Server Optimization

Maximizing Bandwidth in the Virtualized Data Center

Dave Zabrowski

President and CEO

S2io, Inc.

Content

- The data center crisis
- Optimizing the server infrastructure
- Choosing the right interconnect solution
- The I/O fabric convergence

The data center crisis

The data center crisis

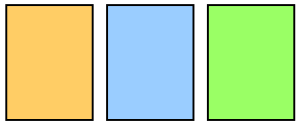

- Corporate IT must meet ever increasing demand
- Budgets are flat or declining to meet the demand

1990's

- Server "sprawl" fueled by economic growth and falling costs of server hardware
- Mission critical applications grew in silos
- Servers individually provisioned to handle peak workloads

Today

- Need to get visibility of all assets
- "Vertical" ownership by lines of business leading to "horizontal" inefficiencies
- Complicated network topologies
- Increased security vulnerabilities

Tomorrow

- Internet and traffic is doubling on an annual basis
- IT must scale to meet demand but with higher ROA
- New architectures are emerging requiring scalable I/O fabrics

“ From email to data warehousing, enterprise application data will increase at a compounded annual growth rate of 30 to 40 percent.”

Salomon Smith Barney Study

“Every Little Gigabit Helps”, November 2003

The outlook is not better

- Enterprise data growing between 30 and 200 percent annually, with new generation of applications (grid databases, video streaming, etc.)
  - E.g. Oracle 10g handles databases up to 8 exabytes
- Growing population accessing the internet
  - 795 million people on internet: 12.3% w/w vs. 69% in NA
  - Only 7% connected in Asia (3.7b potential users) (Internet World Stats, July 2004)
- Growing expectations, but limited if any increases in budget
  - IT budgets expected to remain flat through 2007 (Gartner Dataquest, January 2004)

“ If an environment hosts 200 users and applications are very interactive, plan bandwidth to accommodate 500 users.”

DH Brown Associates

“Enterprise Server Consolidation”, September 2001

Optimizing the server infrastructure

Three keys to optimization

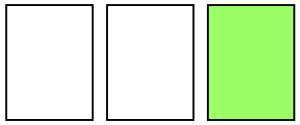

- "Vertical" consolidation
  - Replace several mid-sized systems to a few large servers
- "Horizontal" consolidation
  - Replace several mid-sized systems to a system of servers (Blades, Grid, Modular computing, etc.)
- Virtualization of new or existing assets
  - Move workloads between discrete machines, or re-host applications on larger servers during anticipated spikes

Virtualization Key for an Adaptive Enterprise

An approach to IT that pools and shares resources so utilization is optimized and supply automatically meets demand

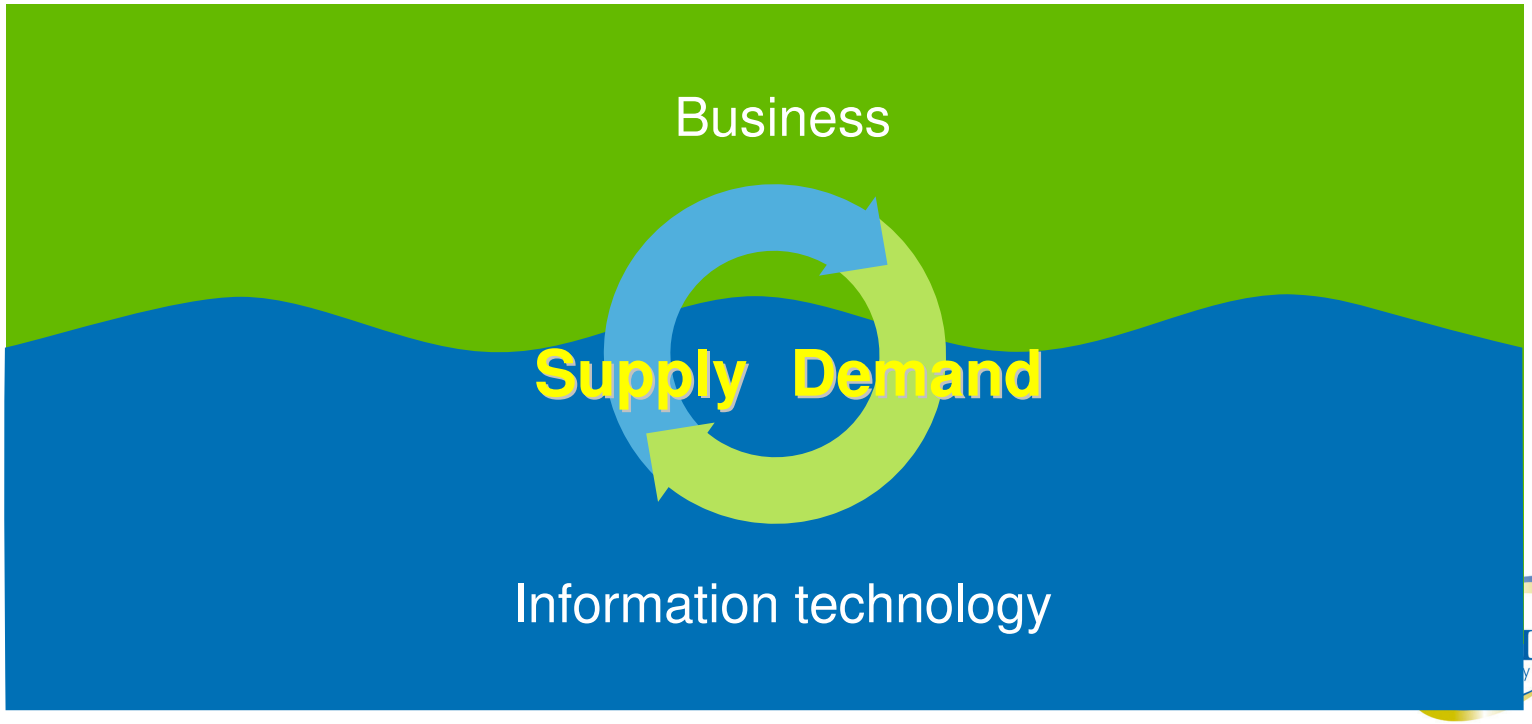

Benefits

- Lower TCO through optimal use of server and storage assets
- Reduced data center management and IT staffing costs
- Greater ability to adapt to changing service level demands
- Simplify interconnect fabric and server management
- Consolidation also leads to reductions in floor space, cabling, power and other physical plant requirements
- Free up expensive server I/O slots to make room for future expansion

Top reasons for datacenter consolidation

% of respondents

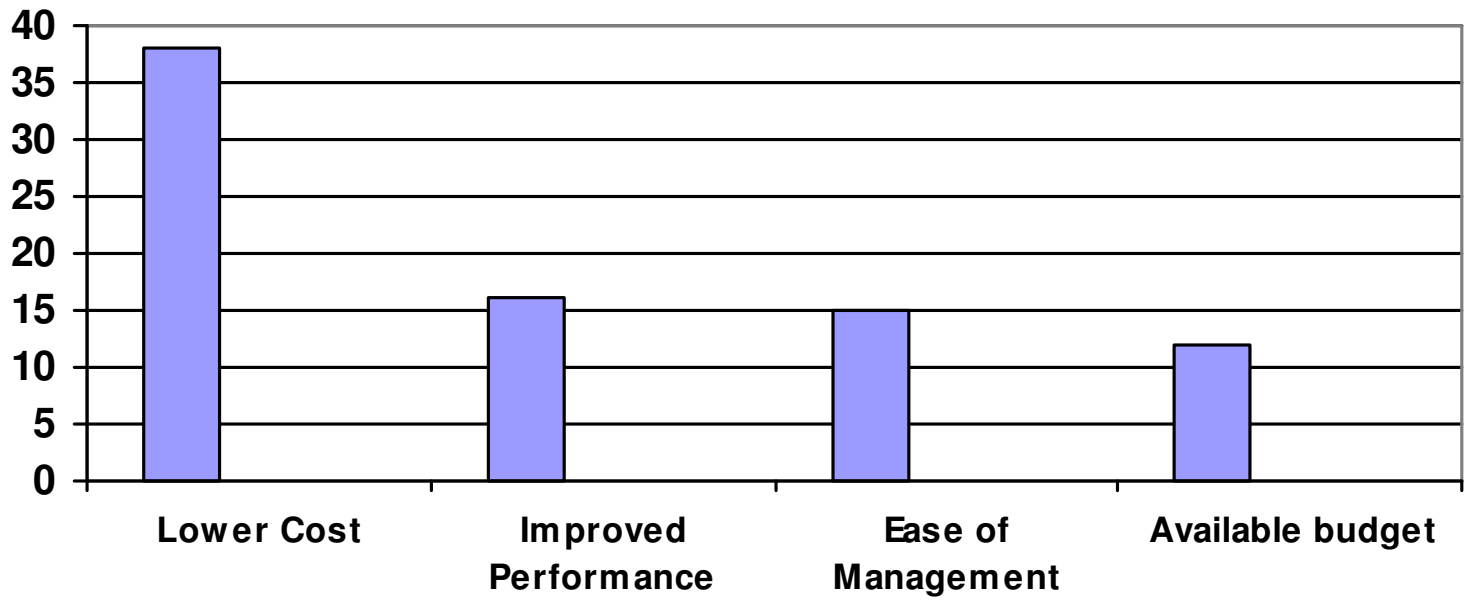

Source: IDC 2002 IT survey

Virtualization: unlocking the value of IT assets

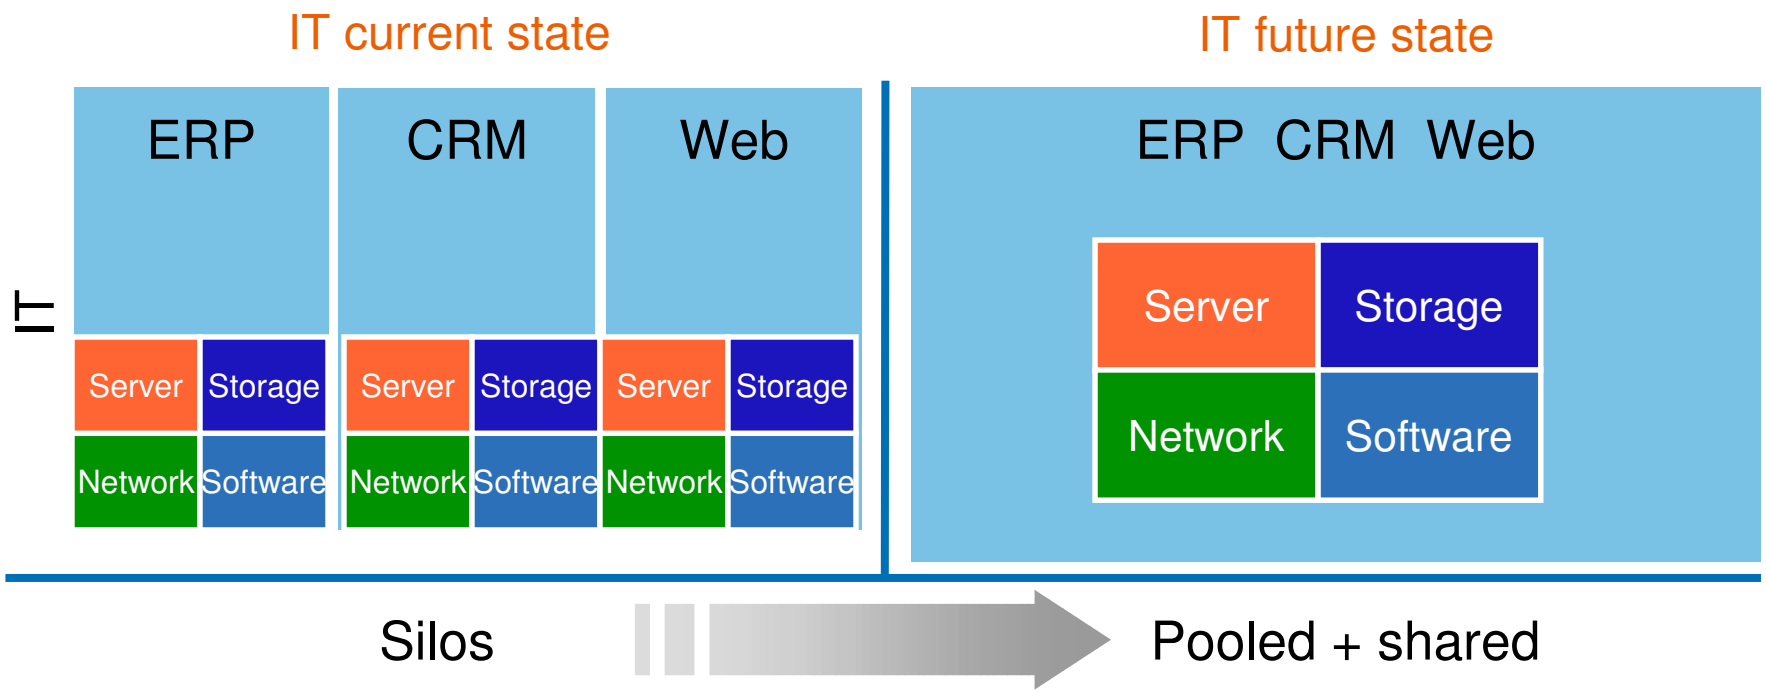

Silo IT – inefficient

- Fixed capacity and cost
- Under-utilized + over-provisioned
- Complex and difficult to change

Virtual IT – agile

- Dynamic capacity – pay per use
- Optimized resources
- Simplified and flexible

Maximizing the Business Value of IT Consolidation with Virtualization

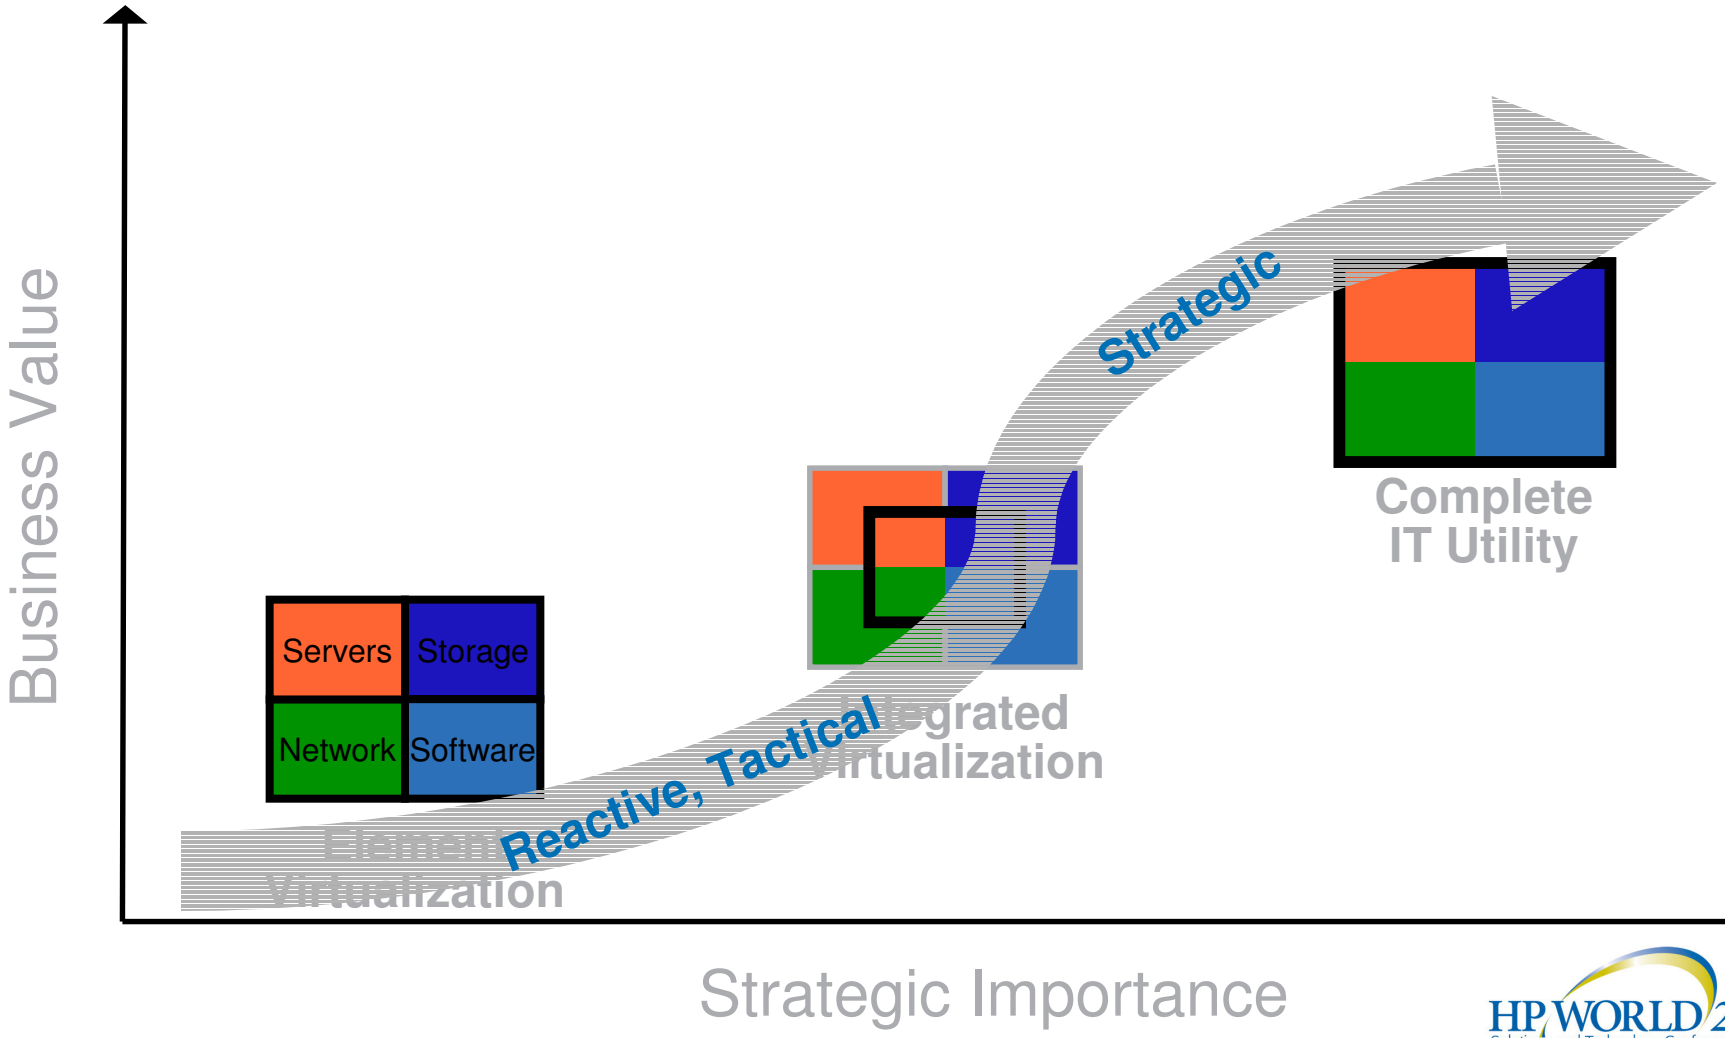

The optimized data center vision

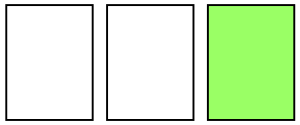

- Simplified
  - Cut administrative and operational costs, which make 90% of the Total Cost of Ownership (Meta Group, July 2004)
- Centralized
  - Centralized storage environments allow administrators to manage up to 20 times more capacity than direct-attached
  - Consolidated storage capacity utilization rate go up to 80% or better vs. less than 40% in direct-attached (Datalink Corp, Oct. 2003)
- Standardized
  - More standardized infrastructure saves labor time and cost, and allows the enterprise to remain agile in the face of constant change
- Strategic
  - From tactical operational cost-center to strategic arm of the enterprise, vital to customer satisfaction and employee productivity

Choosing the right interconnect solution

Interconnect options for the data center

- Infiniband
  - Low latency, high-bandwidth
  - Semi-proprietary, Meager product ecosystem, threat to long-term investment
  - Cost of total solution. Requires specialized tools and training
- Fibre Channel
  - Popular in HPC environments and SANs
  - 1, 2, 4 Gb/s (10 Gb/s in the future)
  - Upgrades require costly retooling
- Gigabit Ethernet
  - Low Acquisition Cost
  - Leverage Ethernet and IP standards
  - Low Performance/low scalability
  - High complexity as network grows leads to high management Cost
- 10 Gigabit Ethernet

The 10 Gigabit Ethernet Solution

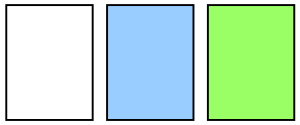

Today & Tomorrow

- High Performance
  - Data throughput is increased 6-8X's
  - 50% latency decrease speeds delivery of data
  - TCP processing overhead reduced by 40% using stateless offload
- Investment protection
  - Simple integration, no OS modifications required
  - Same network administrative tools, no retraining required
  - Low number of industry standard cables
- Highly Leveraged Volume Economics
  - Leveraged volumes with other components
  - Lower cost/bandwidth than 1 Gb Ethernet
  - Lower fundamental cost than Infiniband

Data center optimization – Case Study

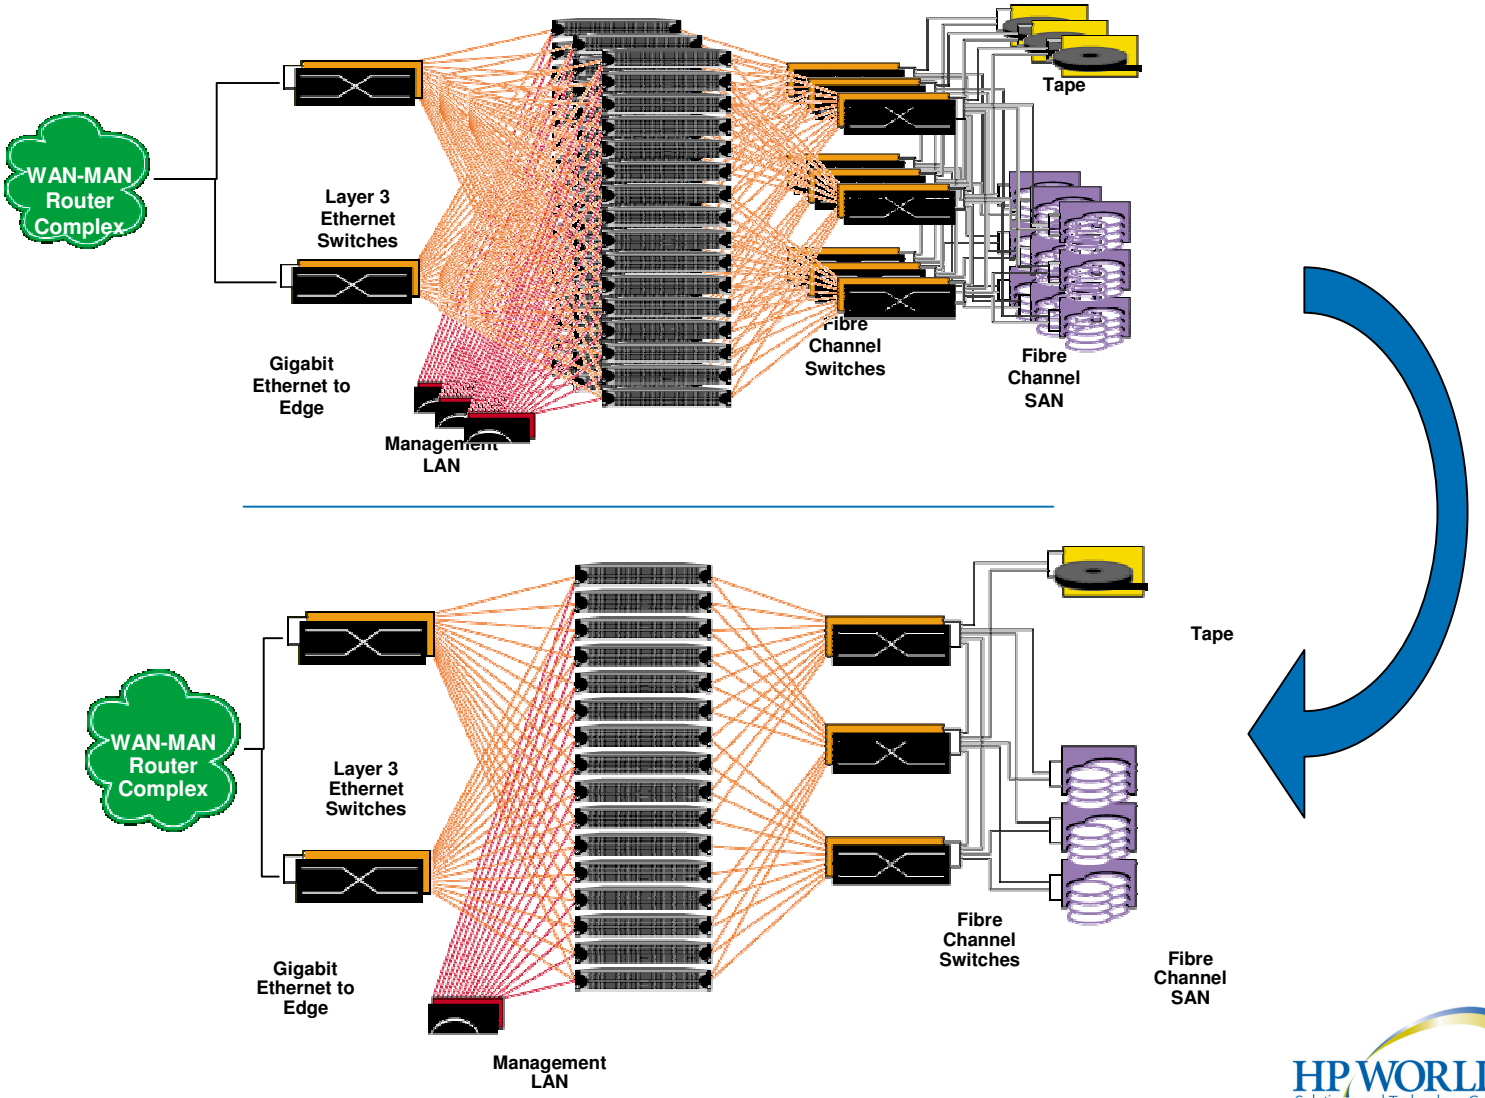

Data center optimization – Case Study Results

- Performance improved dramatically
  - Throughput improved from 1.8 Gb/s to 7.4 Gb/s
  - Latency improved by a factor of 4X which improved application performance
- TCO benefits
  - Management much simpler (6-8X reduction in ports)

Time (hours) to back up 20 TBs

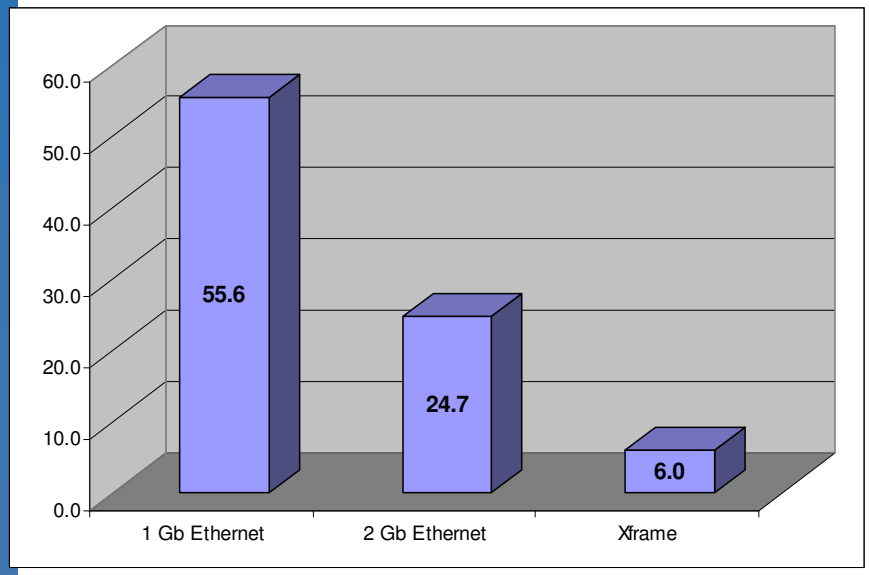

Real World Latency

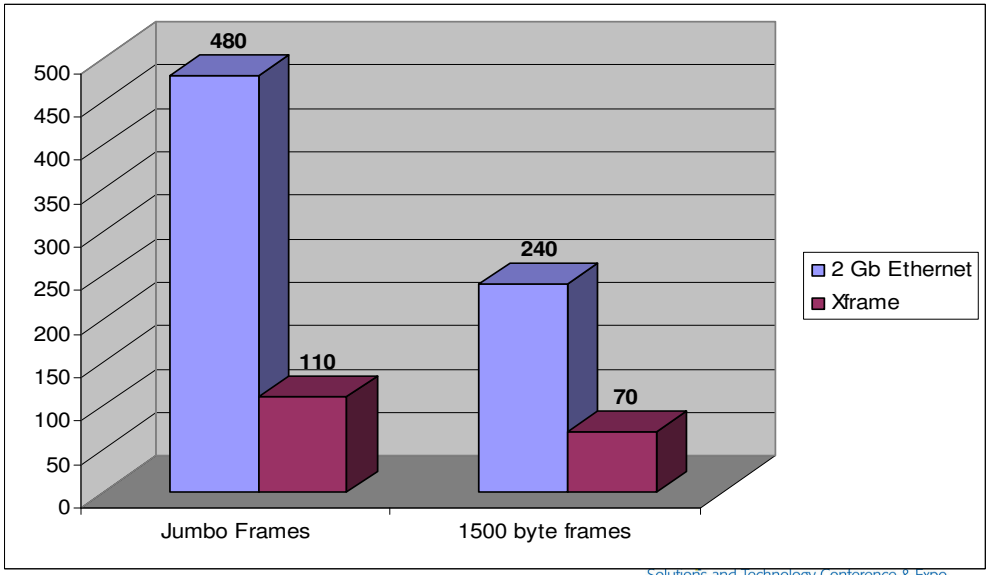

End User Demand: 40% intend to adopt 10 Gb Ethernet within 2 years

Plans to deploy 10 Gigabit Ethernet server network adapters in your network? (n=282)

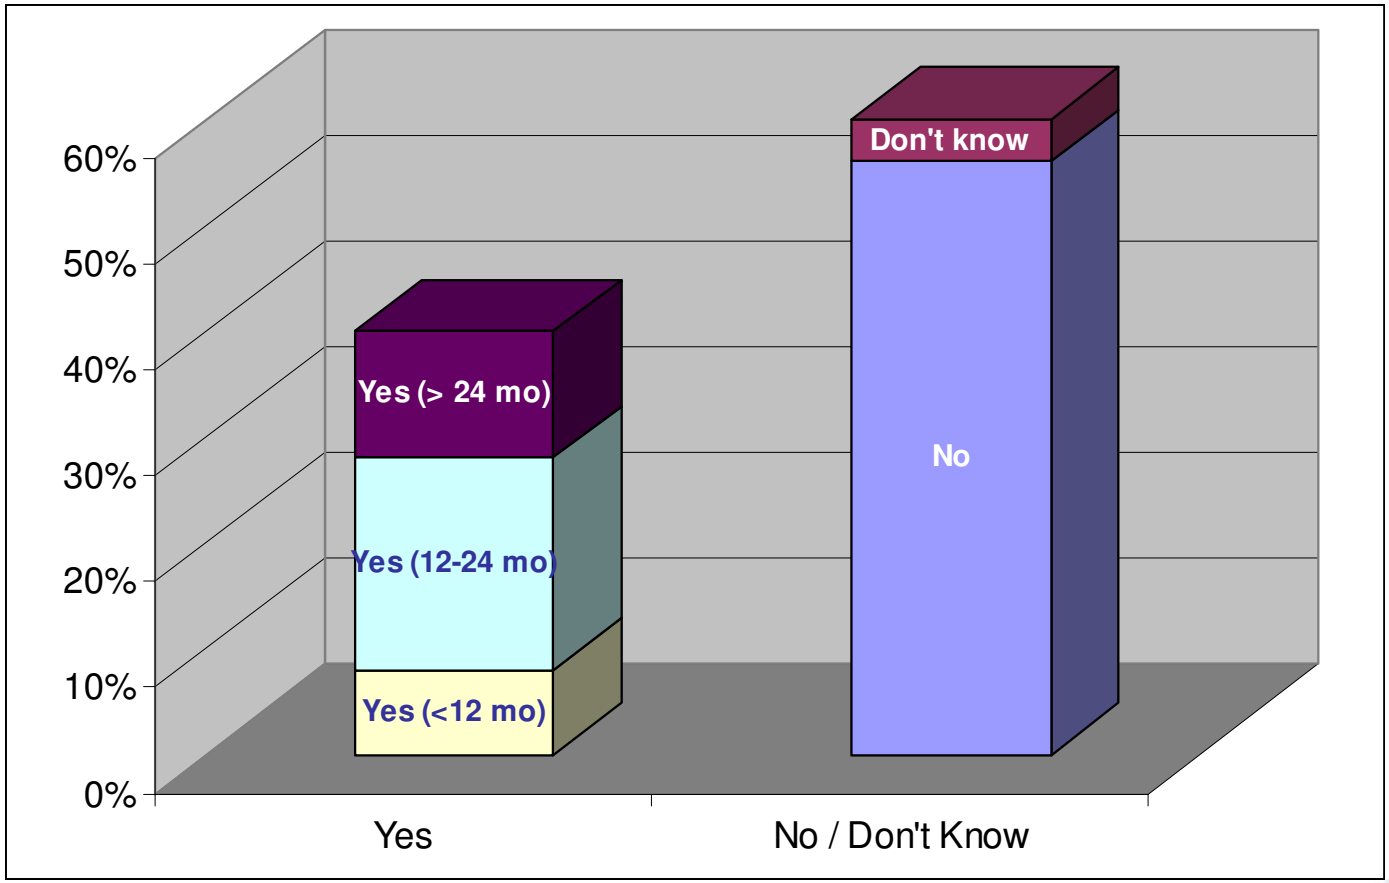

Source: In-Stat/MDR, 05/03.

10 Gb Ethernet Server Adapter Forecast

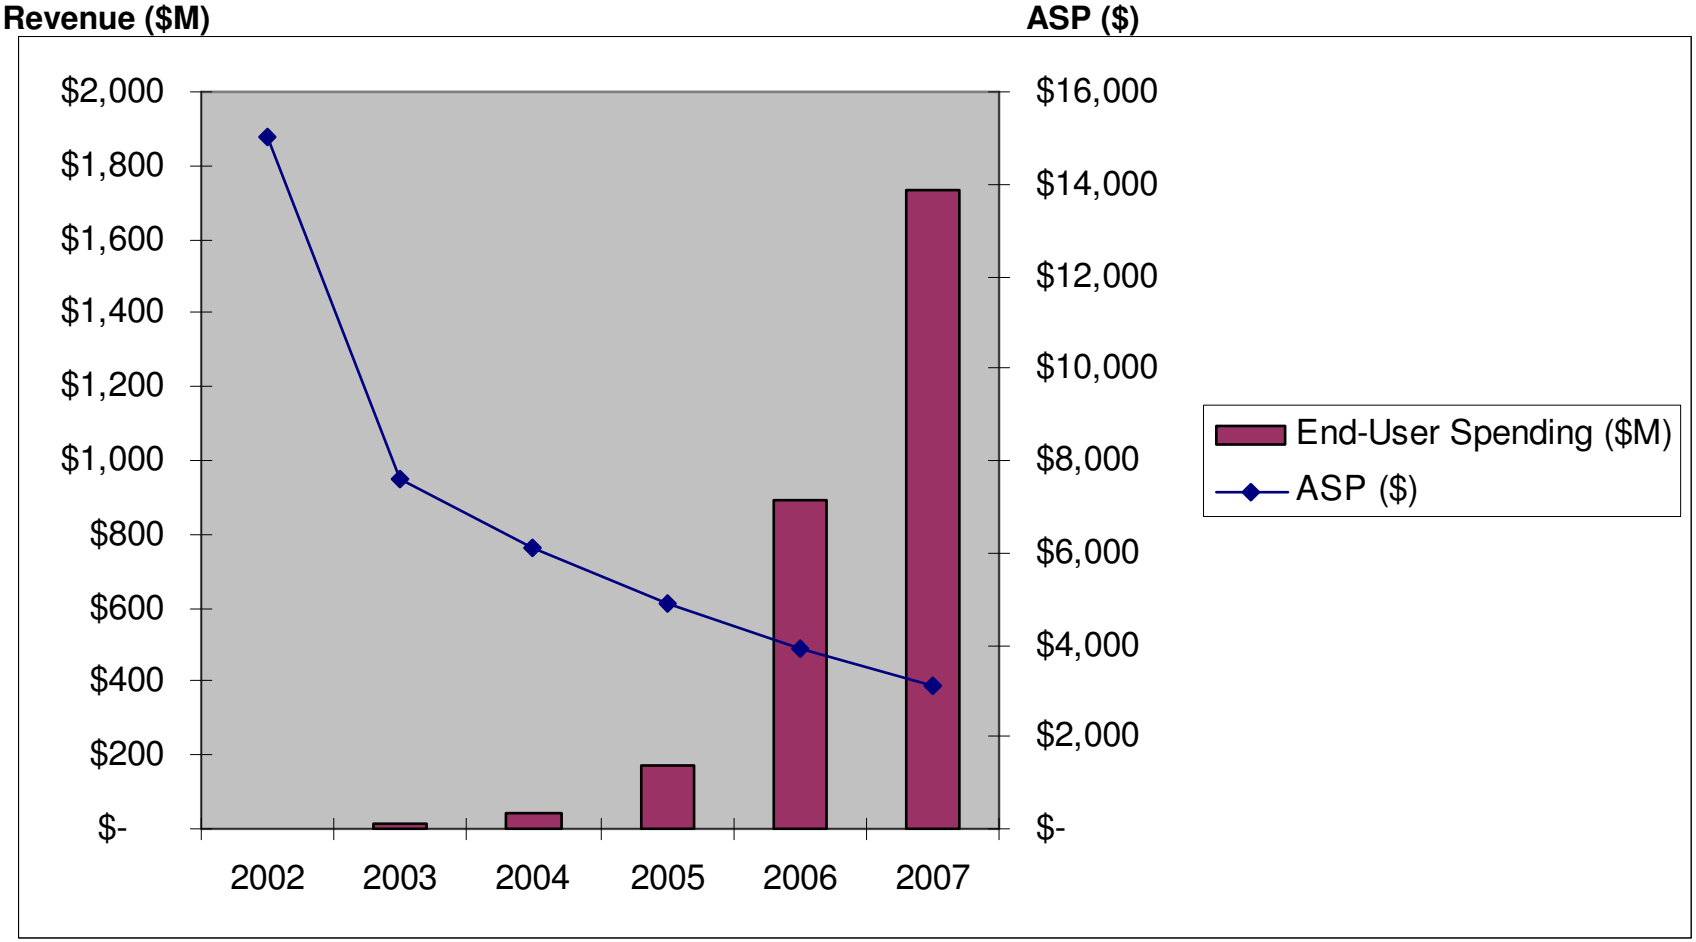

Source: Gartner/Dataquest, September 2003

The I/O fabric convergence

Fabric convergence in the Data Center

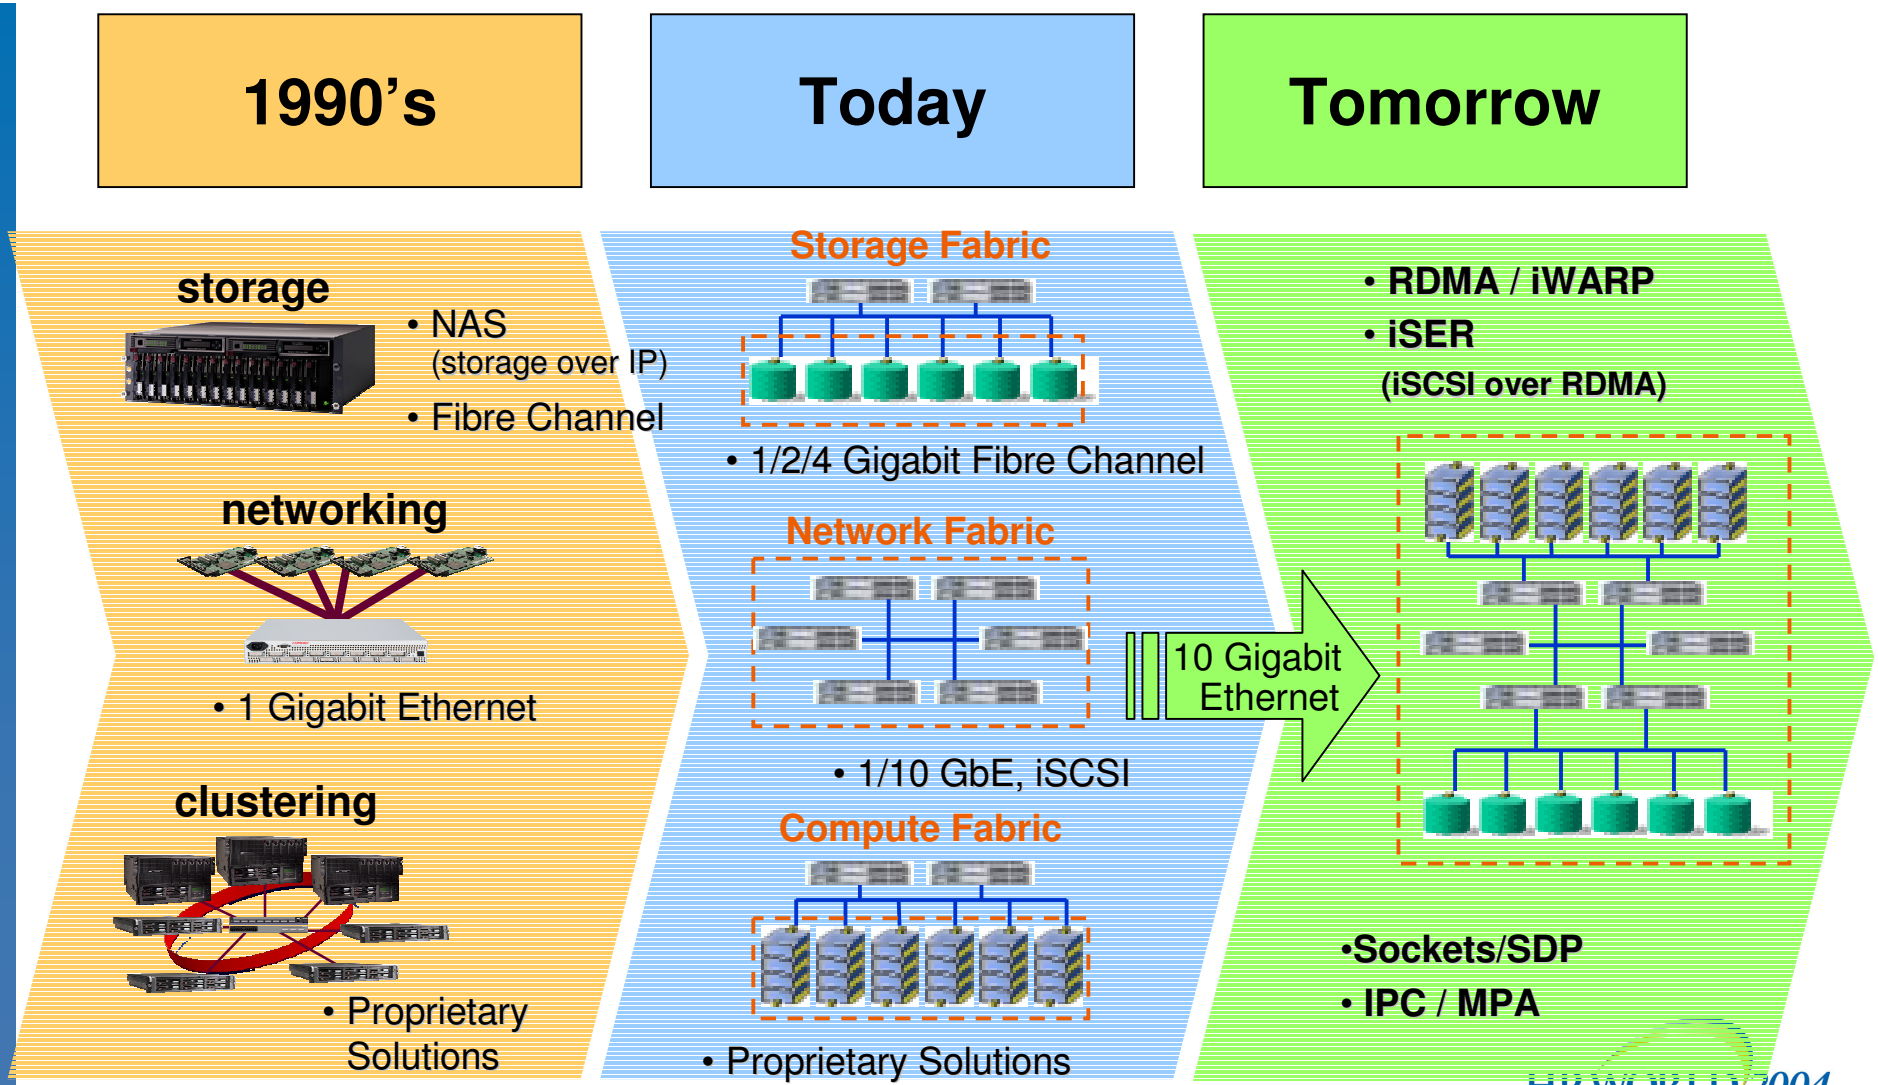

10 Gigabit Ethernet and RDMA (Remote Direct Memory Access)

- Direct Memory Placement
  - Eliminate processor copy
- Protocol offload
  - Eliminate protocol CPU, Memory overhead
  - TCP, iSCSI, SSL, IPSec, iWARP
- OS Bypass
  - Eliminate context switch, reduce latency
- RDMA
  - Combines the above in an industry standard, link independent manner

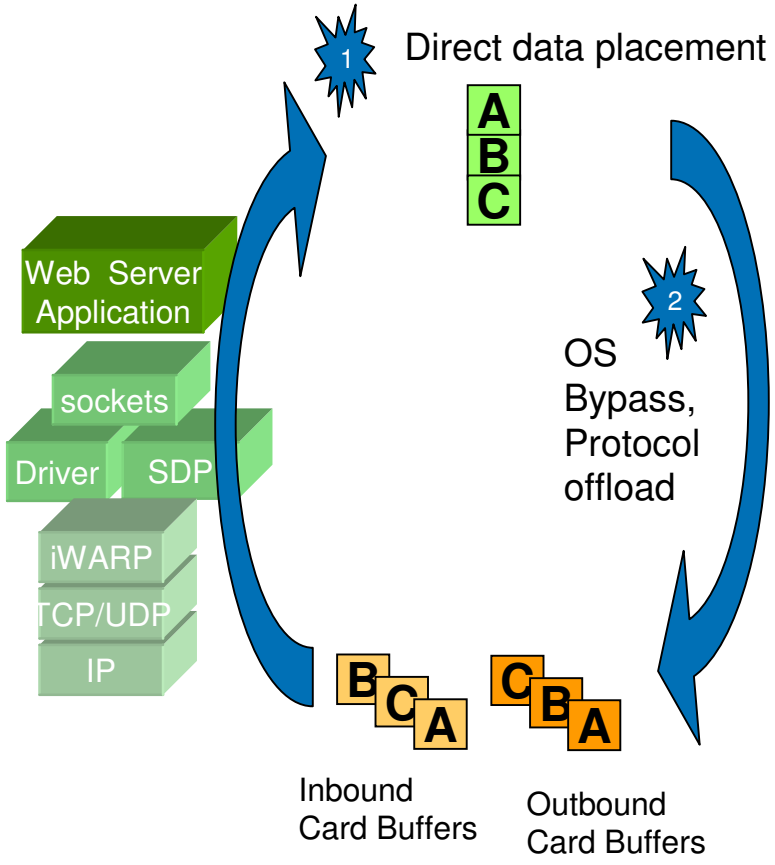

The marriage of 10 Gigabit Ethernet and RDMA will enable fabric standardization

Fabric convergence benefits

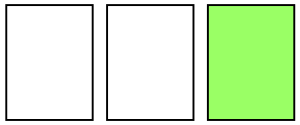

- Standardization, Simplification, Interoperability, Portability

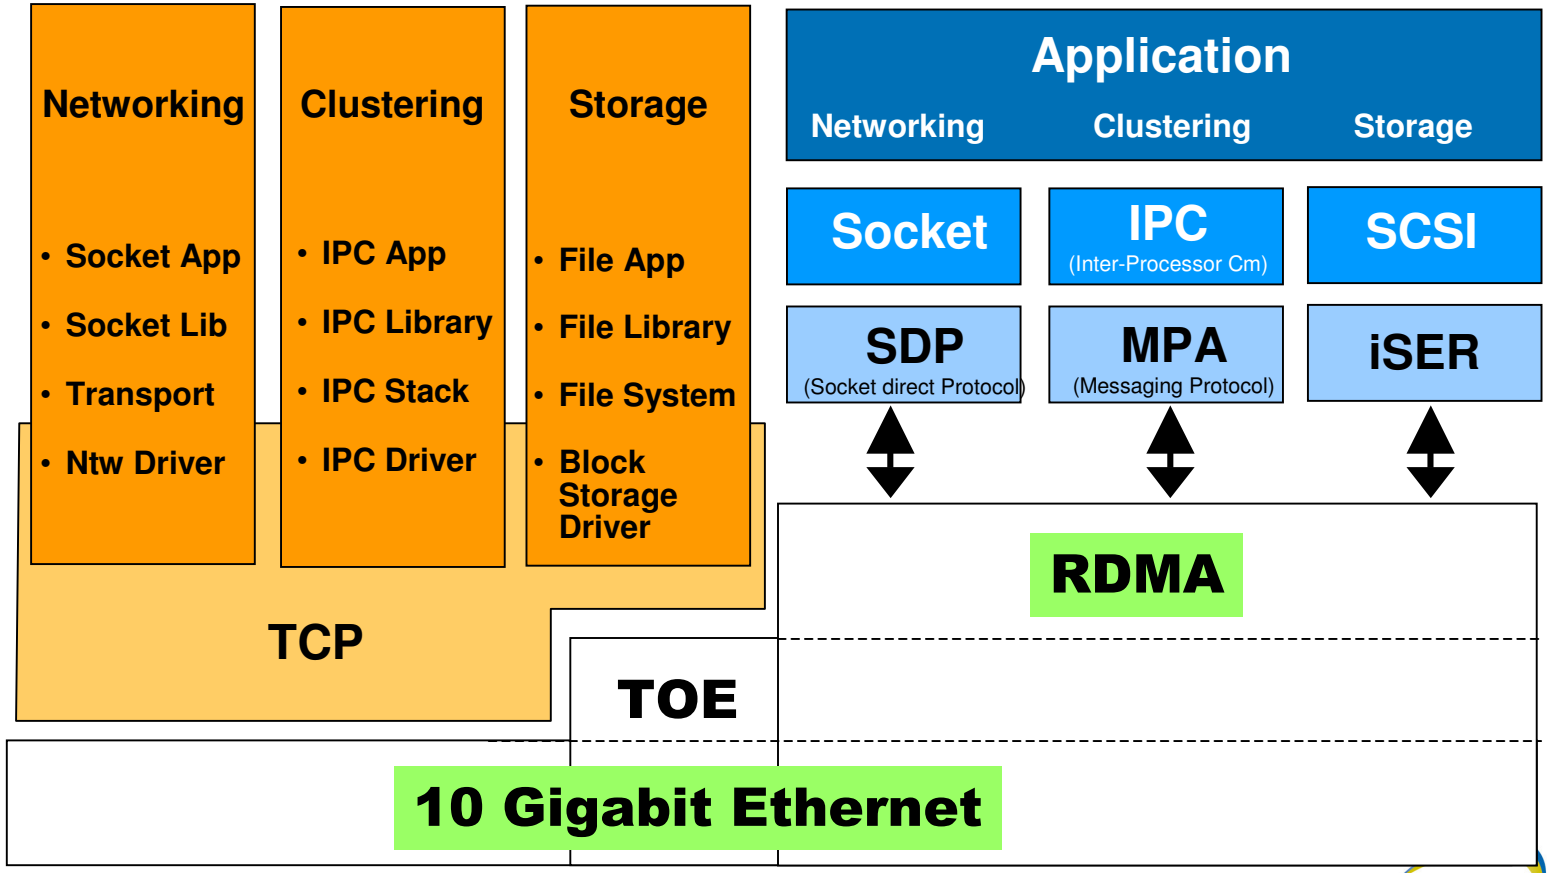

Current Architecture: “Spaghetti net”

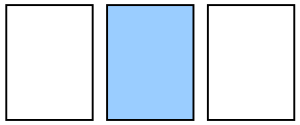

Today

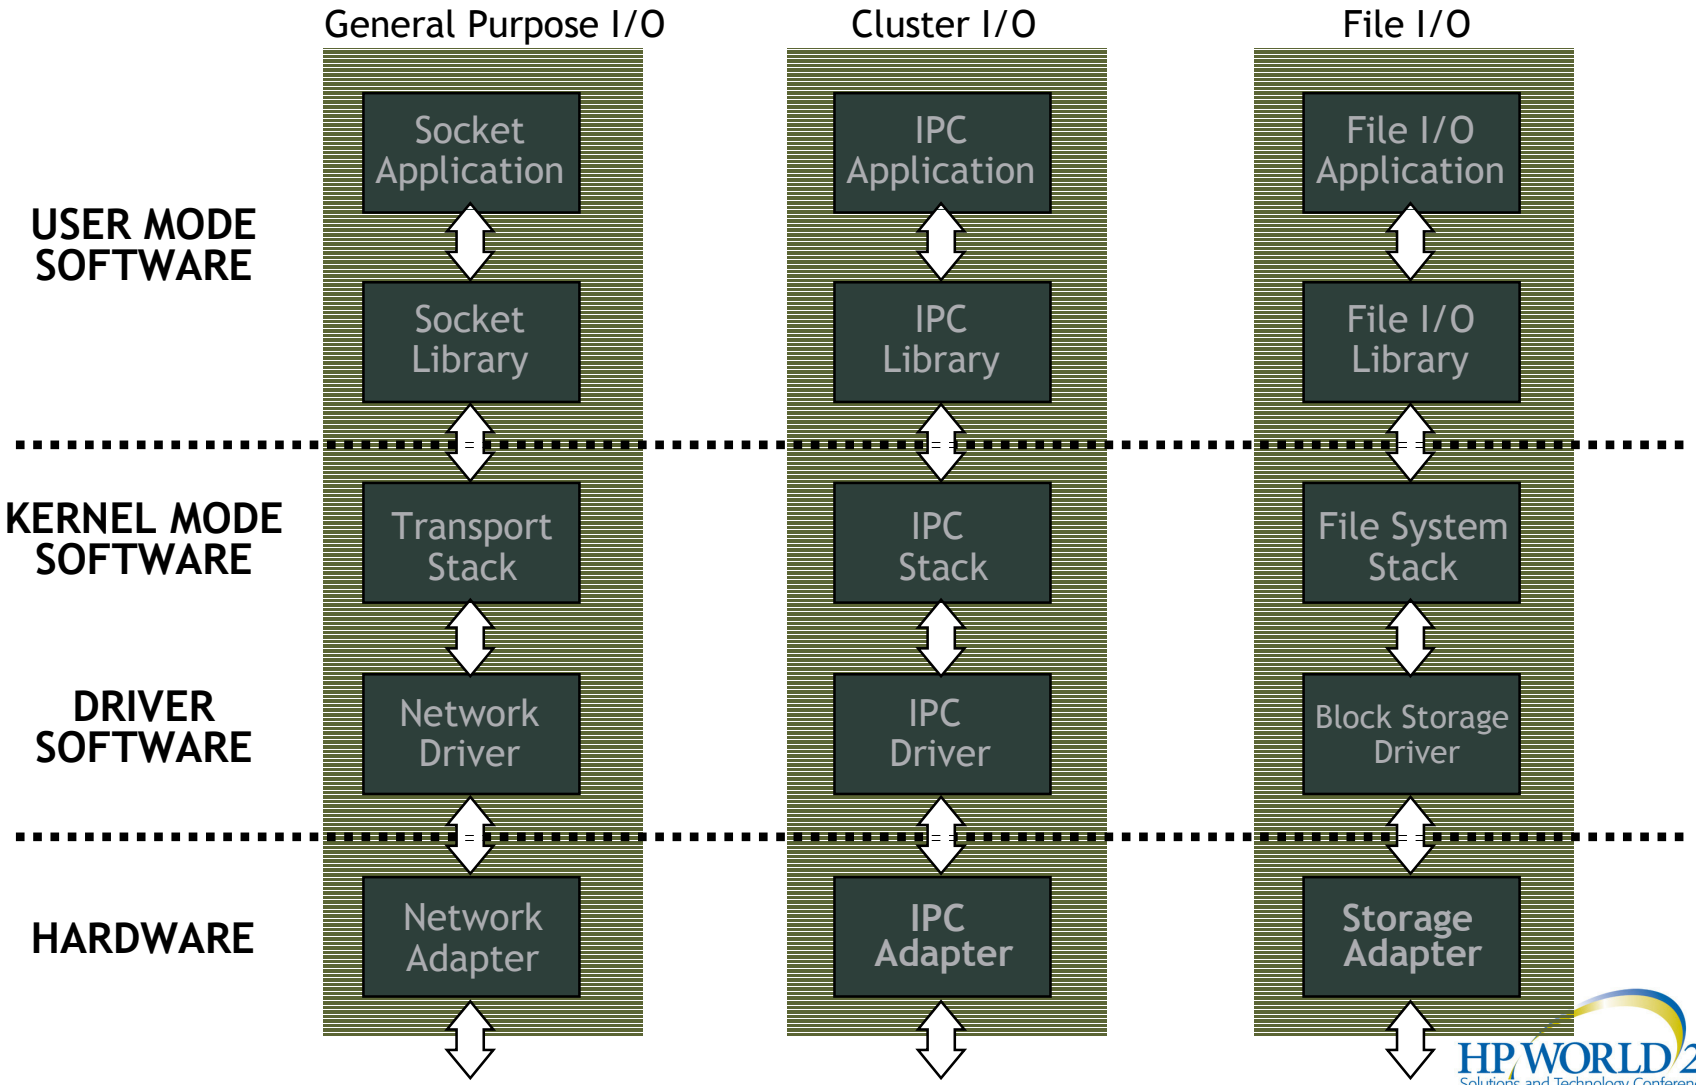

Converged Architecture: Single Fabric

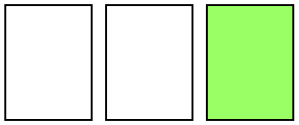

Tomorrow

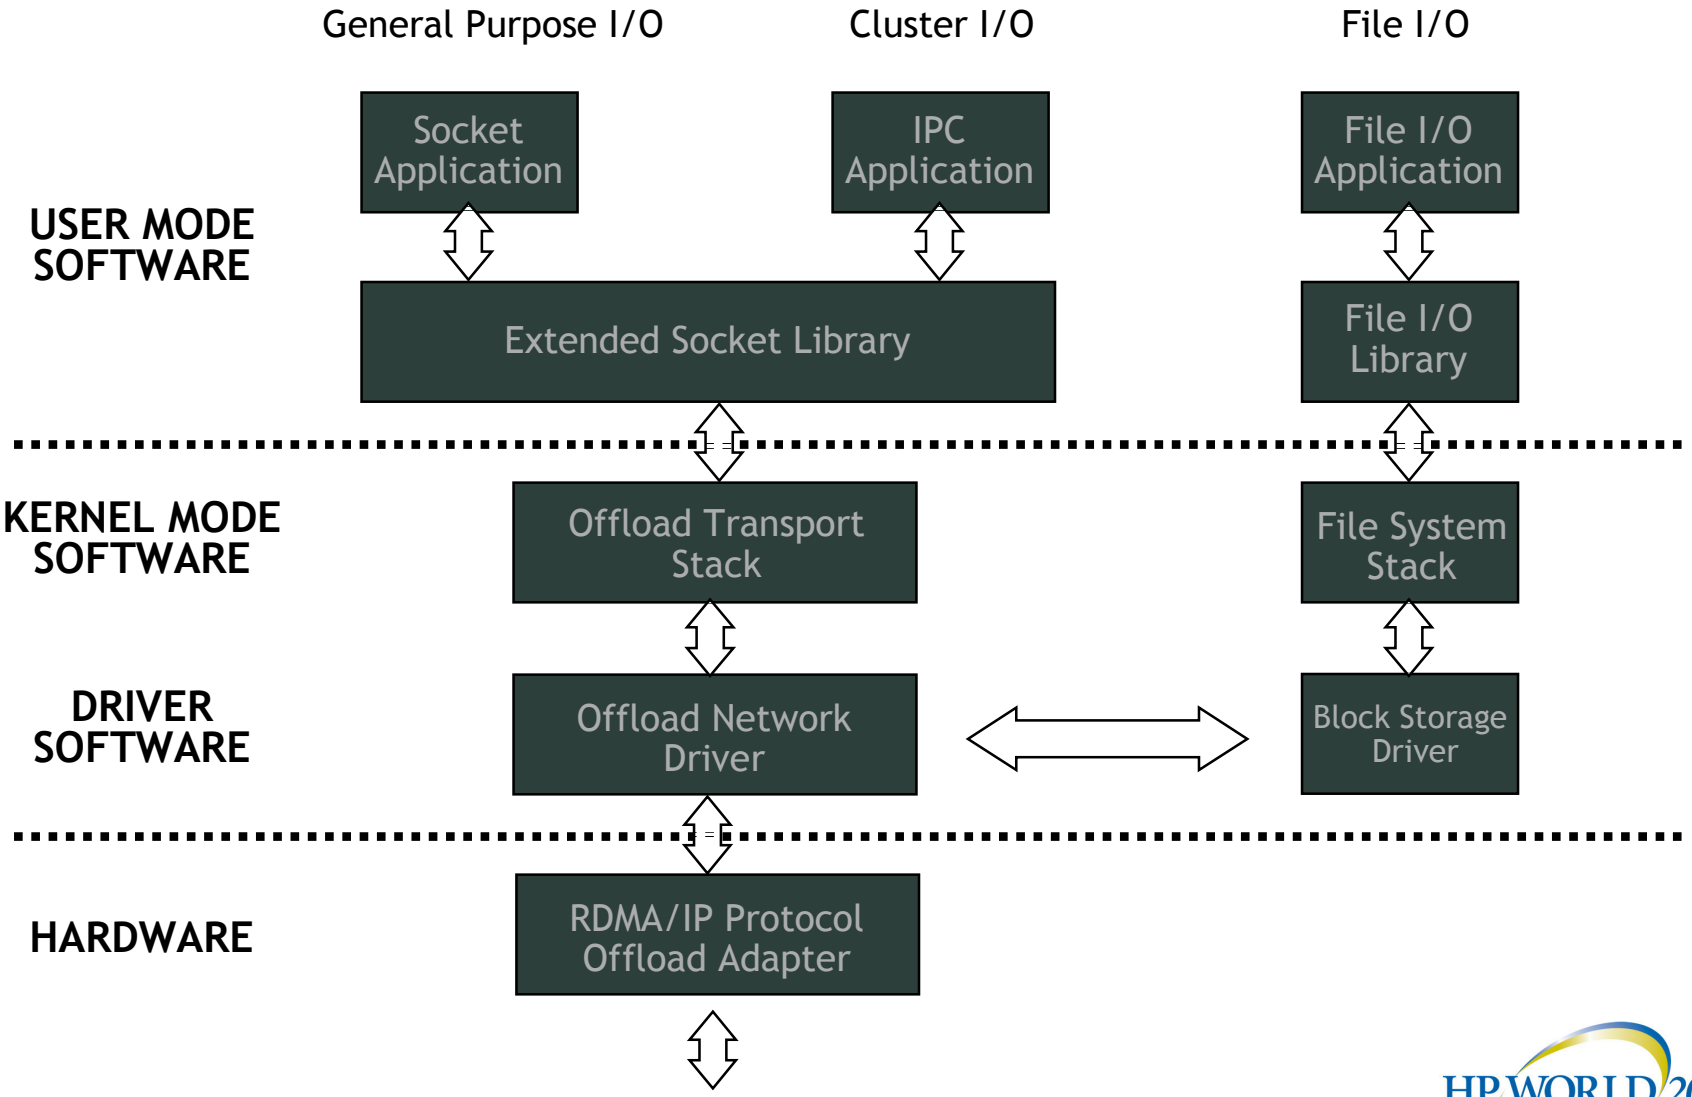

Dave Zabrowski

President and CEO – S2io, Inc.

HP WORLD 2004

Solutions and Technology Conference & Expo

Co-produced by:

interex

shared knowledge • shared power

encompass

AN HP USER GROUP