



Flash Memory Summit

How to Network Flash Storage Efficiently at Hyperscale



Manoj Wadekar





Michael Kagan



ebay Hyper scale Infrastructure



Typical hyper scale servers



Performance Needs


In Memory/Search

Databases


Hadoop

Object Store
FE/Dev

Archival/Cold




Typical hyper scale servers: Design goals



Efficiency:
Utilization, commonality

Growth:
Performance, Capacity

Converged Infrastructure: challenges



Mismatched App Needs:

- Compute/Storage needs can change for different clusters.
- Can result in under-performance or wastage

Inefficiency:

- Spend, infrastructure space, power utilization
- Further challenge to justify high density/high performance drives


Scale Challenge:

- “Cattle” use cases for data-heavy workloads may result in large data movement
- Complicated storage scheduler leads to constrained scaling

Server Platform: Shackled to local storage

What's needed:




Disaggregated Storage



Separate out storage and compute resources

What's needed:

Rack-As-A-Compute



Why:


Rack-As-A-Compute

Right Sizing:


- Clusters can use optimized ratio of compute and storage.
- Allows reducing wastage and improve performance

Independent Scaling:

Compute and storage capacities can be scaled per need



Disaggregated Storage: Interconnect Needs



Throughput:

- Sequential workloads driven by throughput
- Aggregated storage drives higher needs

Latency:

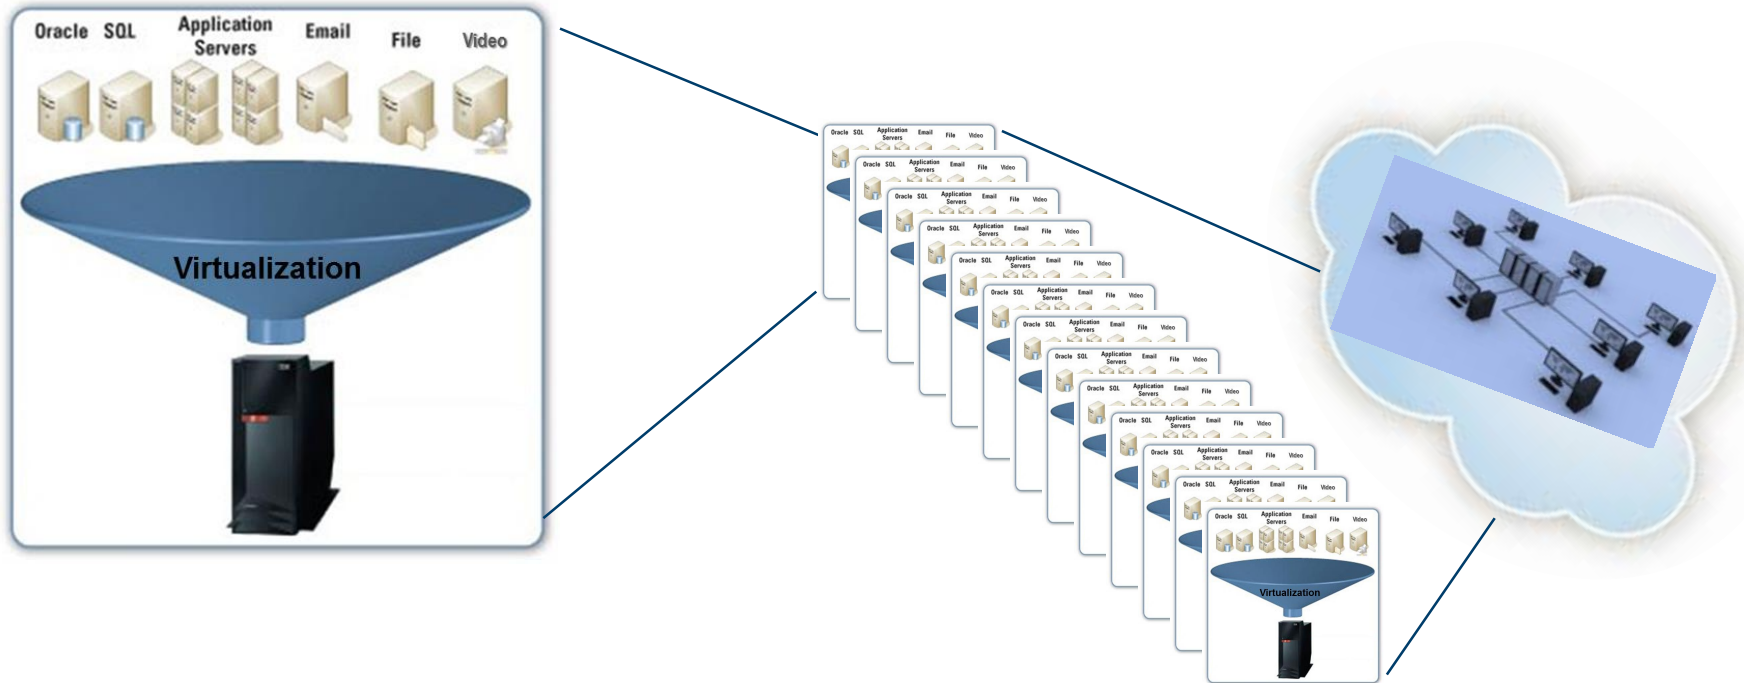
- IOPs sensitive workloads
- Appropriate deployment topologies

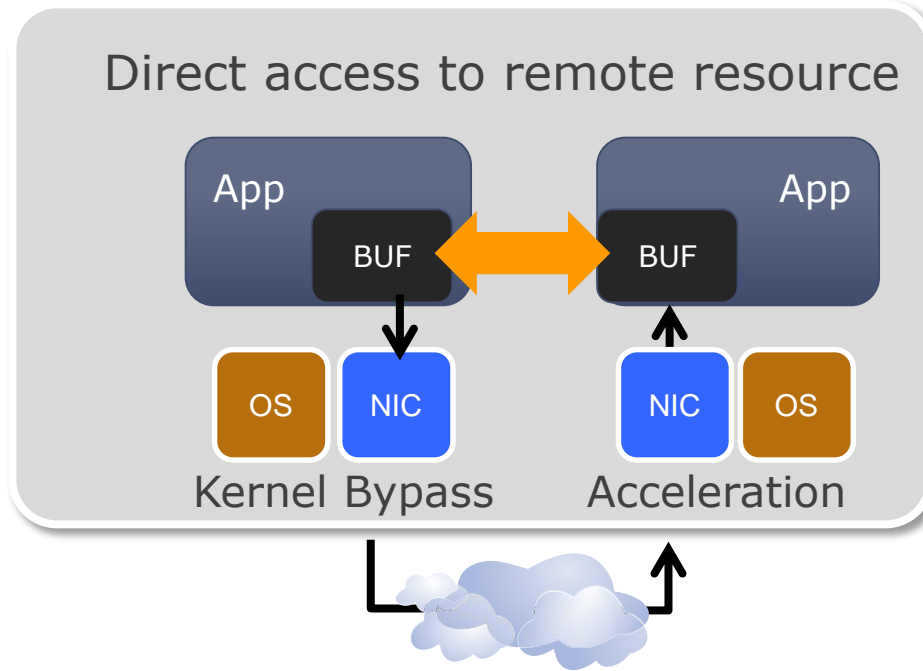
Simplicity:

- Known, ubiquitous network


Is Ethernet Ready?


The Paradigm Shift – Resource as a Service





RDMA over Converged Ethernet – RoCE





Storage Media Access Time

10,000X


Improvement



Networked Storage Continues Growth and Moves to Ethernet



Hyperscale Server SAN & Enterprise Server SAN Revenue Projections 2012-2027




Source: © Wikibon Server SAN Research Project

Networked Storage (SANs)

- Better utilization: capacity, rack space, power
- Scalability, management, fault isolation



Ethernet & Fibre Channel Switch Port Shipment



Crehan Research Long Term Forecast, July 2017

- Ethernet growing very rapidly driven by:
 - Cloud & Hyper Converged Infrastructure (HCI)
 - No Fibre Channel in the Cloud
 - NVMe Over Fabrics
 - Software Defined Storage

Ethernet Storage Fabric





*Everything a Traditional SAN Offers but ...
Faster, Smarter, & Less Expensive*

Performance

Efficiency

Ubiquity

Local NVMe vs. Remote NVMe-oF



200 Gb/s of throughput



BlueField

More than 5 million IOPS



ConnectX[®]



TILE


Industry-Leading **ConnectX** Intelligent Offload
Tile Multicore ARM Architecture

ASIC




- Family of products
 - 4, 8, 16 cores

Storage Controller




- NVMe storage solutions
- Customized per-OEM

Adapter



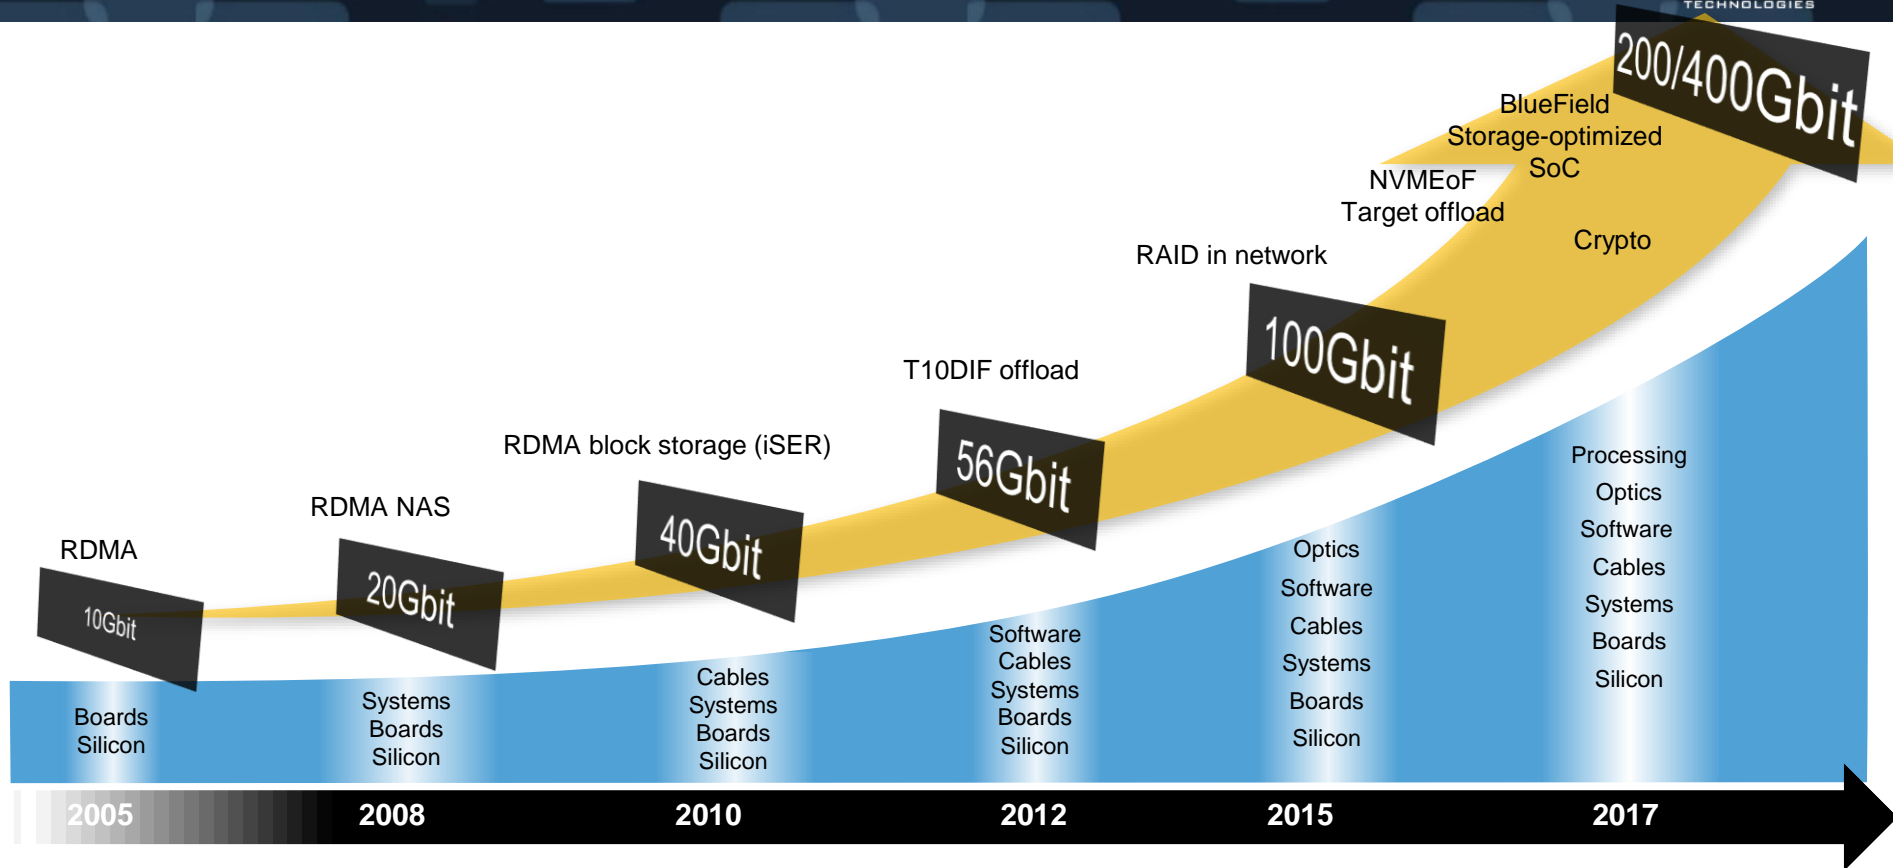
- Smart NIC
- Standard PCIe slot

Open Platform Software – Standards Based
BlueOS Linux – Built from kernel.org
Standard Linux development tools to build your application

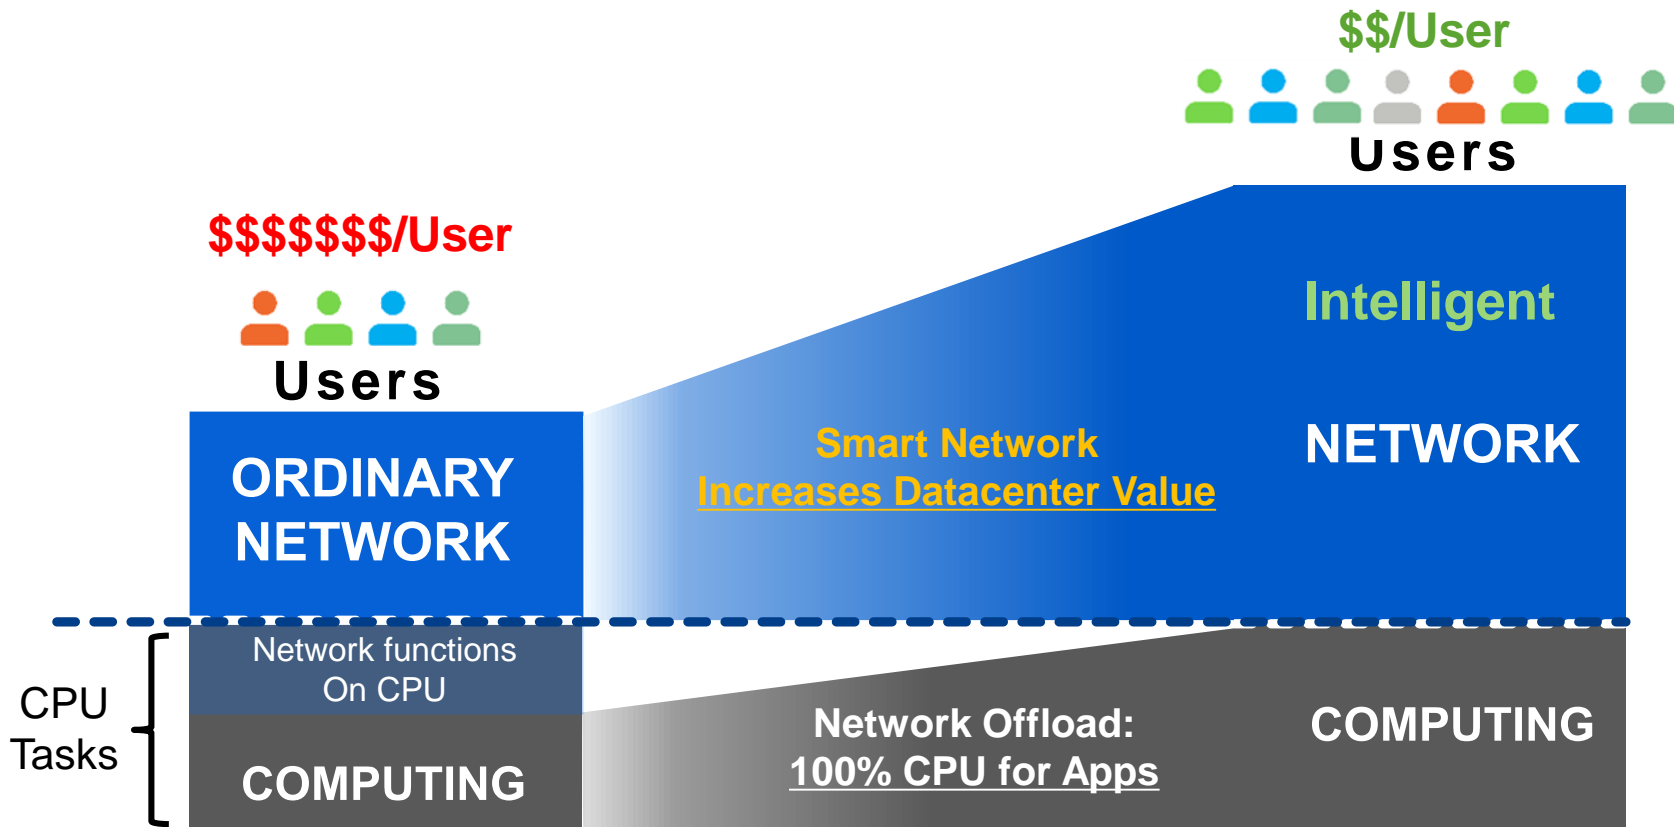


SCM Over Fabrics Definition WIP in IBTA

Technology Leadership



Network is a Computer – Deliver Value





Flash Memory Summit

Questions?

Manoj Wadekar



Michael Kagan

