Update on Storage and Interconnects

ISC22
June 2022

www.HyperionResearch.com
www.hpcuserforum.com

Mark Nossokoff

System Interconnect Preferences

InfiniBand edges ahead of Ethernet in aggregate adoption

- **Initiated distinction between storage and system interconnects**
- **Ethernet**
  - Adopted at 44% of sites surveyed
  - 10 Gbit most widely adopted
- **InfiniBand**
  - Deployed at 46.1% of sites surveyed
  - EDR 100 Gbit most widely deployed
- **Omni-Path**
  - Some adoption across all sectors
  - 2nd largest deployed in academic sites (13.9%)

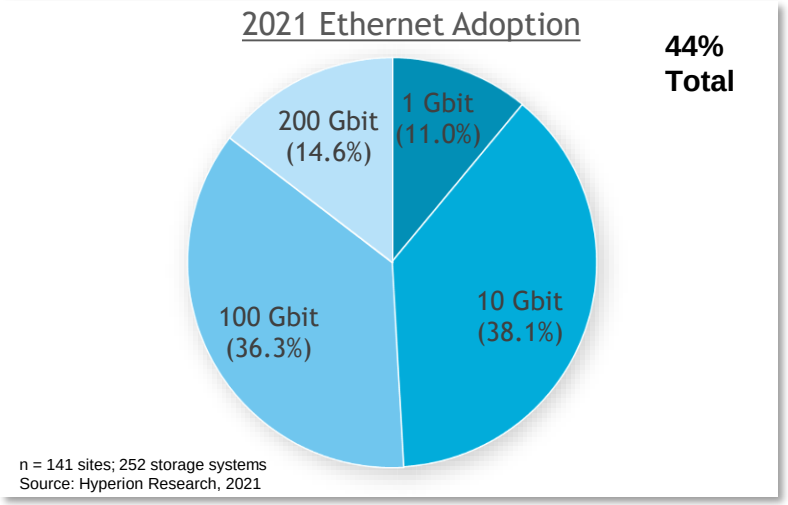

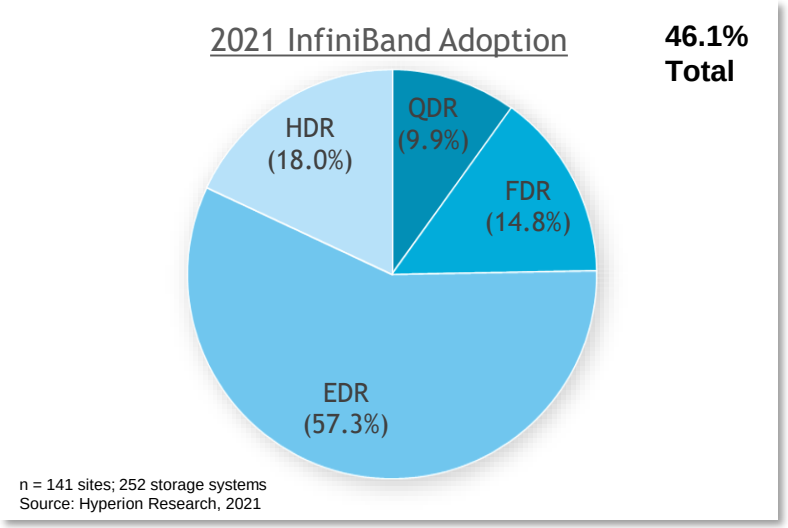

Storage Interconnect Preferences

Ethernet emerged as preferred storage interconnect

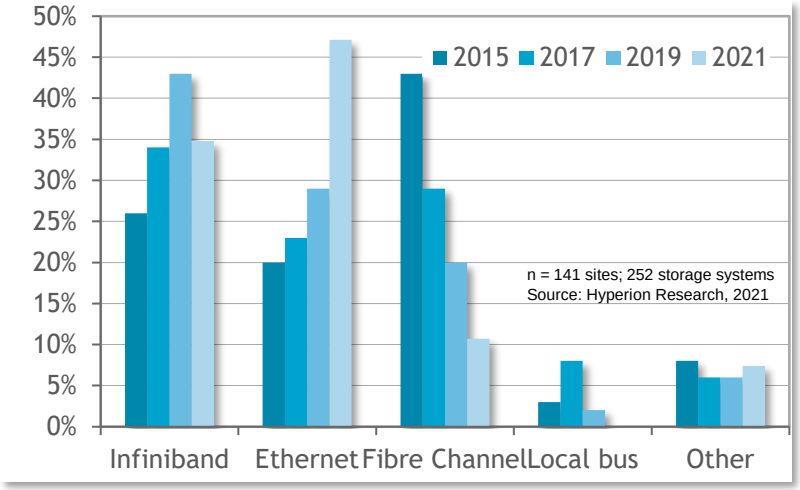

- **Sector preference differences**
  - Industry leans ethernet
  - Government split between ethernet and InfiniBand
  - Academia leans InfiniBand
- **Omni-Path is large part of other**

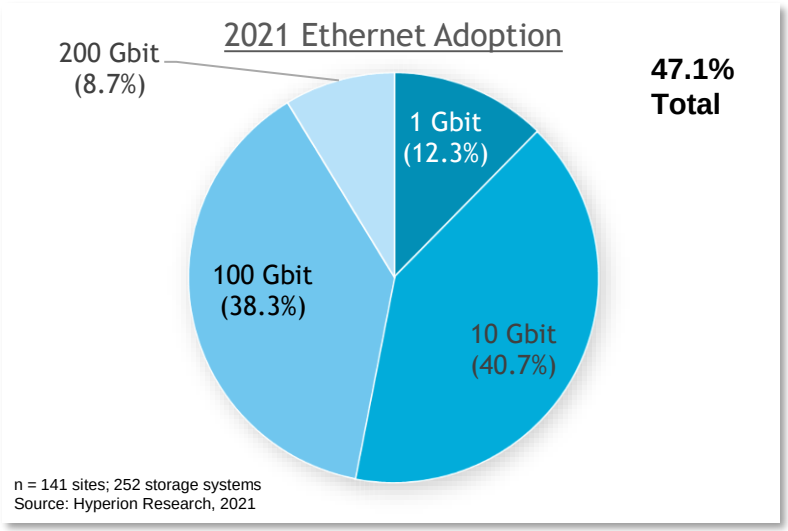

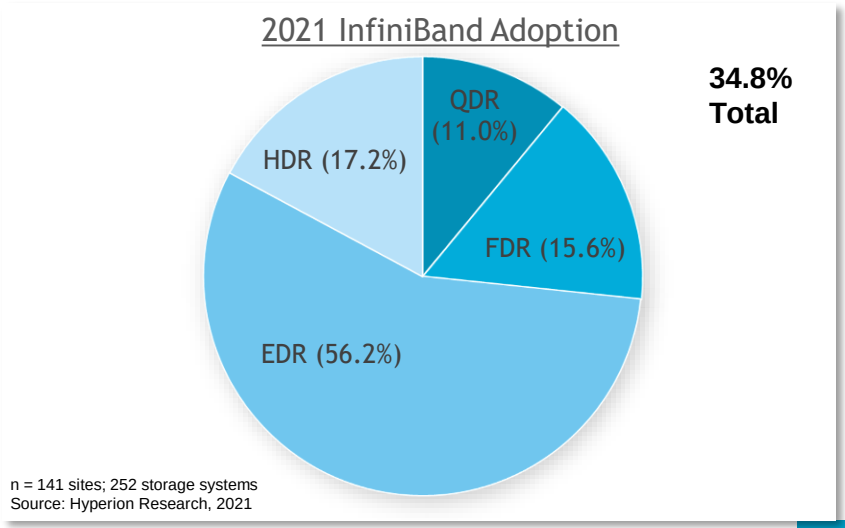

Interconnect Topics

System scale and heterogenous workloads are pushing existing system interconnects to their limits

- **Ethernet and ethernet-based**
  - Additional HPE Slingshot deployments
  - Rockport Networks emerging from stealth
- **Infiniband**
  - Continued investment in new features
  - Focus on datacenter scale and operating environment
- **OPA (OmniPath Architecture)**
  - DOE NNSA adoption of and investments in Cornelis Networks
    - Networking partner for Tri-labs CTS-2
    - $18M for NG-HPCN

Interconnect Topics

Composability, disaggregation driving interest in new interconnects and protocols for dramatic latency improvements

- **Emergence of DPUs**
  - Offload compute complex of data processing functionality
  - Mellanox, Fungible, and Pensando (soon-to-be-AMD) notable developers of DPUs
- **Industry-wide support of CXL**
  - Composable systems to dynamically pool resources depending on specific application and workload needs
  - Marvell acquires Tanzanite
  - Rambus acquires PLDA (Jun 2021) & Hardent (May 2022)
- **Google Aquila**
  - Convergence of ultra-low latency disaggregation and HPC/ML workload adoption by traditional datacenters
    - Cost-effective topology: direct-connected Dragon Fly
    - Ultra-low latency: shallow buffer cell-based GNet L2 protocol
    - Unified fabric: integrated RMA primitives with ToR-in NIC (TIN)
    - Co-existing in larger ethernet network: ethernet at the edges of Aquila-connected cliques

Interconnect Topics

Growing instantiations of chiplets are exposing the need for inter-die, intra-packaging interconnects

- **Collection of dies using different geometry nodes from different vendors**
- **Portability and re-use of older node-based functions**
- **Standardization efforts**
  - UCIe – Universal Chiplet Interconnect express
  - OCP ODSA – Open Compute Project Open Domain-Specific Architecture

HYPERION RESEARCH

PLEASE SHARE YOUR THOUGHTS!

mnossokoff@hyperionres.com

Our ISC22 Briefing Agenda

All registered attendees will receive the slide deck

- **Market Update and Forecasts**
- **Some Perspectives on European HPC**
- **Sustainability: No Longer a "Nice to Have"**
- **HPC and AI Talent Challenges**
- **Exascale Update**
- **Cloud Update**
- **Quantum Update**
- **AI Update**
- **Update on Storage & Interconnects**
- **Conclusions and Wrap-up**