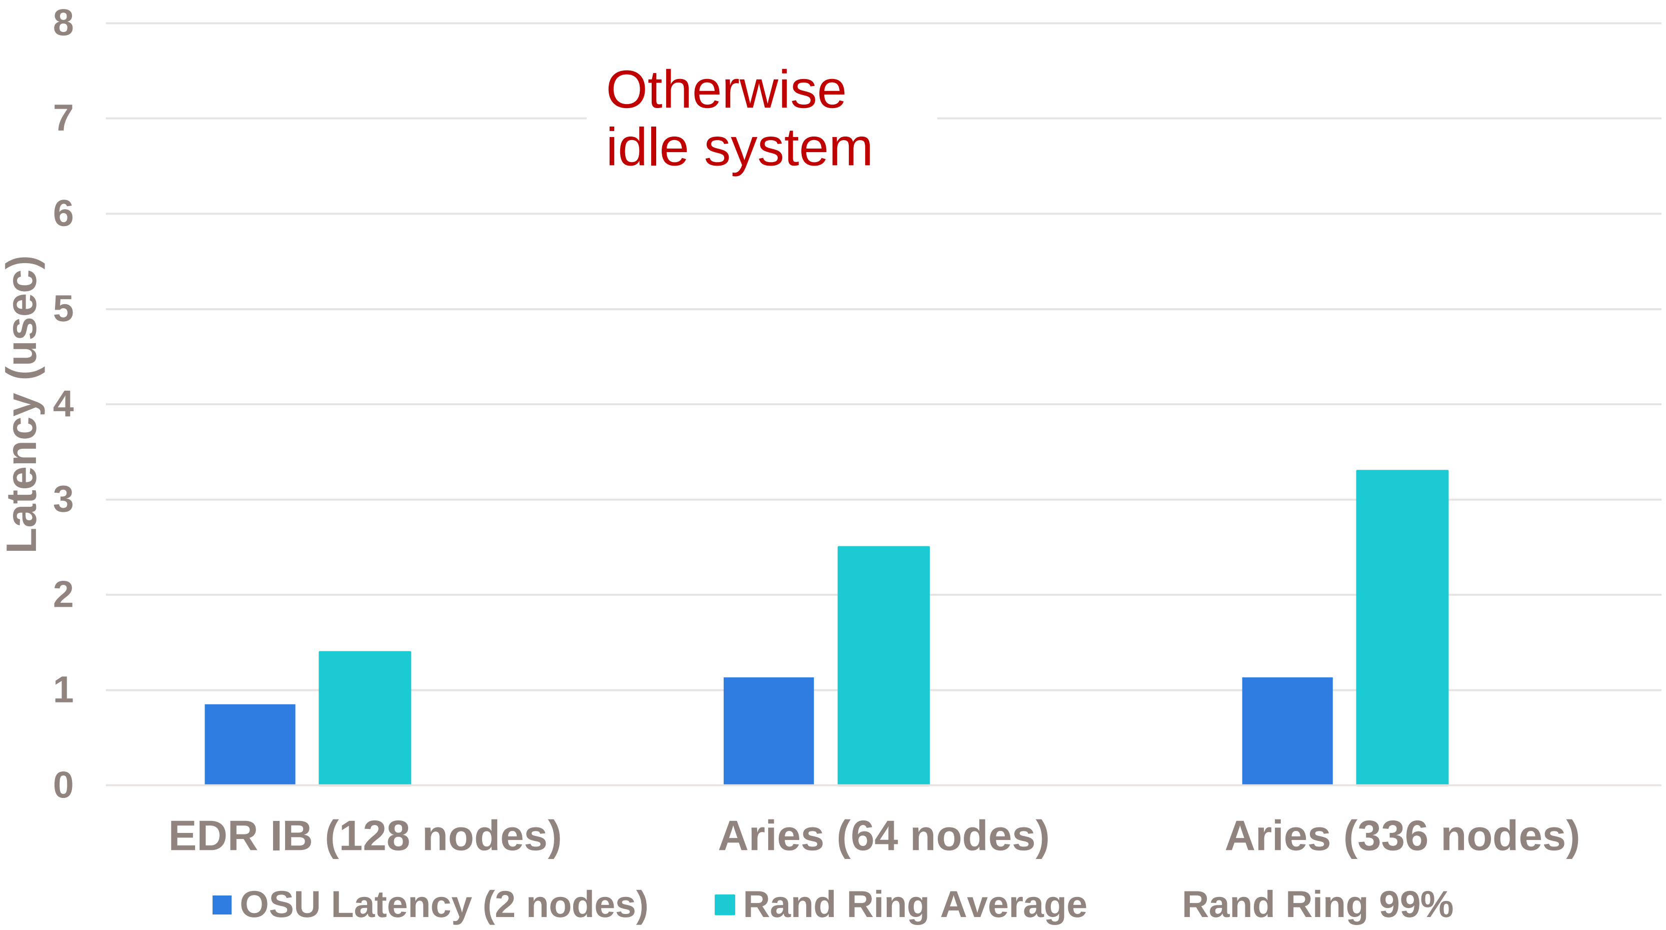

* Aries system: 336 dual-socket BDW nodes in a single dragonfly group

* EDR system: 128 dual-socket BDW nodes in full tree

MEASURED LATENCIES – WHAT LIMITS APPLICATION PERFORMANCE?

- OSU ping-pong latency on 2 nodes (1 PPN)
  - ~ 1 usec
- Average Random Ring latency (16 PPN)
  - Depends on hop count distribution
  - ~ 1 – 3 usec
- Random Ring 99% tail latency (16 PPN)
  - Also influenced by contention and other delays
  - Over 5 usec

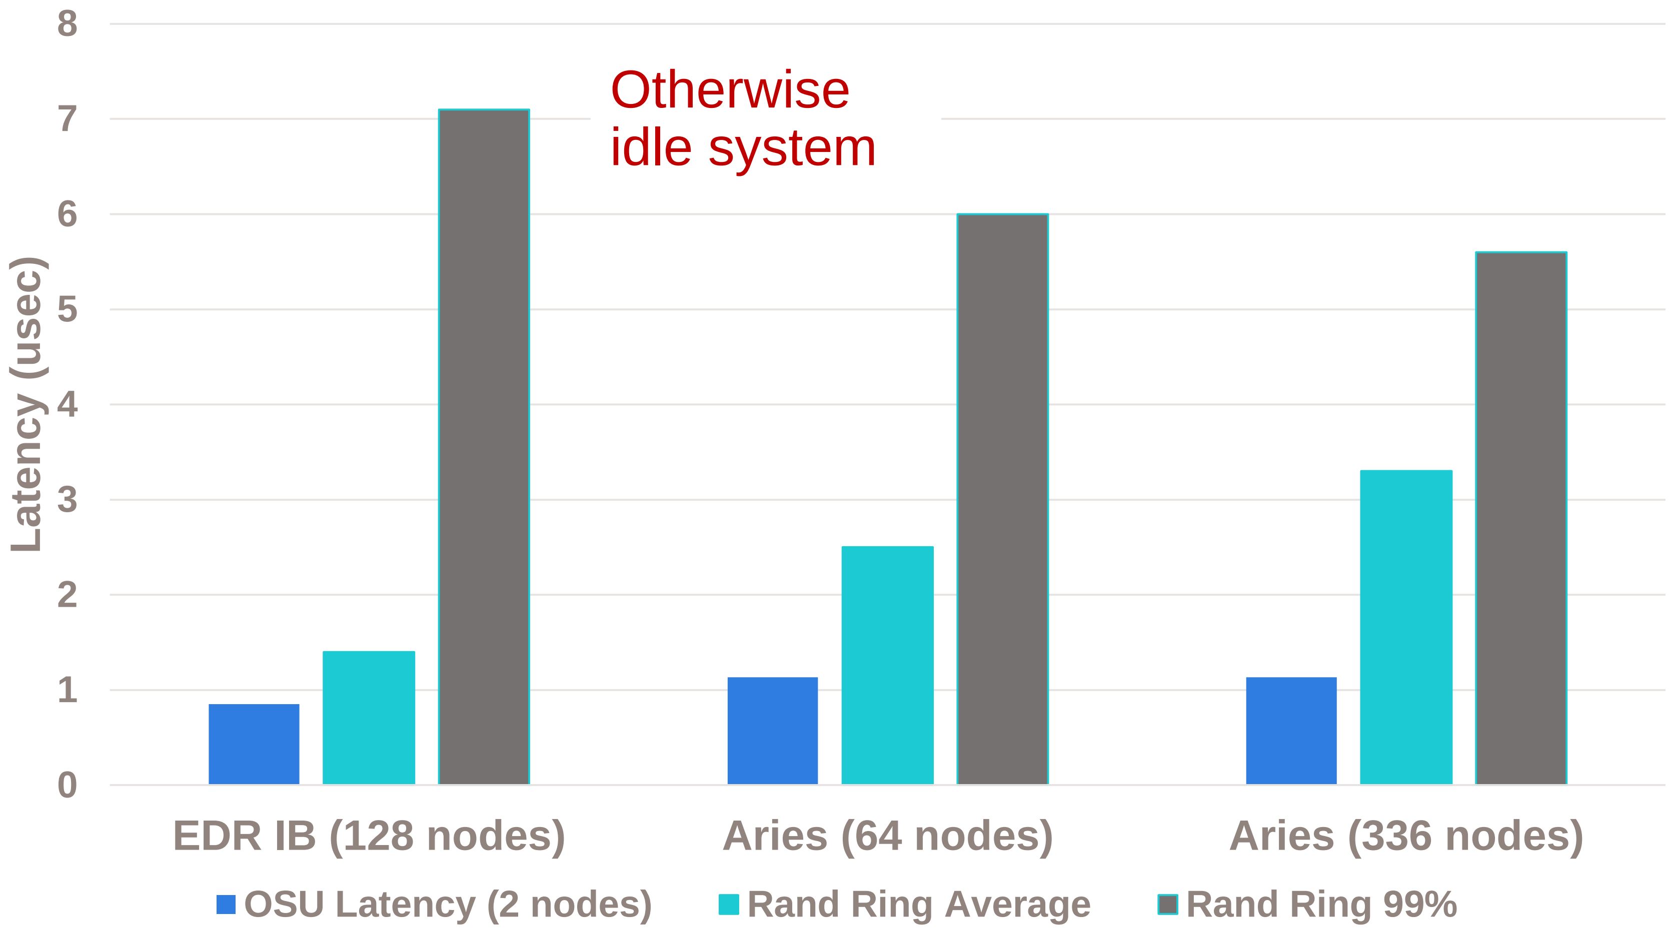

* Aries system: 336 dual-socket BDW nodes in a single dragonfly group

* EDR system: 128 dual-socket BDW nodes in full tree

MEASURING CONGESTION

HPC WORKLOAD INTERACTIONS

- Network contention and congestion occurs on HPC networks from
  - Demanding applications
  - Cumulative effects of a diverse set of applications and system services

- Communication patterns that can induce congestion
  - In-casts (e.g., many-to-few)
  - Many-to-many (e.g., all-to-all)

- Use the proxy patterns from the improved network tests as “canaries” for measuring the effects from congestion

CONGESTORS

- Extend benchmark to include several patterns known to cause application interference / variability
  - All-to-all
  - Point-to-point in-cast (e.g., MPI_Isend)
  - One-sided in-cast (e.g., MPI_Put)
  - One-sided broadcast (e.g., MPI_Get)
  - *Infrastructure allows others to be included*

- Benchmark divides nodes into 5 randomly selected sets
  - 20% running a selected application proxy
  - Four sets of 20% each running a congestor

- The benchmark runs these in coordination to measure
  - Baseline performance of application proxy
  - Performance of application proxy with congestors running

CONGESTION – LIMITING APPLICATION PERFORMANCE

- Random ring 99% tail latency increases
  - 10 – 100X on EDR and Aries
  - *Impact depends on number of nodes*
- Would observe significant slowdown of latency sensitive application
- Preliminary Cray Slingshot results
  - No impact observed from congestion

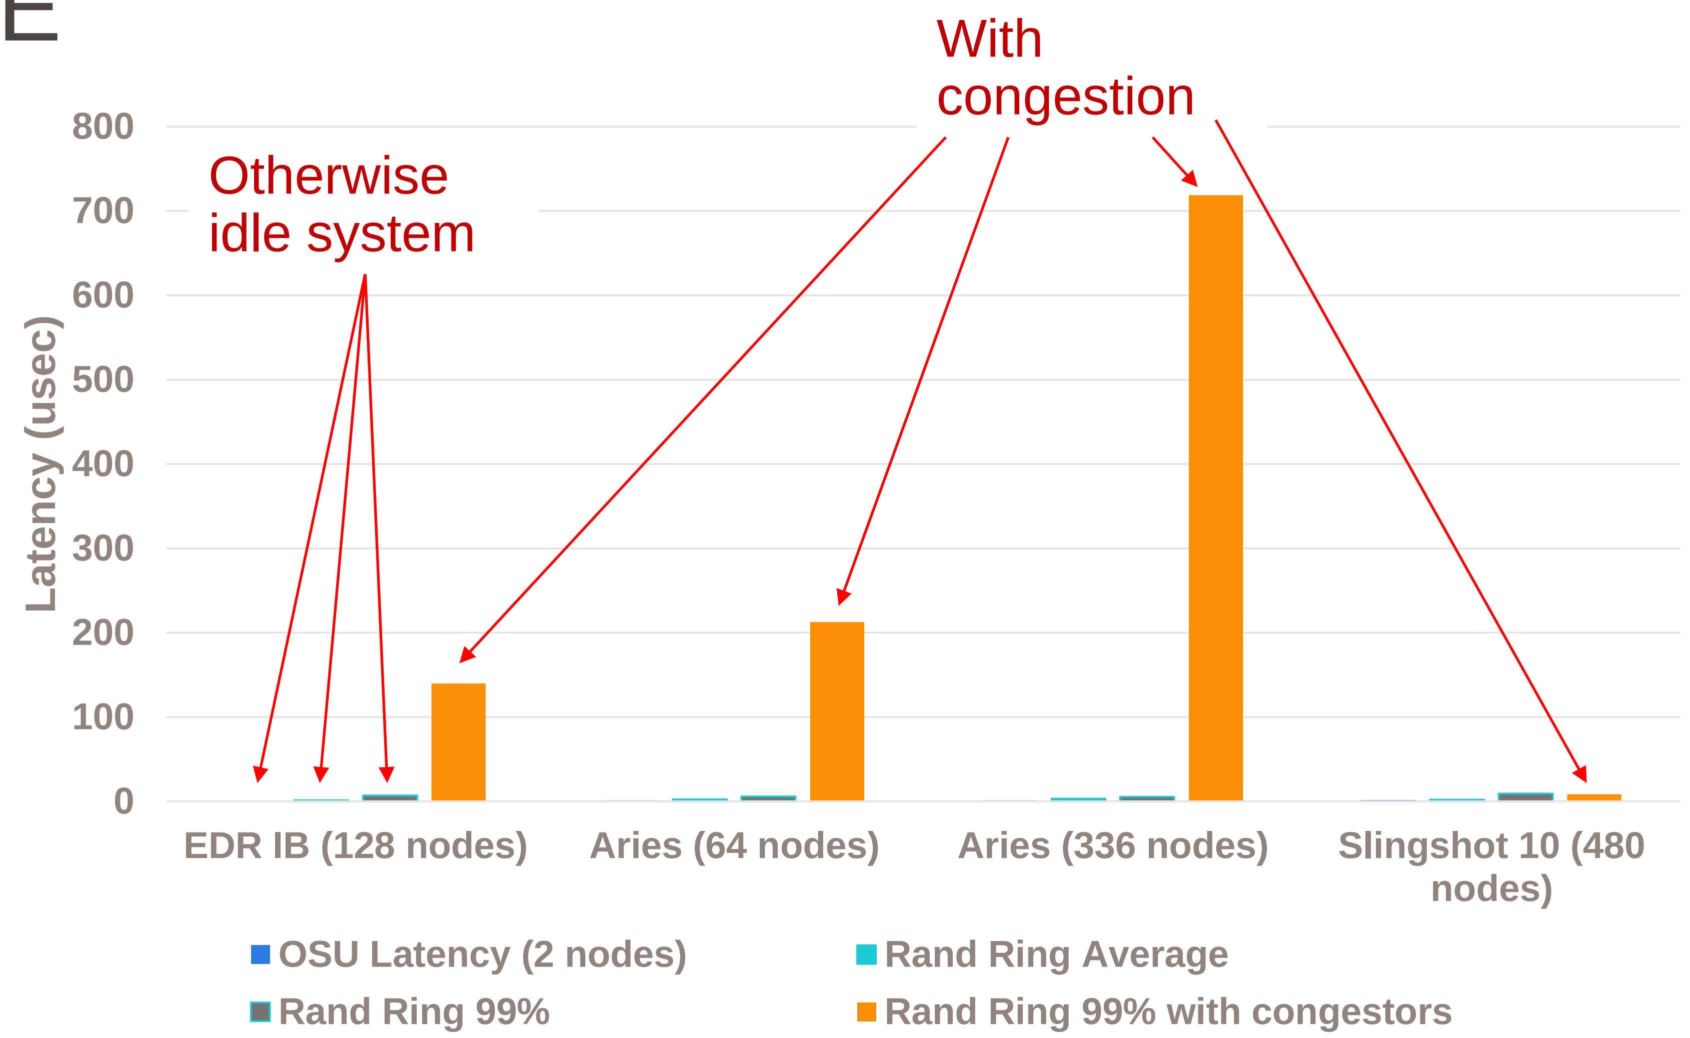

* Aries system: 336 dual-socket BDW nodes in a single dragonfly group

* EDR system: 128 dual-socket BDW nodes in full tree

* Slingshot system: 485 dual-socket SKX nodes in four dragonfly groups

CONGESTION – LIMITING APPLICATION PERFORMANCE

- Same plot now with log y-axis
- OSU ping-pong latency not a meaningful measure for network performance
- Preliminary Cray Slingshot results
  - Has sophisticated congestion management
  - Effectively contains congestor traffic
  - Congestors still at ideal performance
    - Incasts at single node ejection bandwidth
    - Alltoall near injection bandwidth

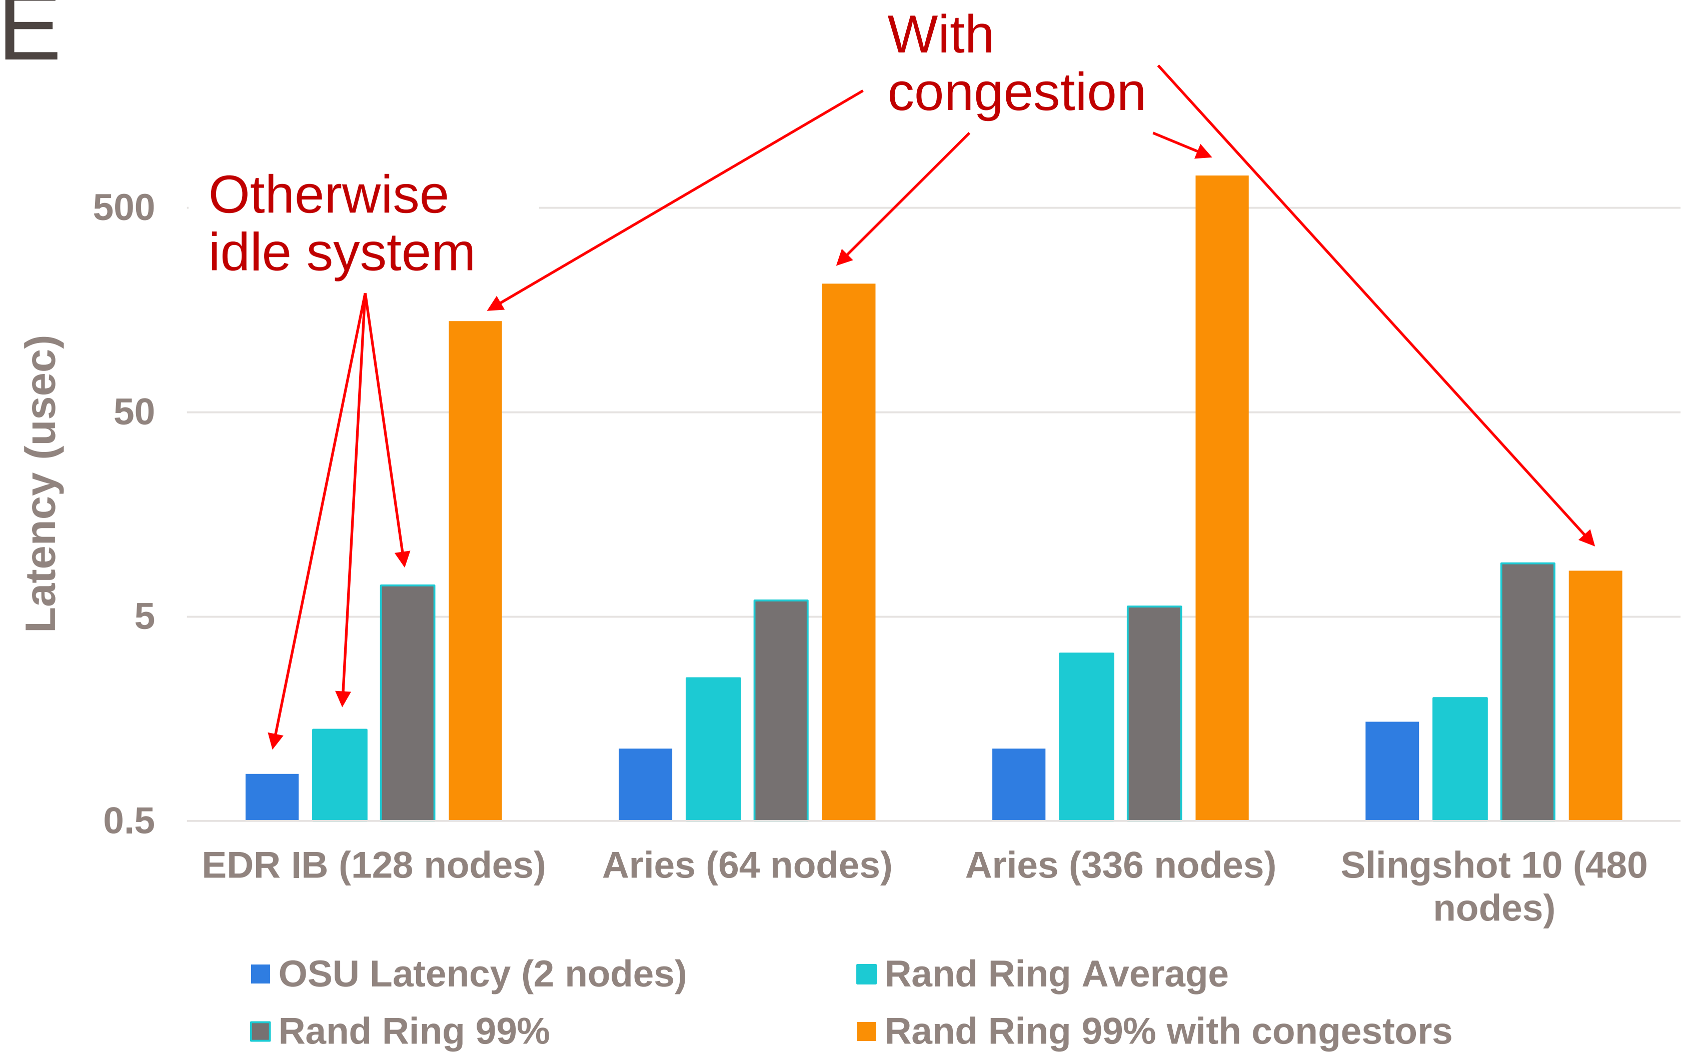

* Aries system: 336 dual-socket BDW nodes in a single dragonfly group

* EDR system: 128 dual-socket BDW nodes in full tree

* Slingshot system: 485 dual-socket SKX nodes in four dragonfly groups

GPCNeT BENCHMARK

- Presented overview of new methods for
  - Characterizing HPC network performance
  - Inducing and measuring impact of congestion

- Methods implemented into new benchmark called GPCNeT
  - Global Performance and Congestion Network Tests
  - Developed in collaboration with NERSC and ANL
  - Publicly available at
    - https://xgitlab.cels.anl.gov/networkbench/GPCNET

- Promising new tool for
  - Understanding HPC networks
  - Measuring capability of congestion mitigation features of a network
    - E.g., adaptive routing or congestion management
  - Designing future systems and networks

SUMMARY

- Common MPI benchmarks (e.g., OSU and IMB)
  - Useful for measuring performance of MPI library components
  - Not sophisticated enough to measure capabilities of modern HPC networks
  - Not able to measure the impact of congestion

- GPCNeT is a benchmark designed to
  - Measure HPC network capabilities
  - Impact of congestion

- Demonstrated congestion effects on EDR and Aries networks

- Cray Slingshot network
  - Has sophisticated congestion management features
  - Able to detect and isolate congestors
  - Little to no impact measured on sensitive network traffic

SAFE HARBOR STATEMENT

This presentation may contain forward-looking statements that are based on our current expectations. Forward looking statements may include statements about our financial guidance and expected operating results, our opportunities and future potential, our product development and new product introduction plans, our ability to expand and penetrate our addressable markets and other statements that are not historical facts.

These statements are only predictions and actual results may materially vary from those projected. Please refer to Cray's documents filed with the SEC from time to time concerning factors that could affect the Company and these forward-looking statements.

THANK YOU

QUESTIONS?

pjm@cray.com

cray.com

@cray_inc

linkedin.com/company/cray-inc-/