“All Grid – All the Time”

Stephen Perrenod, Ph.D.

Group Mgr., HPTC Marketing

Sun Microsystems

www.sun.com/hptc

Sun microsystems

We make the net work.

All Grid – All the Time

- You say you want a Revolution?
- N1, the Network Operating System
- Grid Computing: Vision & Strategy
- Compute Grids / Data Grids / Graphics Grids
- Examples: Grids & Partner Projects
- Future of Grid and Utility computing

“System” Is Changing, Again

Historical Cycle of Computing Industry, every 10 – 15 years

System	When	Tech	Scale
MF	'60s	Trans/Core	1
Minis	'70s	SSI/MSI	1
PC/WS	'80s	uP / LSI	1
SMP	'90s	VLSI / DRAM	100
Grid	'00s	Server-on-a-Chip/ Fiber Optics	1000's

NEW *BFWTS, N1, Grid, Utility, Horizontal...*

The Revolution

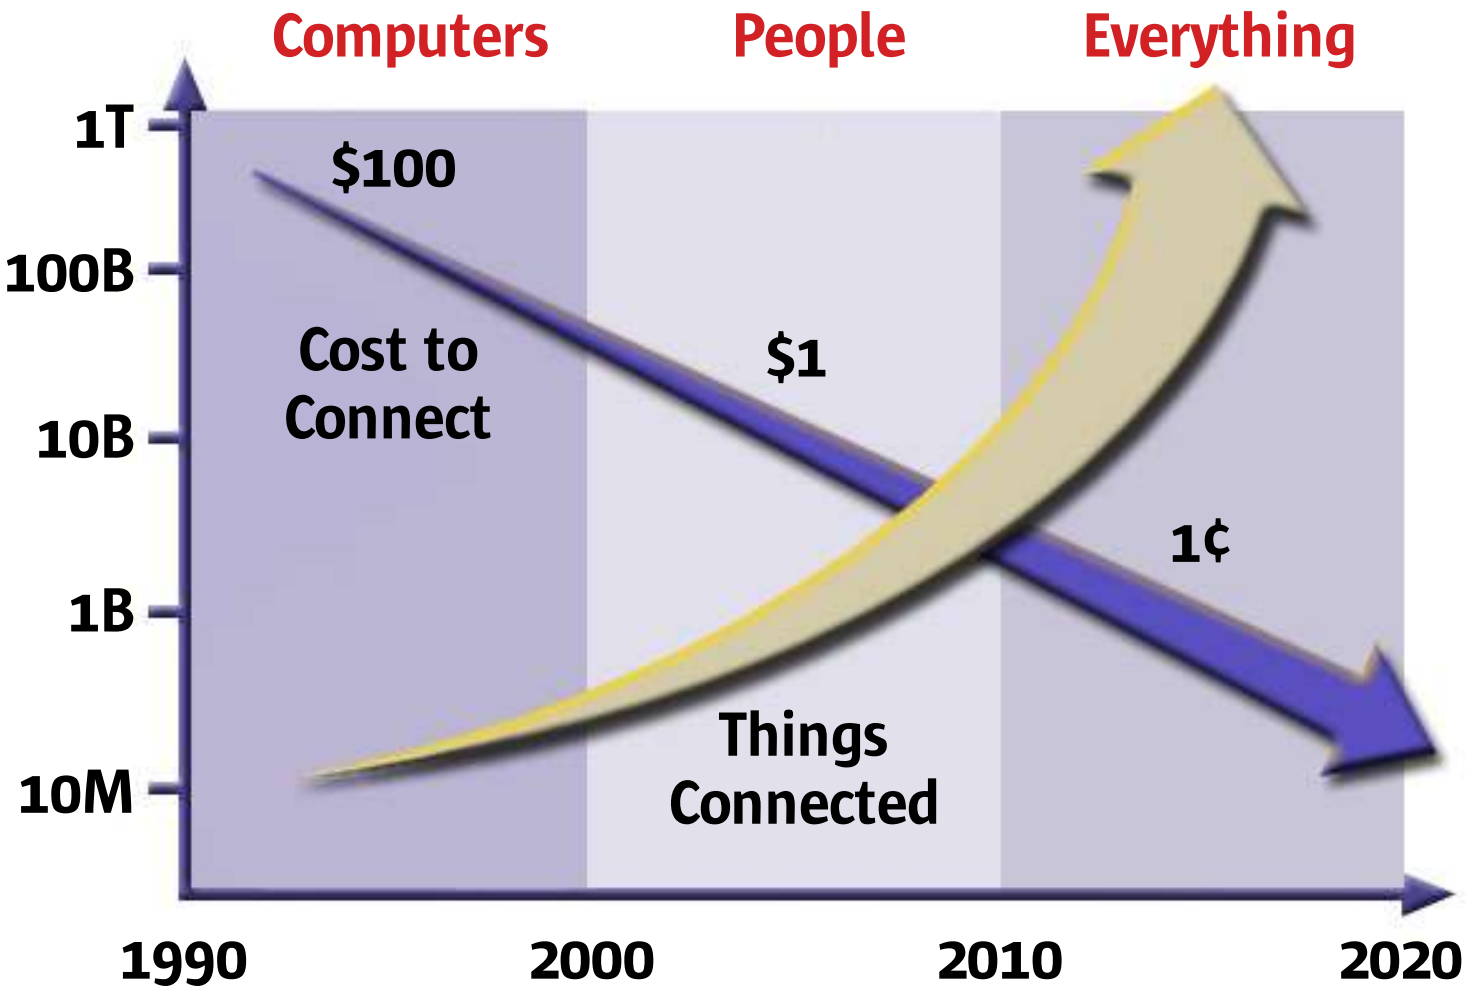

Frontiers of Automation

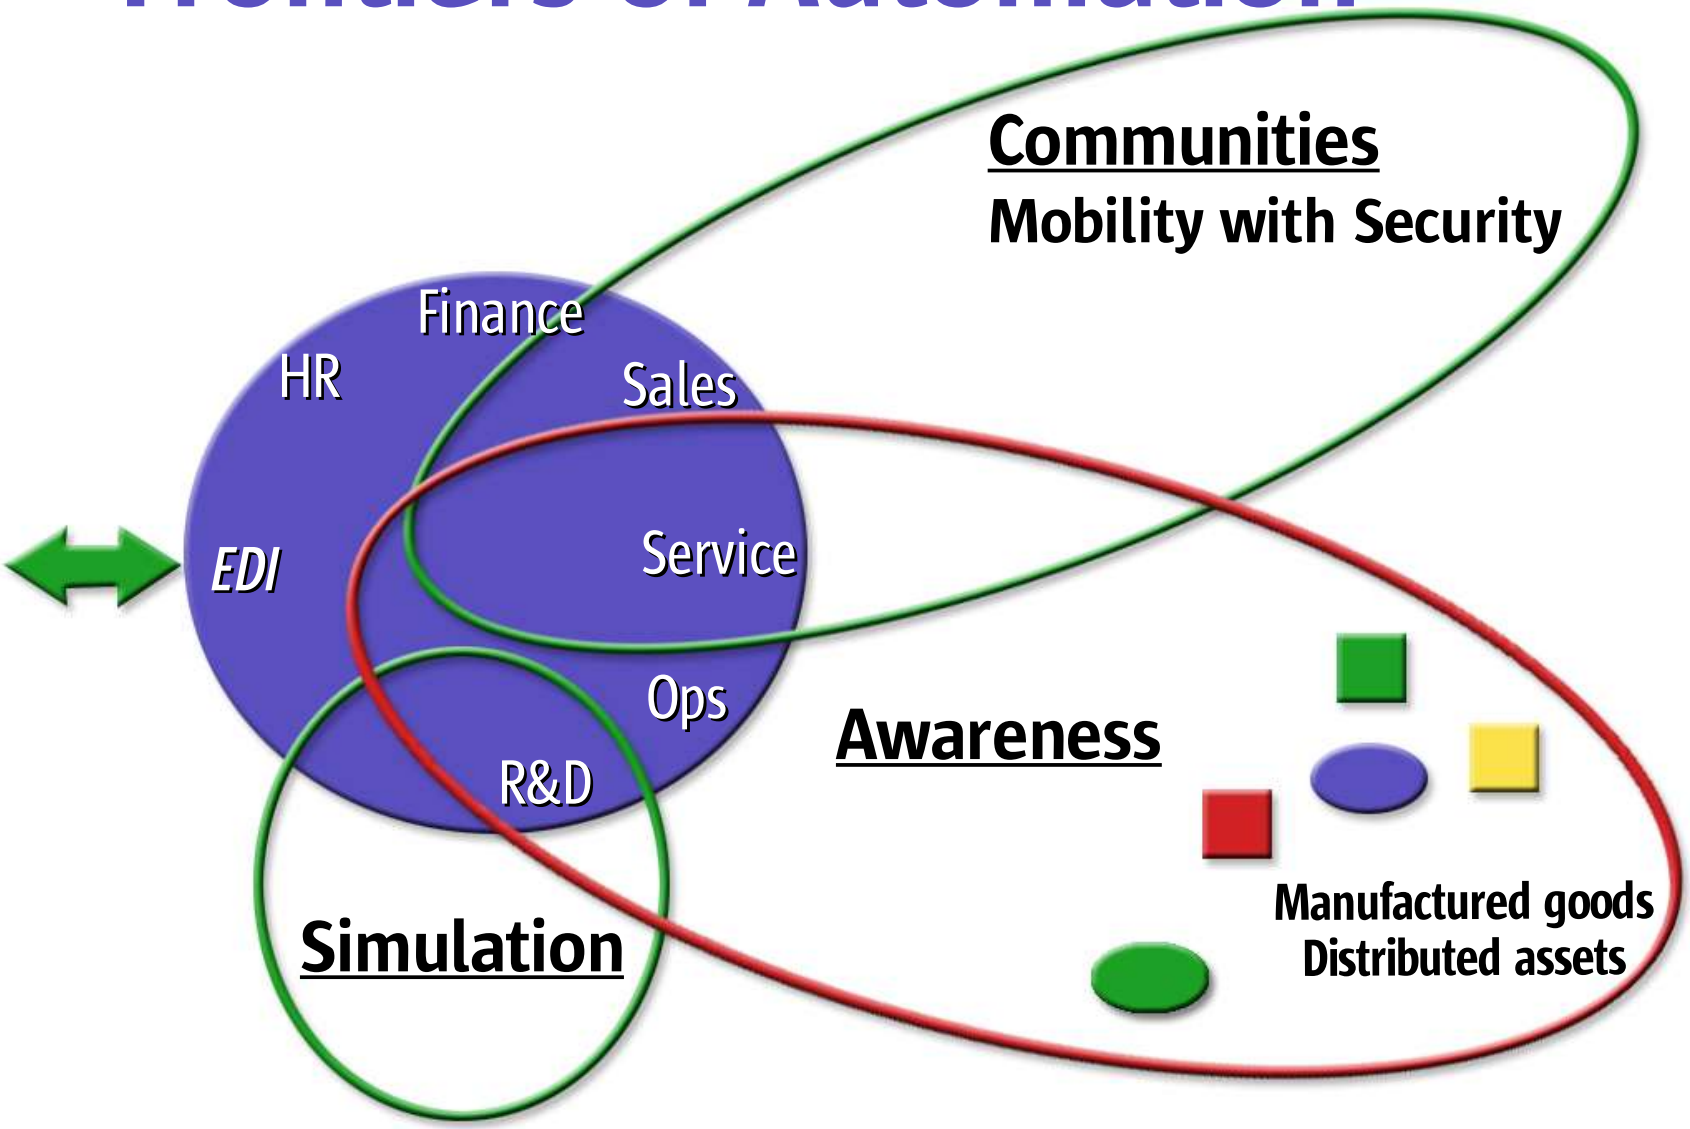

A New Meaning of "System"

What we did inside the F15K box...

We are doing to the network...

- True scalability:
  Add performance without adding management complexity!
- "Soft configuration" and " Soft cabling"
- Multiple, secure domains
- But with a big difference:
  - Heterogeneous elements
  - Network becomes like SMP backplane

Old systems are components in the new one

Server connected to networks

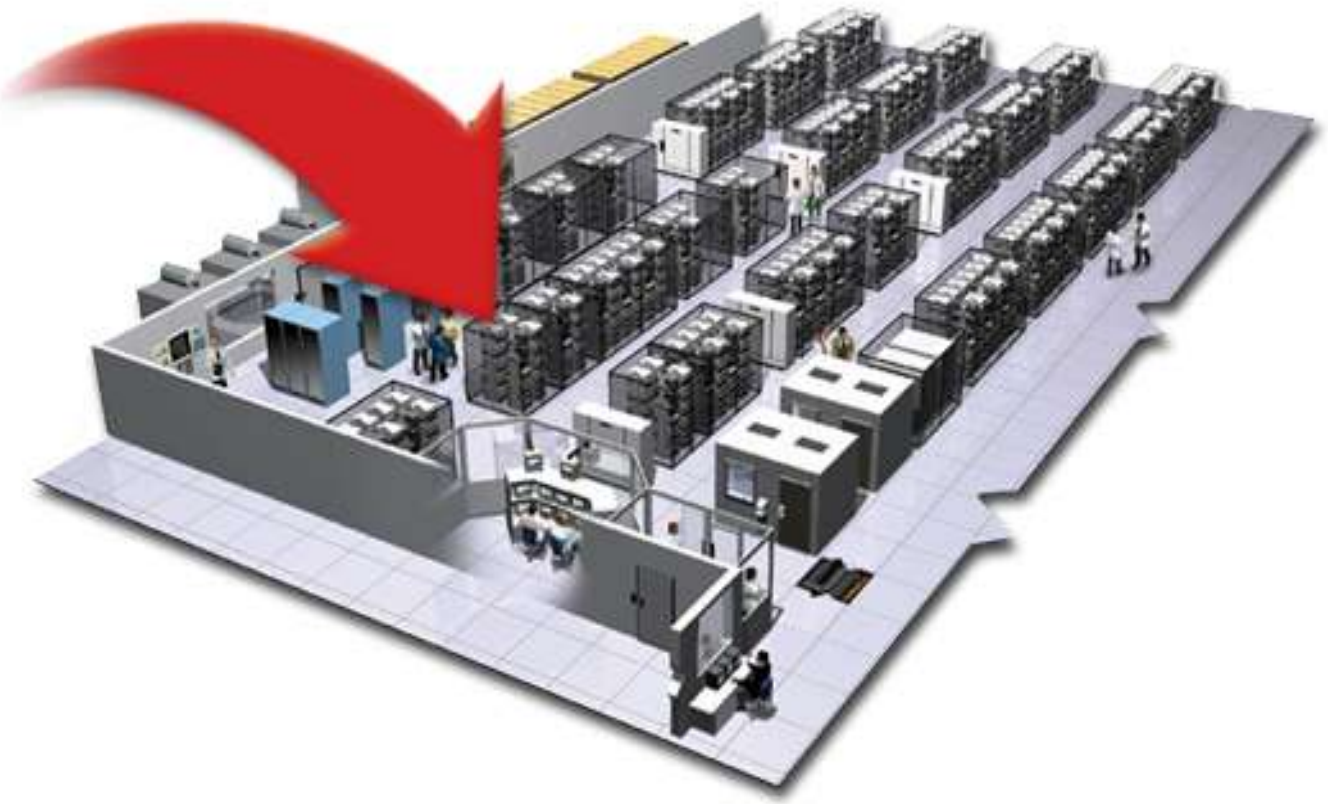

What is O/S for this new system?

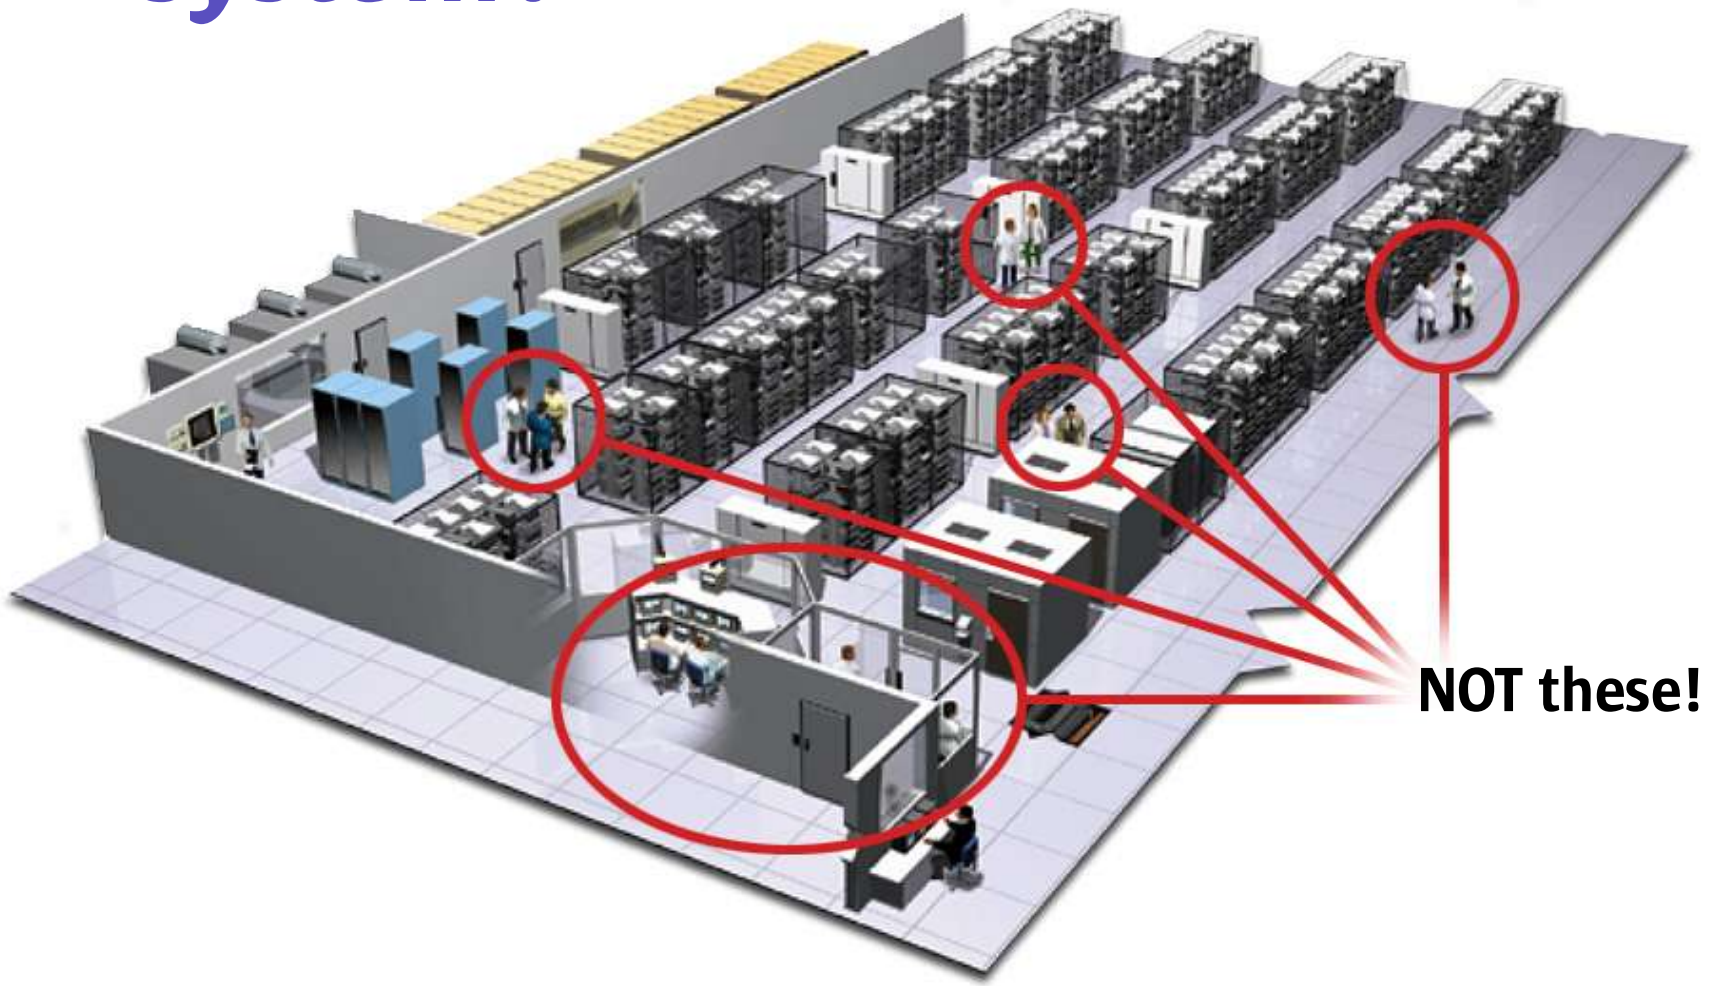

The New System

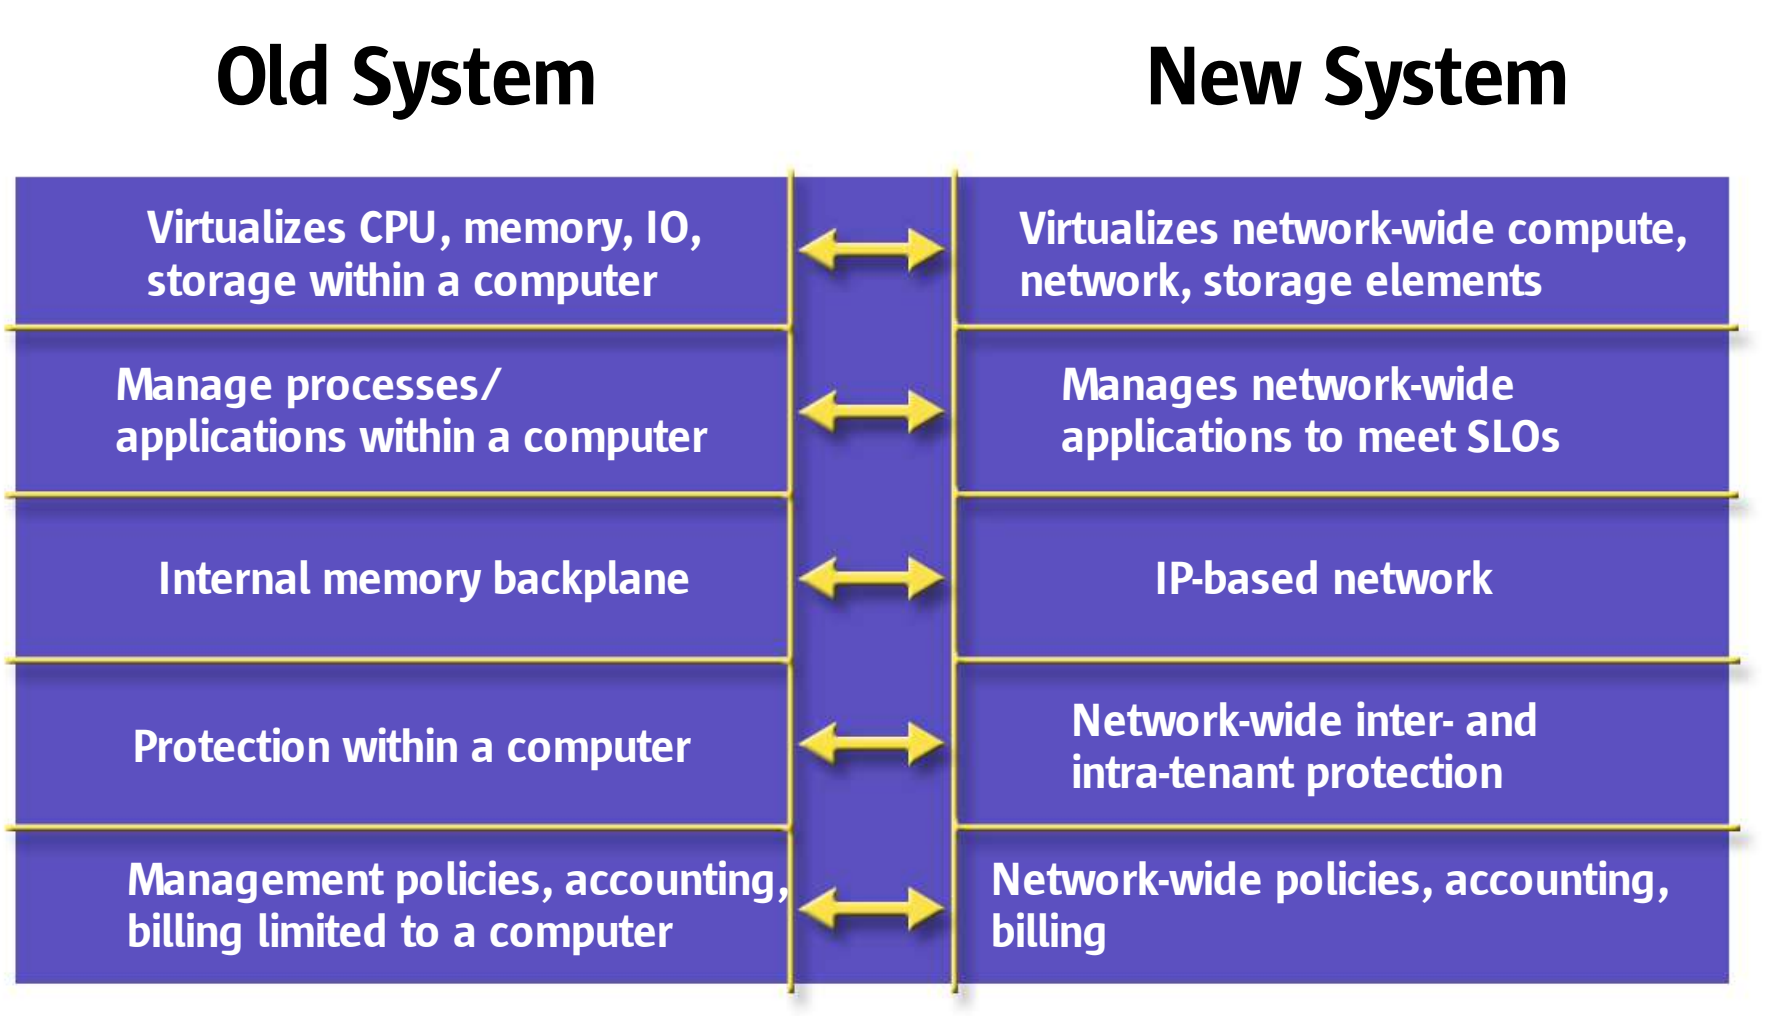

Traditional OS Virtualization

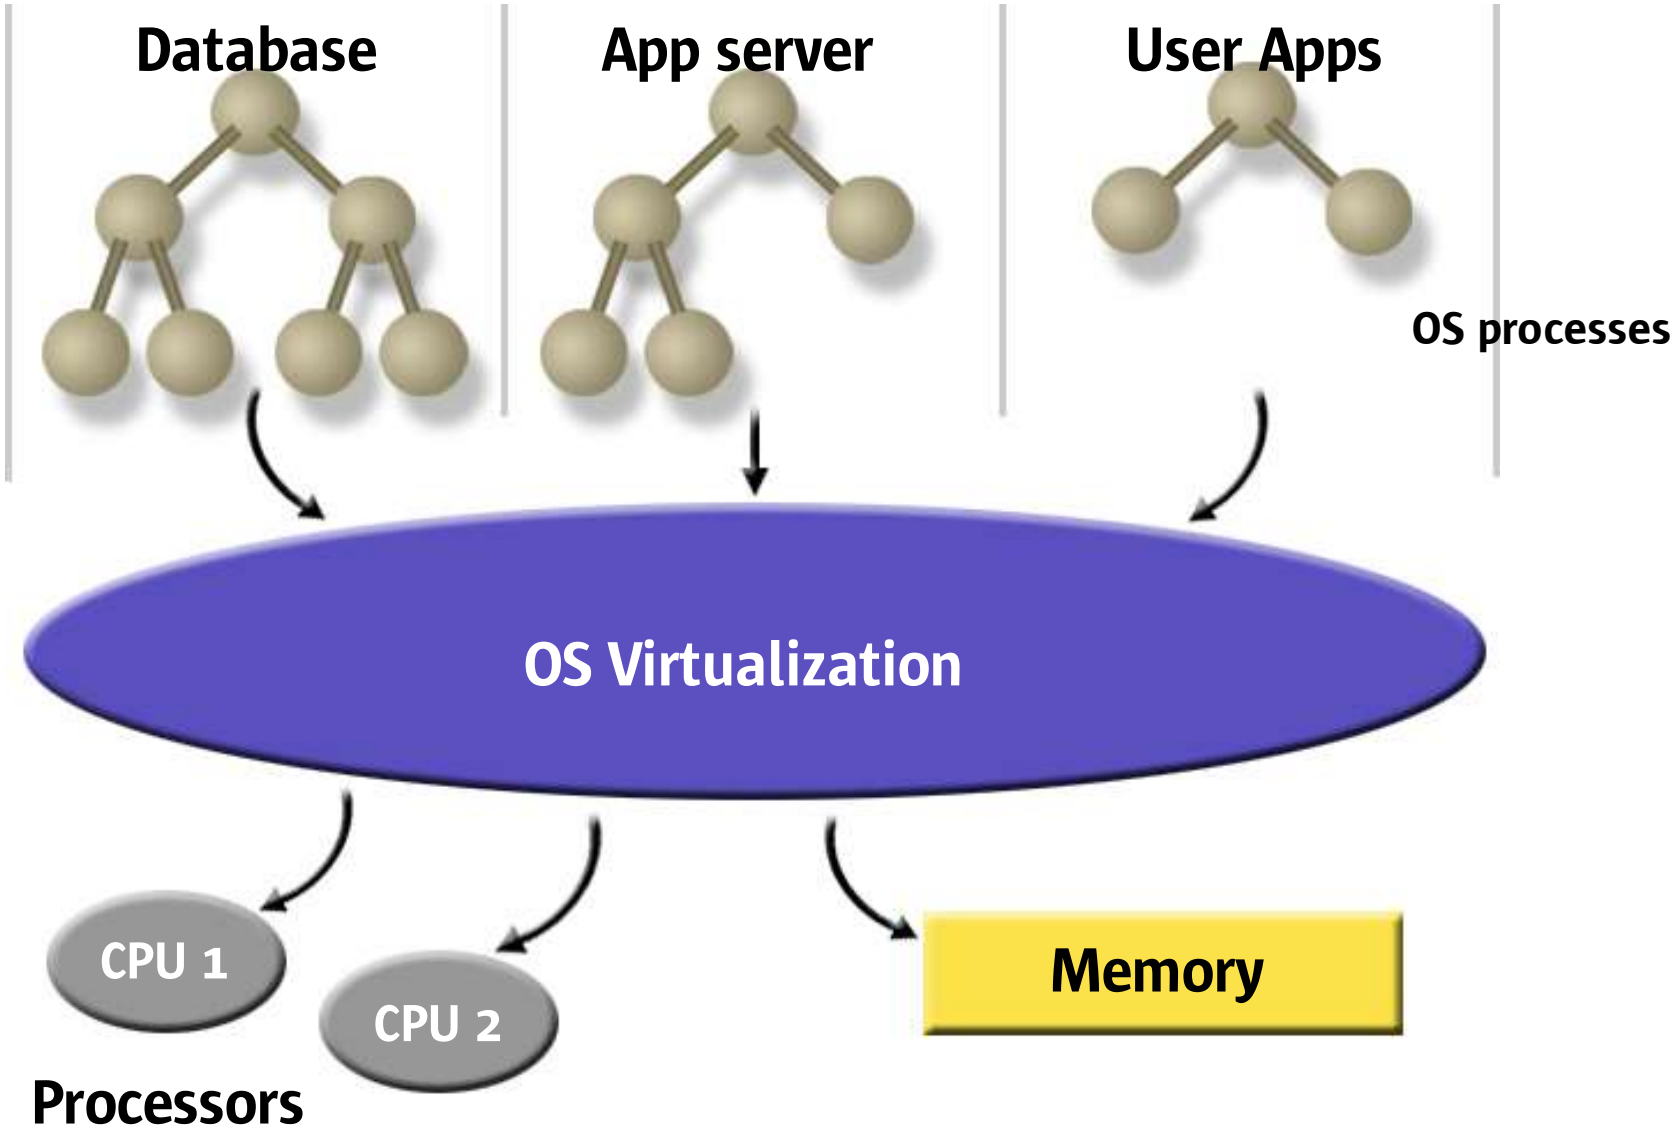

The New Software

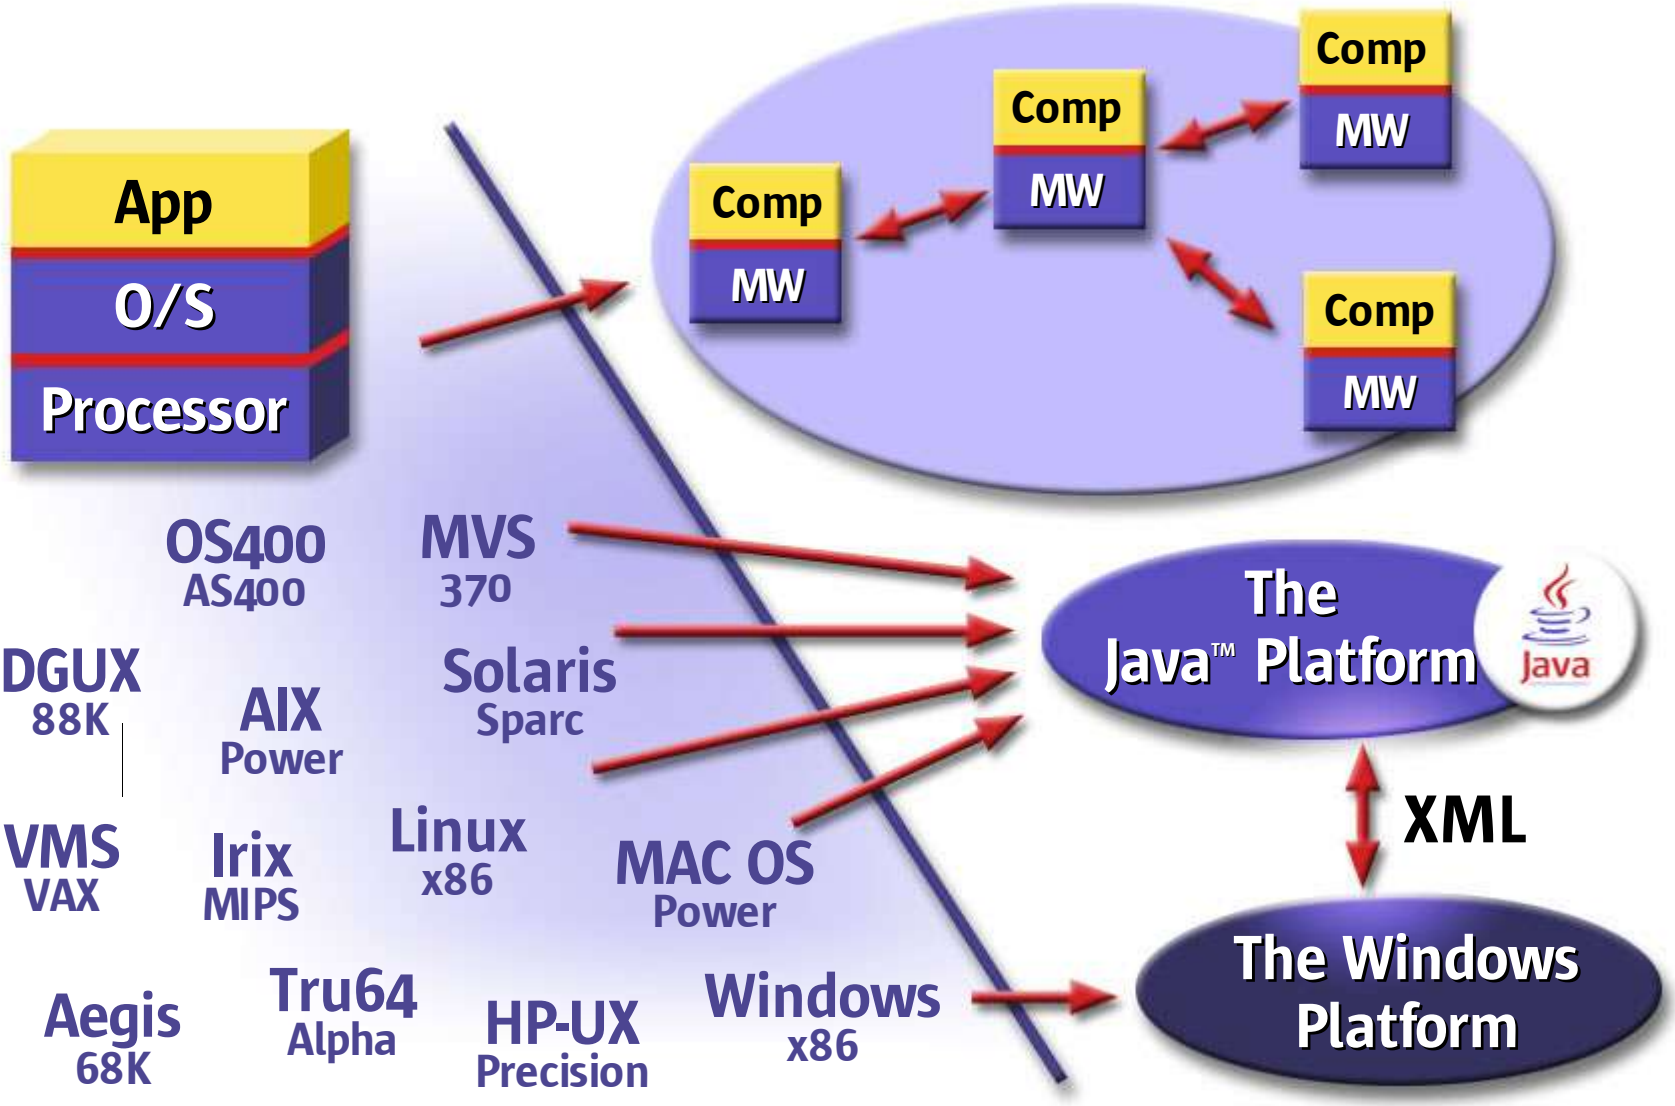

N1 Macro-Virtualization

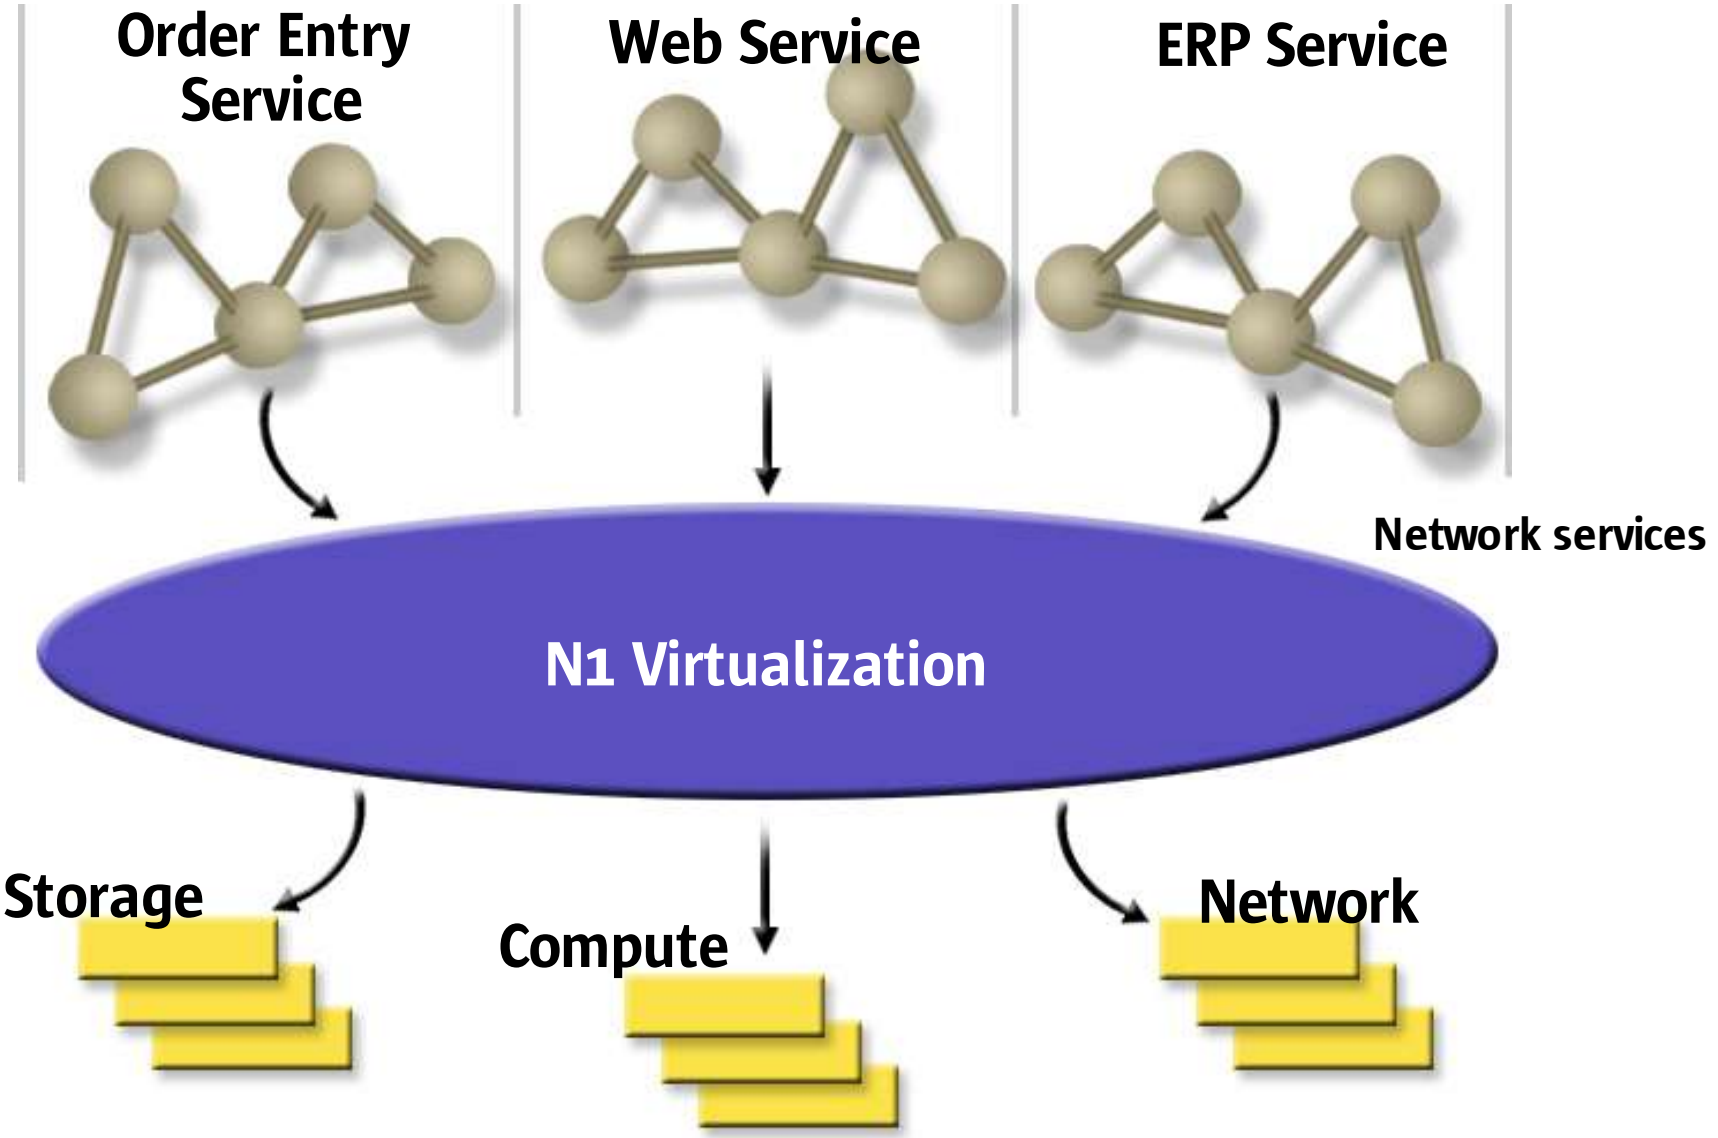

The N1 Effect on Efficiency

Radical improvement in costs & uptime

- Sys Utilization
- Server/Admin
- Terabytes/DBA
- Ports/Admin
- Availability
- Time to Deploy

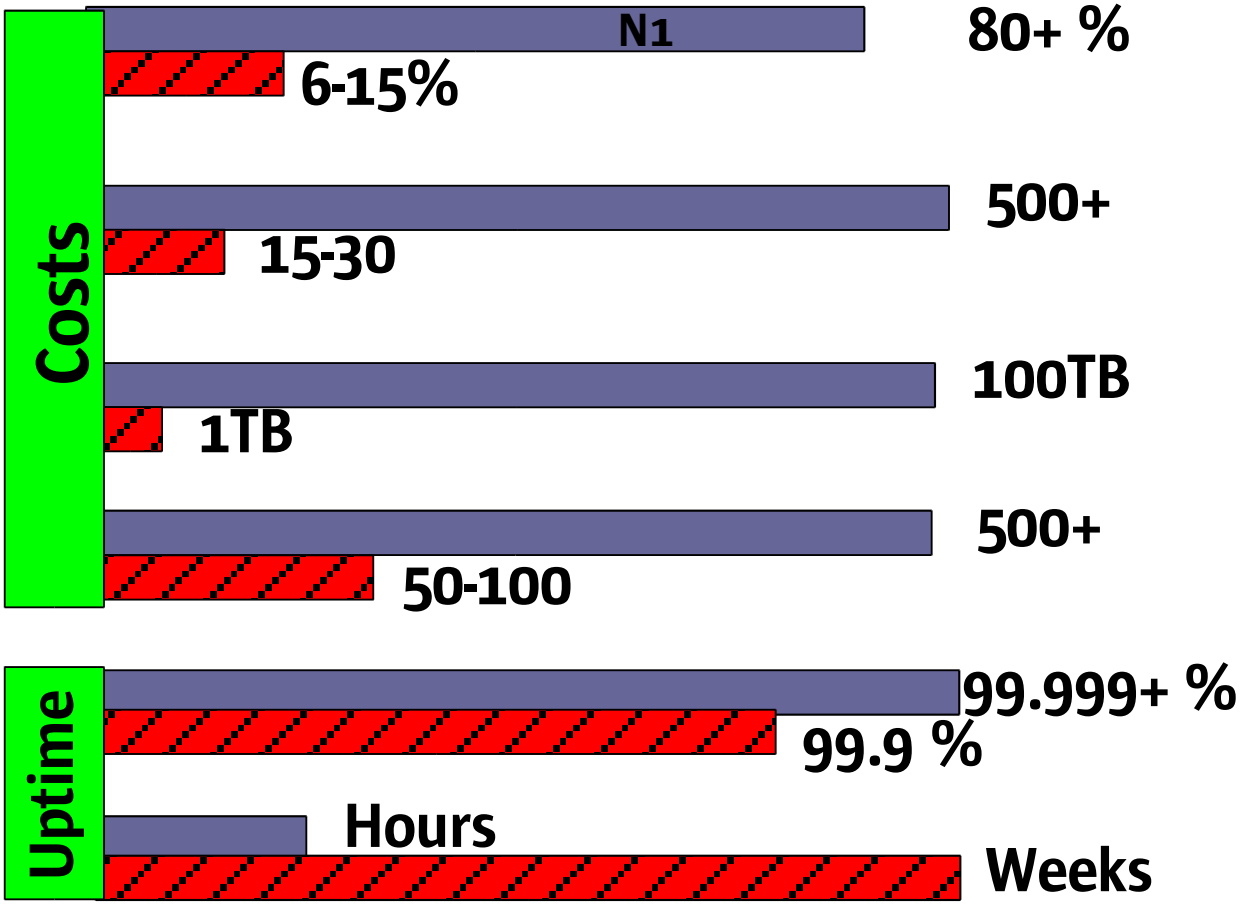

Java Web Services: The New API

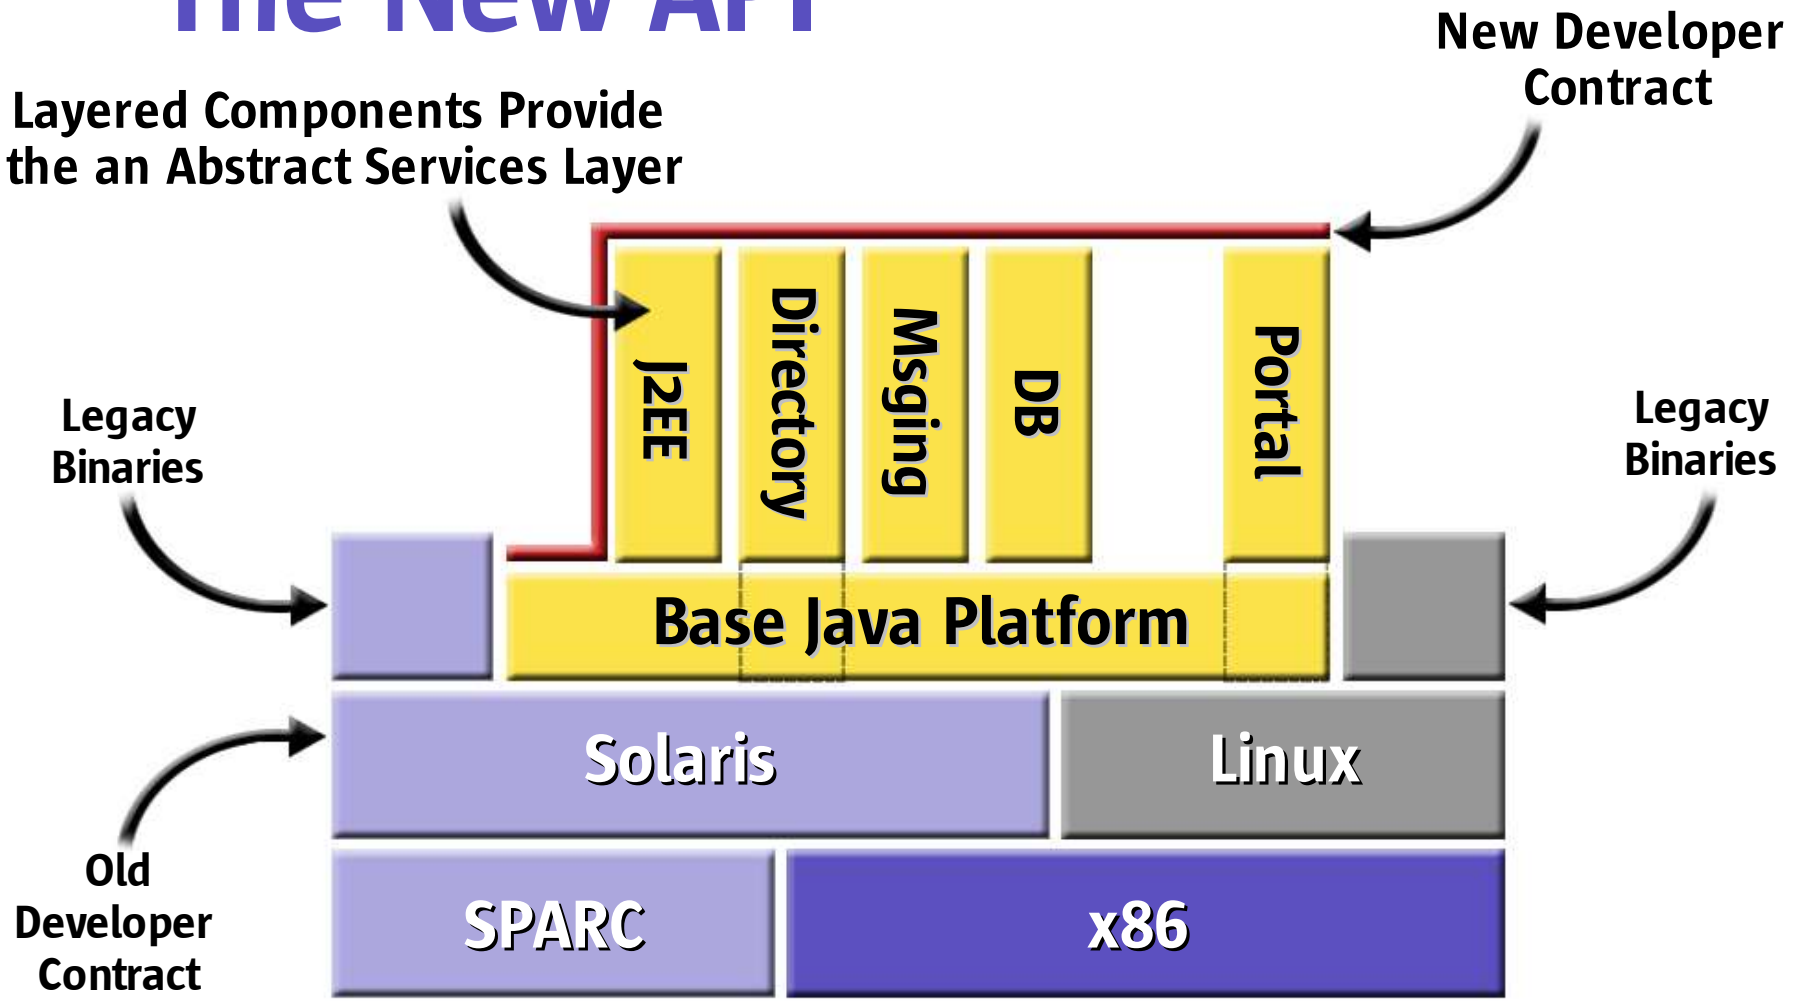

Sun's HPTC Strategy: "All Grid, All the Time"

Data Grids

Compute Grids

Graphics Grids

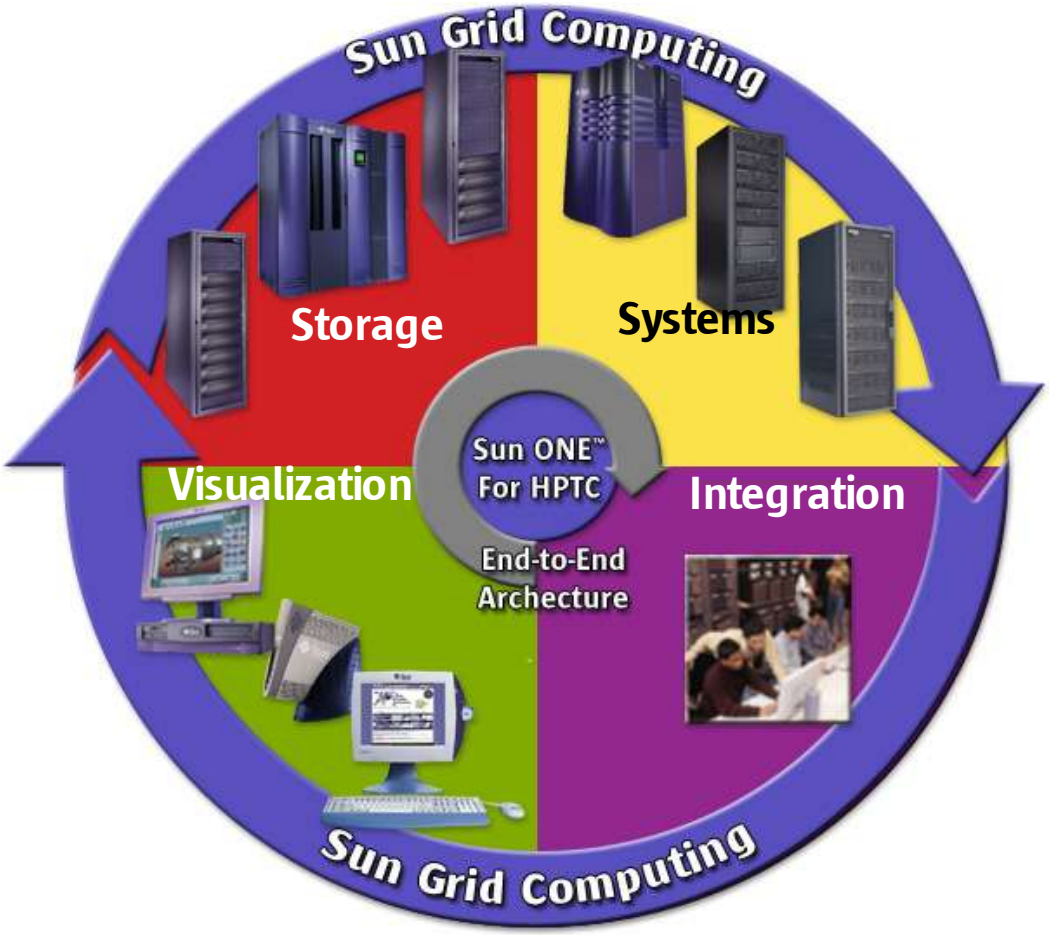

Example: Sun Grid Services Environment

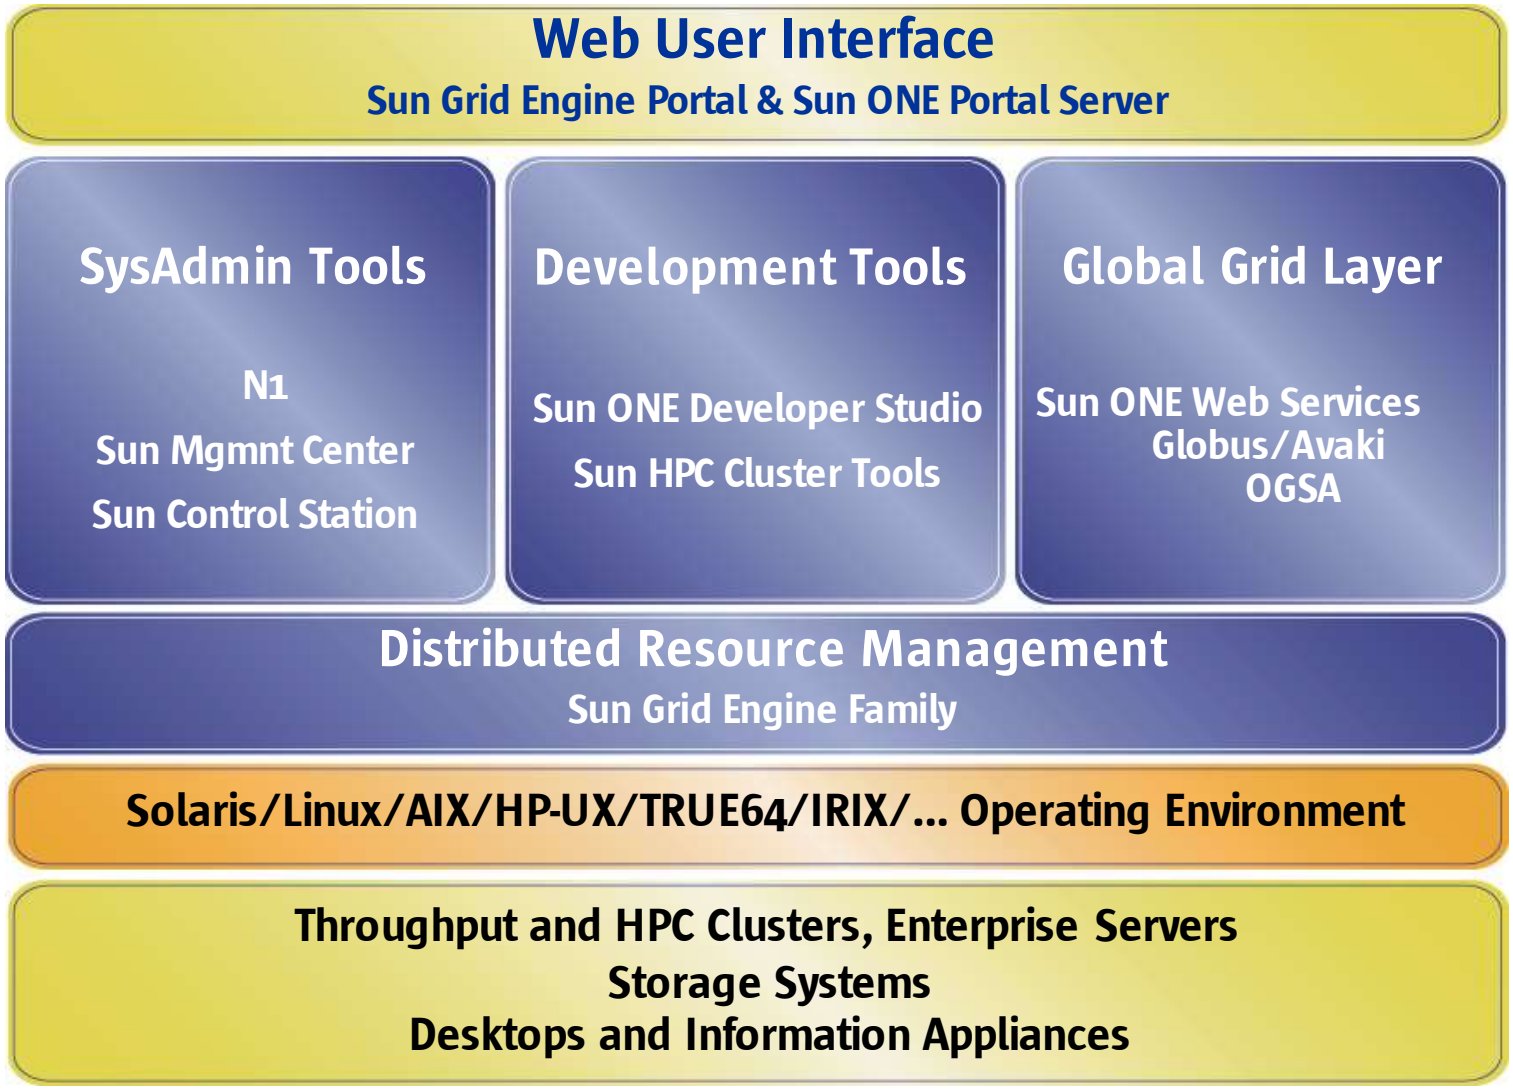

Sun Grid Engine Family

Distributed Resource Management in Cluster & Enterprise Grids

- Multi-platform, open source, standards
  - 8000+ grids today (departmental, enterprise, global grids)
  51% Solaris, 25% Linux, 24% Mix
- Sun Grid Engine, SGE, free Web downloads for Solaris & Linux
  - Identifies best-suited, least loaded resource for your work
  - Queuing, prioritizing, scheduling
- Sun Grid Engine, Enterprise Edition
  - Equitable, enforceable sharing between groups and projects
  - Alignment of resources with business goals via policies

HPC Grids vs Commercial Grids

HPC Research & Eng. Grids
• High Productivity • High Performance • High Performance • Collaboration • Interpretation

Commercial Grids
• Scalability • Manageability • Availability • Reliability • Transparency

HPTC Market:

Moving to Clusters

Grand Challenge, Extreme Scale:
ASCII, HPCS, Earth Simulator

High Performance Community:
Research Capability, Productivity

Technical Computing:
Production throughput and embarassingly parallel

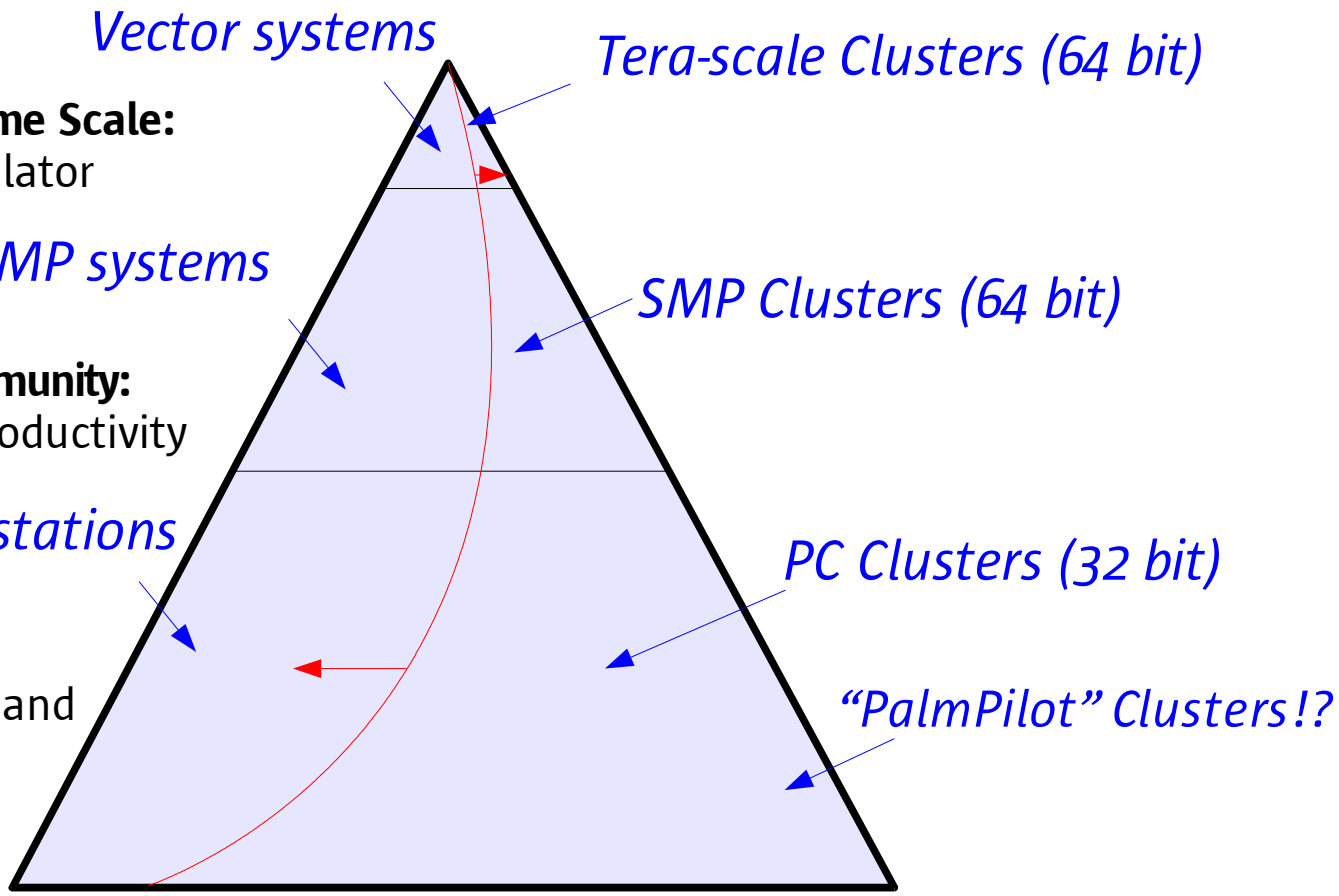

Commercial Grid Requirements

- Performance, robustness, manageability, availability
- downtime unacceptable
- purchased applications, standardized solutions
- tightly integrated with other enterprise systems
- focus on enterprise software, databases and transactions
- quantifiable ROI
- *Market mostly unknown, on the verge of strong growth*

GRID as an Operational Concept

Workflow Service Definition & Mapping

Capability Computing Services

Data Storage, Management and Access Services

Capacity Computing Services

Collaboration, Graphics, and Visualization Services

UK e-Science Grid

$ 180 & 180 Mio in 3 & 3 years for science and engineering

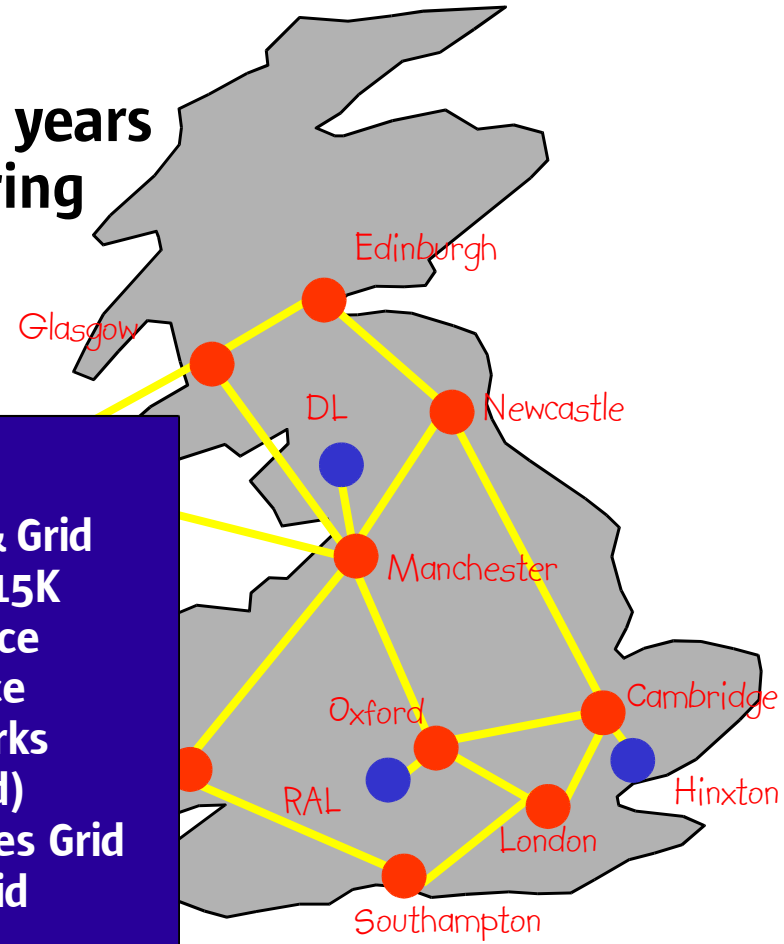

Our Grid Centers in UK:
Edinburgh EPCC, Sun CoE HPC & Grid
Cambridge, 2TeraFlops 10 SF15K
Oxford, Computational Finance
London IC, Sun CoE e-Science
London UCL, Sun CoE Networks
Manchester, MyGrid (BioGrid)
Leads, Sheffield, York: White Roses Grid
Durham: Cosmology Engine Grid
....

Example: White Rose Grid in England

- Leeds, York + Sheffield Universities
- Deliver stable, well-managed HPC resources supporting multi-disiplinary research
- Deliver a Metropolitan Grid across the Universities

WRG Hardware

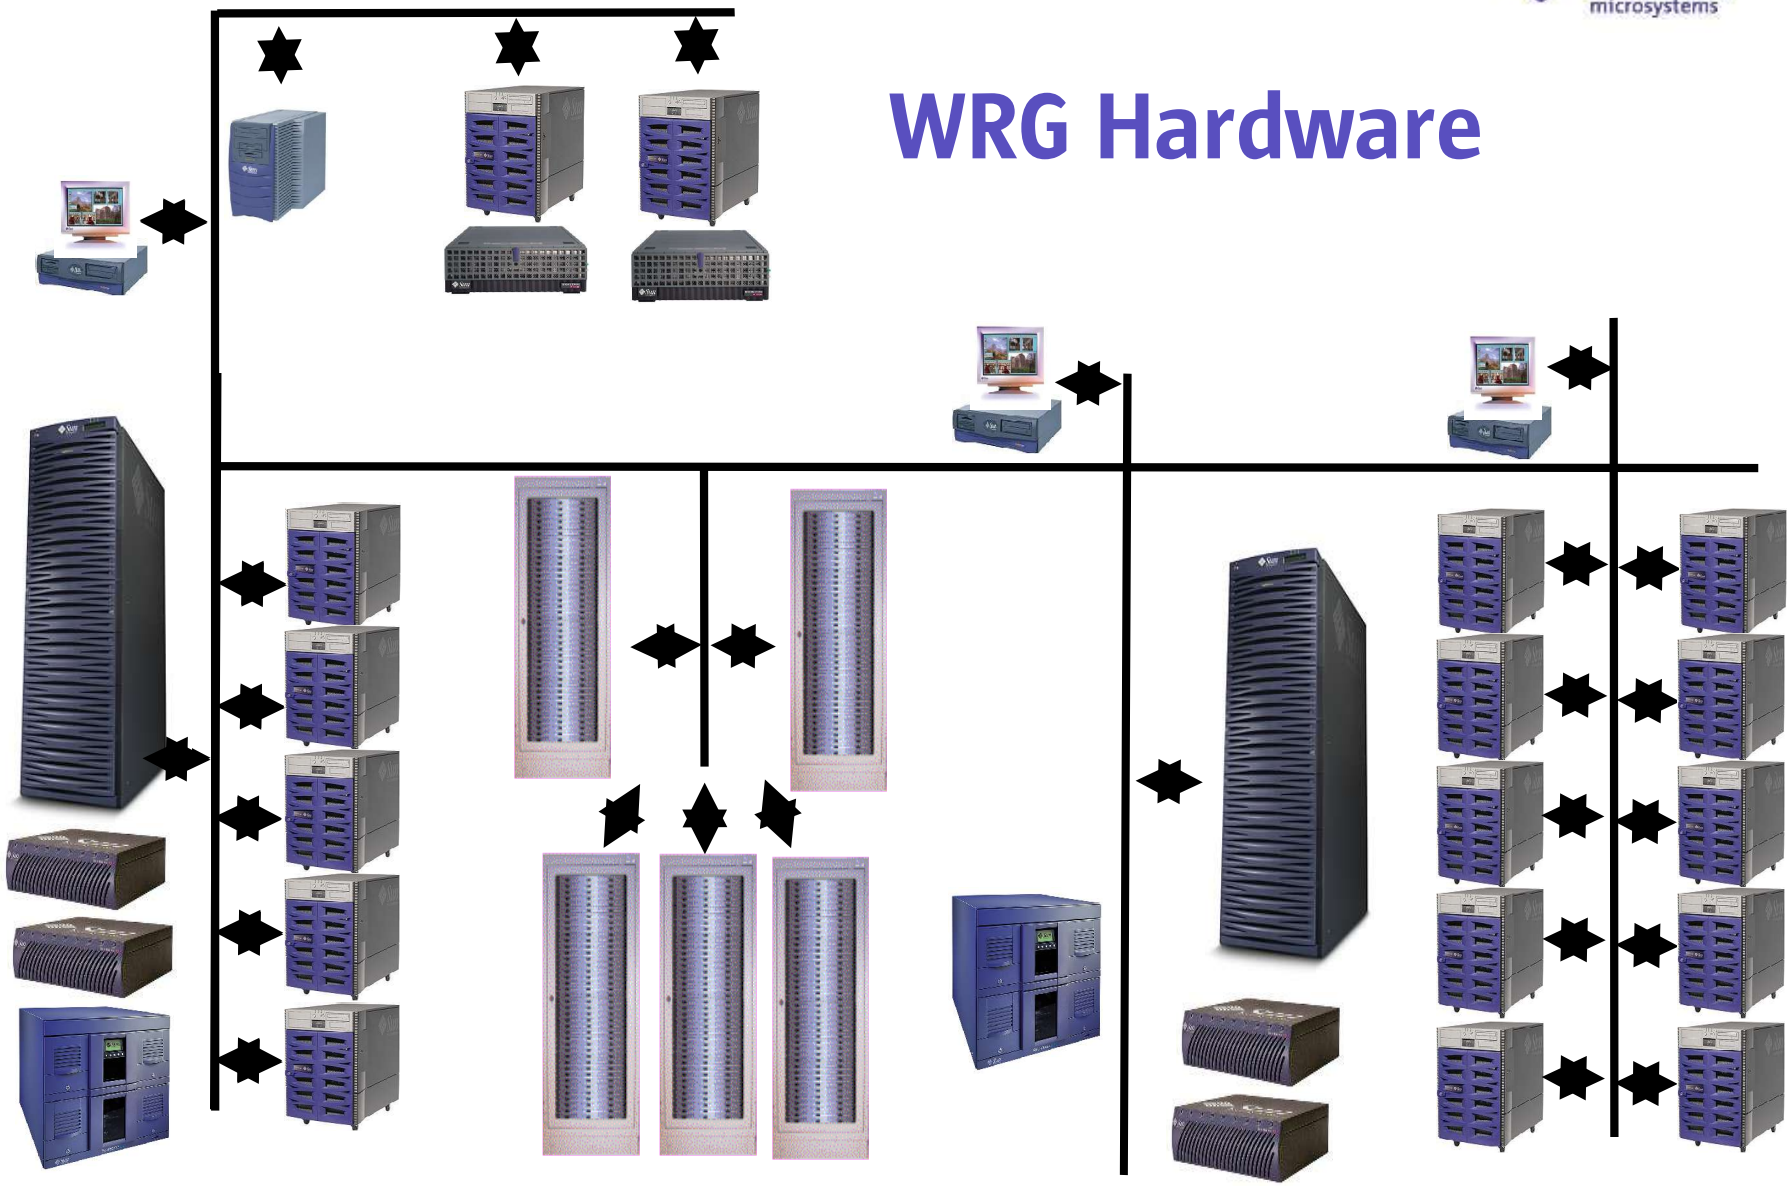

WRG Architecture Overview

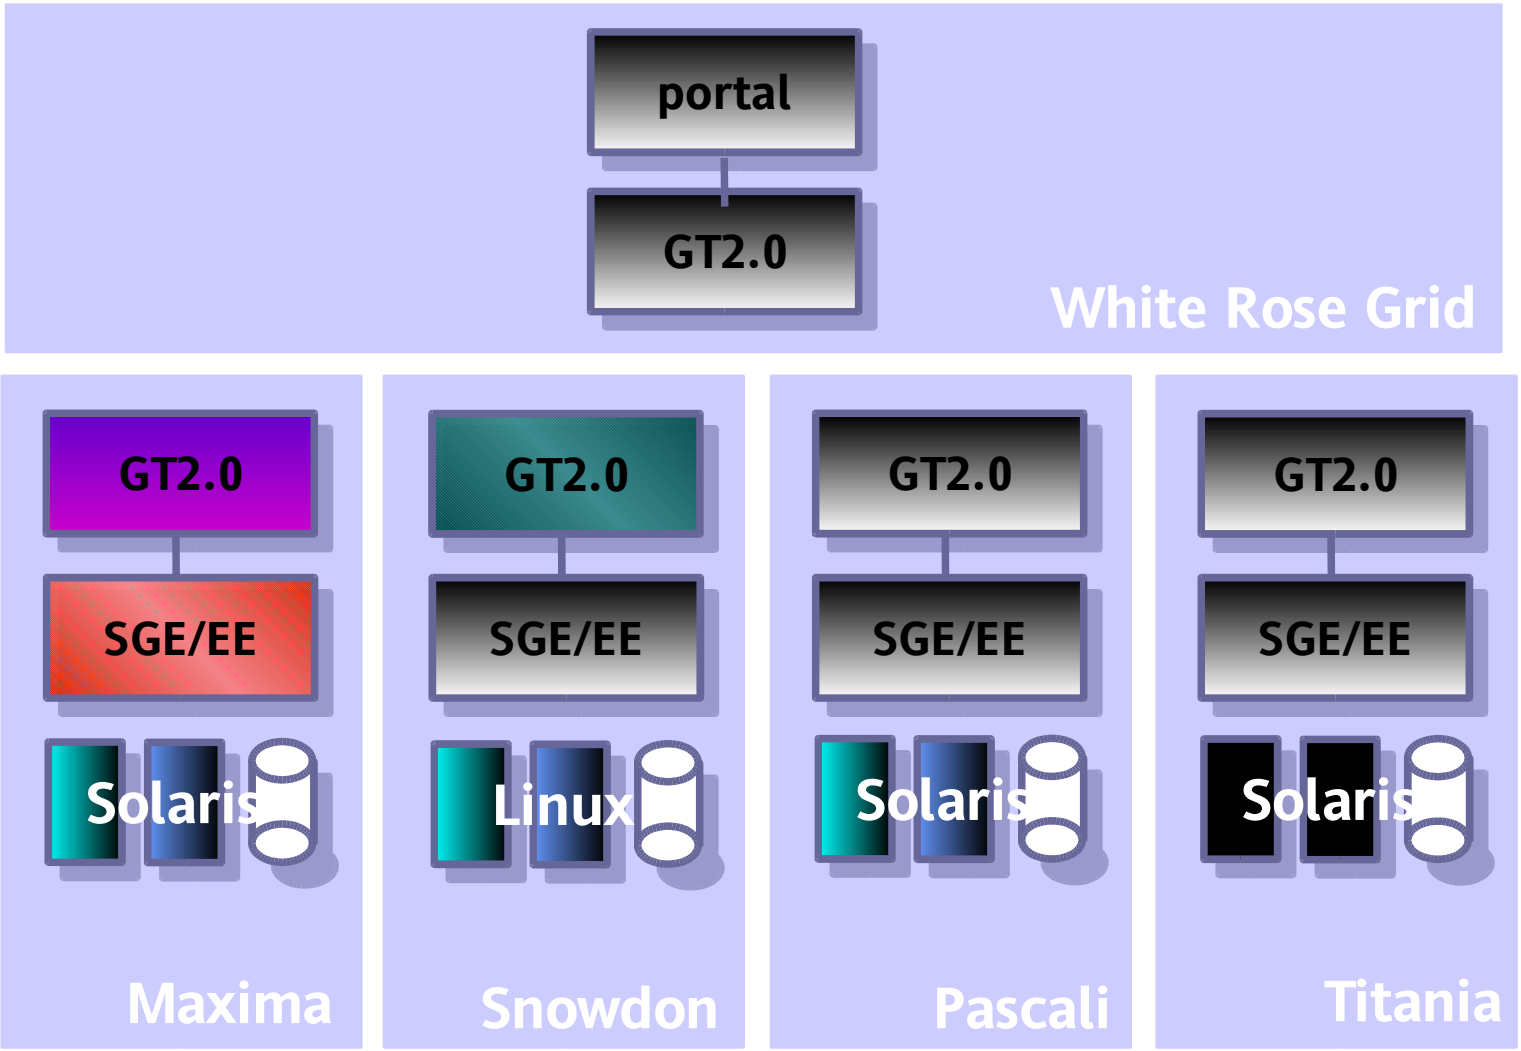

Data Grid: SAN, NAS, DAS

Storage Access, Management and Archival

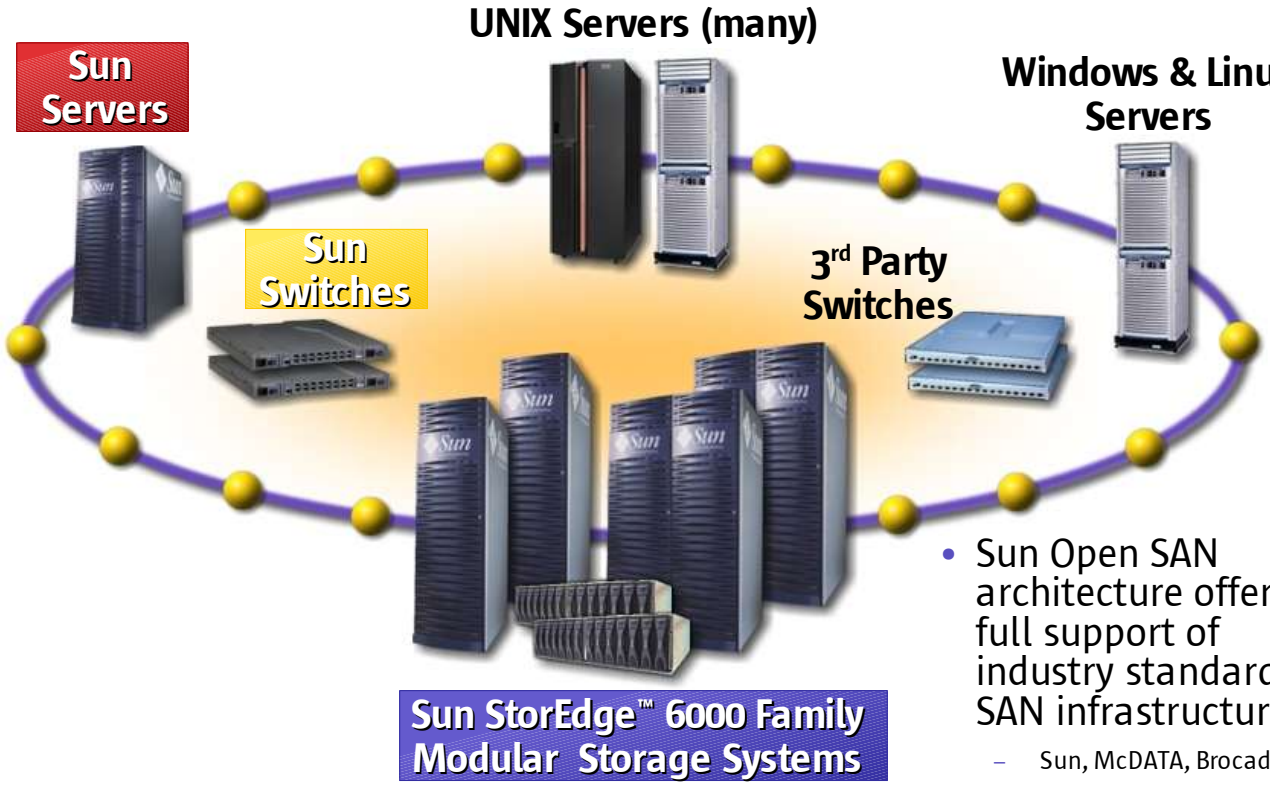

- Sun Open SAN architecture offers full support of industry standard SAN infrastructure
  - Sun, McDATA, Brocade

StorEdge 3000
Low Cost per GB → **StorEdge 6000**
Modular Flexibility ← **StorEdge 9000**
Low Complexity
High Performance

San Diego Supercomputer Center

- SDSC Problem
  - Data-intensive requirements: storage management, complex scientific applications, relational databases and data mining
  - Mixed/heterogeneous compute environment
  - Provide a gateway for grid distributed data to develop NSF funded TeraGrid
- Solution – Sun HPC SAN
  - 500TB's of new Sun Storage
  - Single point of data, filesystem and storage management
  - Sun Professional Services support
- New Capabilities
  - >3.2GB/sec with Sun StorEdge™ 3910 industry's fastest movement of data across TeraGrid network
  - Reduction from days to hours in the transfer of multi-terabyte datasets

Graphics Grid:

Access for More Users to Visualization Services at Required Visual Quality and Performance Levels

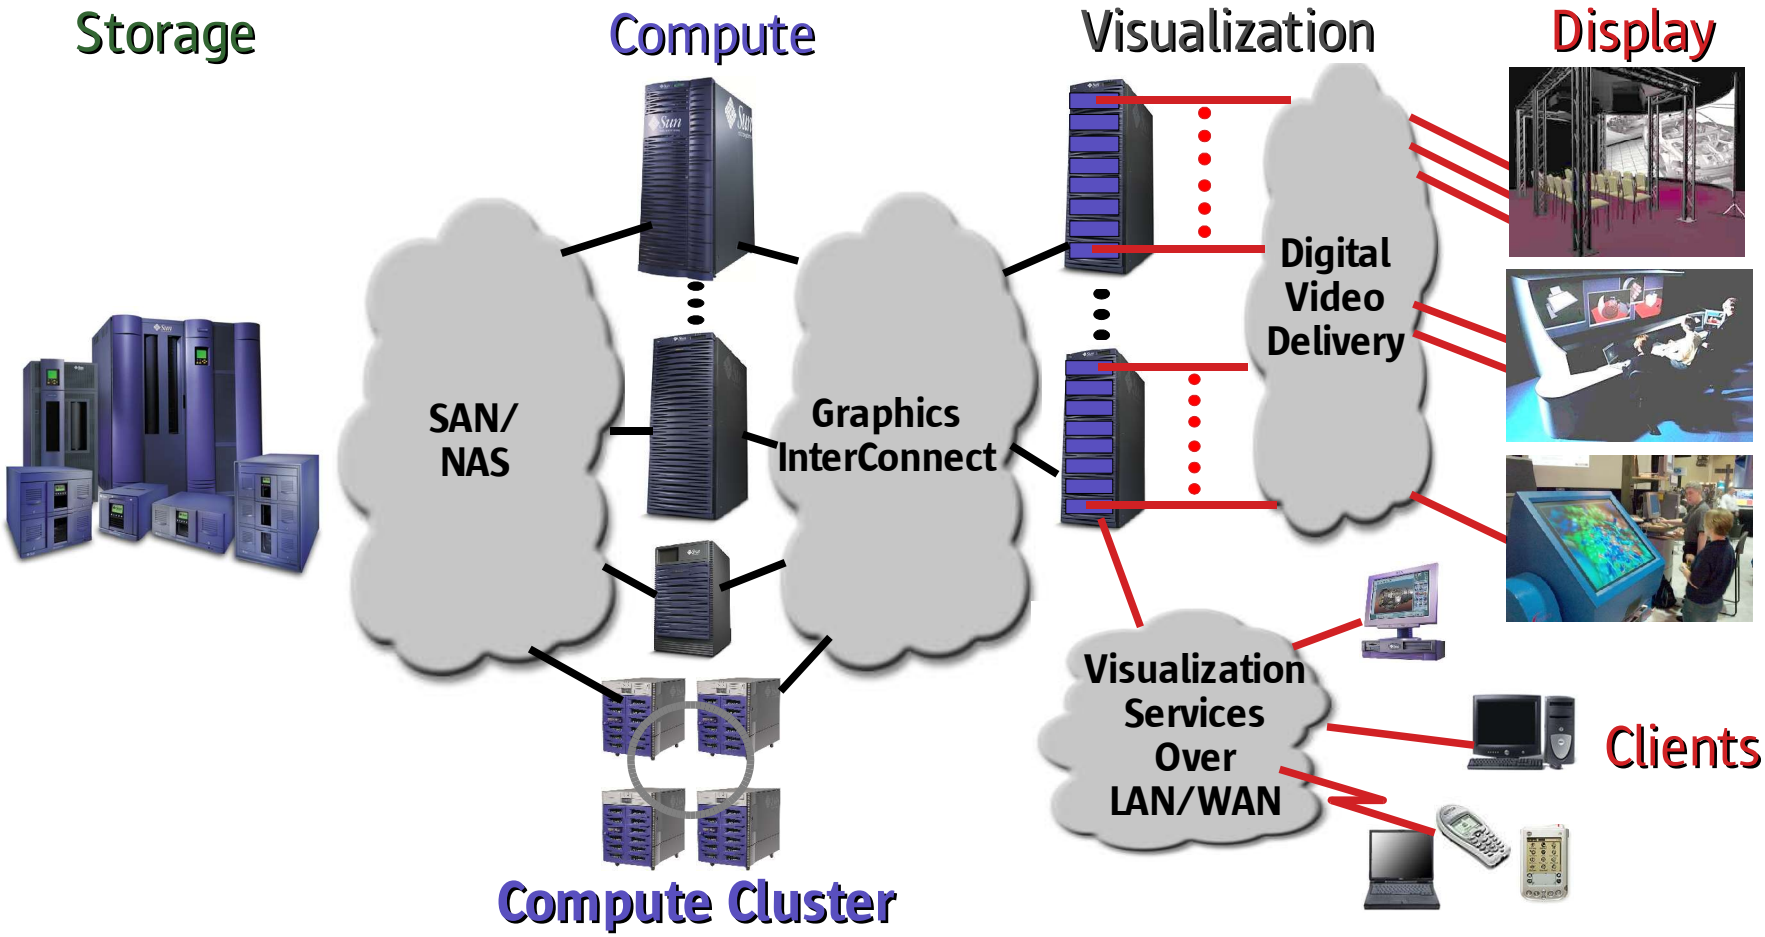

Scalable Distributed Sun Fire Visualization Solution

- Scalable computing power for visualization applications
- Scalable to larger numbers of displays with multiple Sun Fire V880z rendering engines
- No application modifications necessary

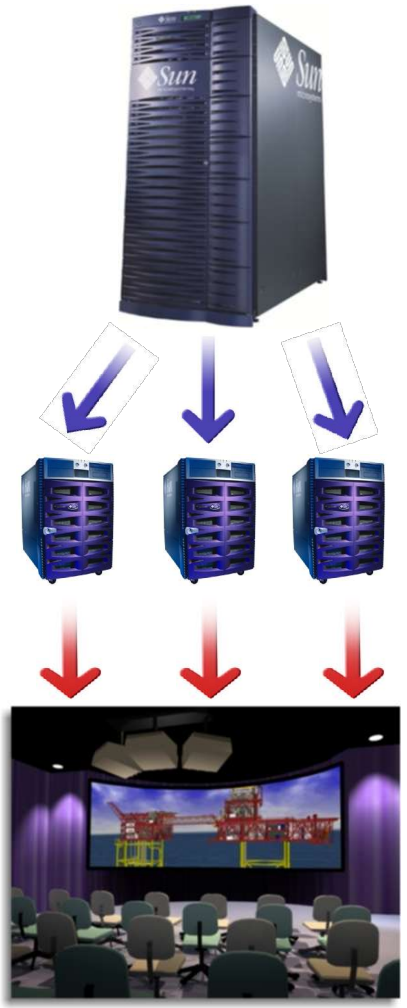

Large Screen, Tiled Displays, CAVE, etc.

The University of Texas at Austin

Visualization at the ACES Visualization Lab

- Visualization for a range of applications including computational fluid dynamics, bioinformatics and geophysical modeling
- Complex rendering and multiprocessing of graphics and images for researchers in a diverse range of fields
- Sun Fire V880z and a Scalable Distributed Sun Fire 6800 visualization system

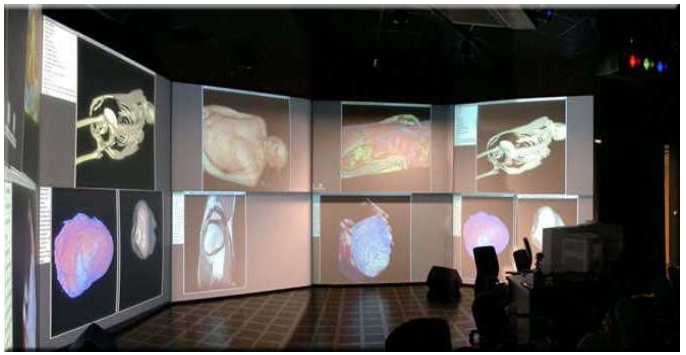

"The speed of graphics and superior image quality of the Sun Fire V880z visualization system will play a significant role in these key areas of technology and research."

Dr. Kelly Gaither, Associate Director of the Texas Advanced Computing Center

Future of the Grid

- To what extent will the commercial sector be involved?
  - It has to solve their day-to-day problems and help them drive up efficiency and drive out cost without adding complexity to their operations. Fortunately, Grid resource management software does this. Security and level of comfort is the big issue. Enterprise grids are happening but even inside the firewall there are a lot of political and organizational issues. Crossing the firewall is significantly more challenging for commercial customers – it's a cultural shift – but everyone is learning the new culture in this Web-centric age.
- What is division of effort between funding agencies and commercial sector?
  - The funding agencies should encourage industrial and commercial users to take advantage of the expertise in the academic community and government labs. The funding agencies and public research community should encourage industrial/commercial participation in joint R&D activities which use Grid infrastructure to advantage. This will help grow their experience and comfort levels.

Future of the Grid

- Will there be support for R&D and operations?
  - Will R&D and operations work together? We will instrument operations and inventory and the customer experience will be better quantified to drive requirements back into the R&D process.
  - As business processes acquire massive amounts of data and become more complex, the use of HPTC techniques to optimize business operations will be commonplace. HPTC is not just for R&D anymore.
- Will there be adequate networking infrastructure?
  - It seems inevitable. Speed-of-photon latency limitations will have to be dealt with via, for example, smart caching and replication.

The Revolution

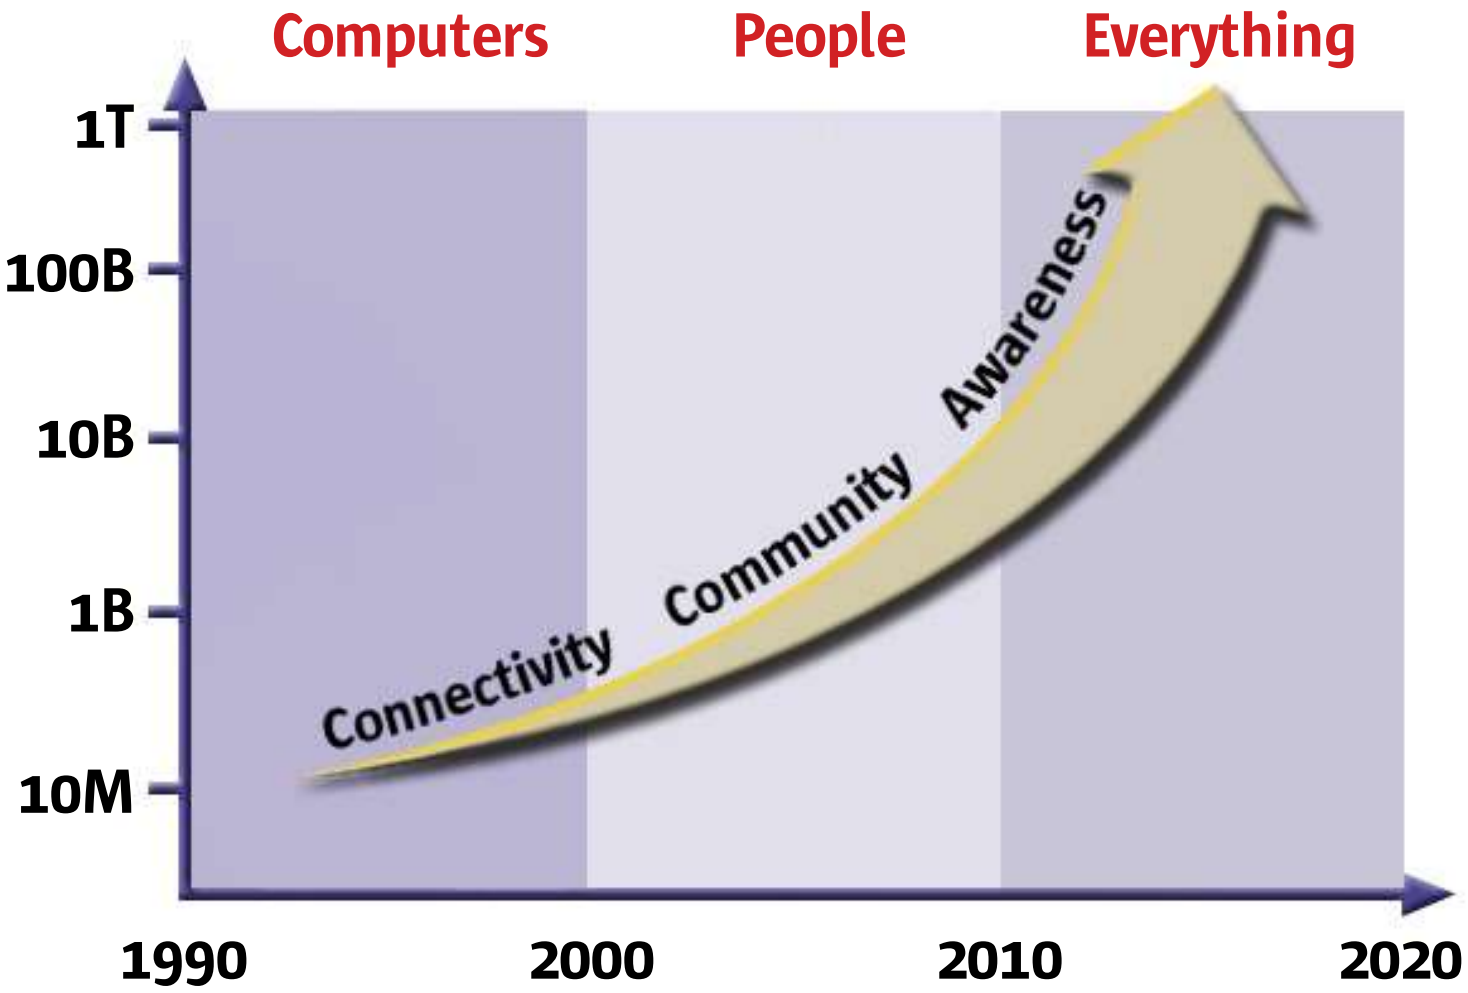

“All Grid – All the Time”

Stephen Perrenod

stephen.perrenod@sun.com

www.sun.com/hptc
www.sun.com/grid

Sun microsystems
We make the net work.

Frontiers of Automation

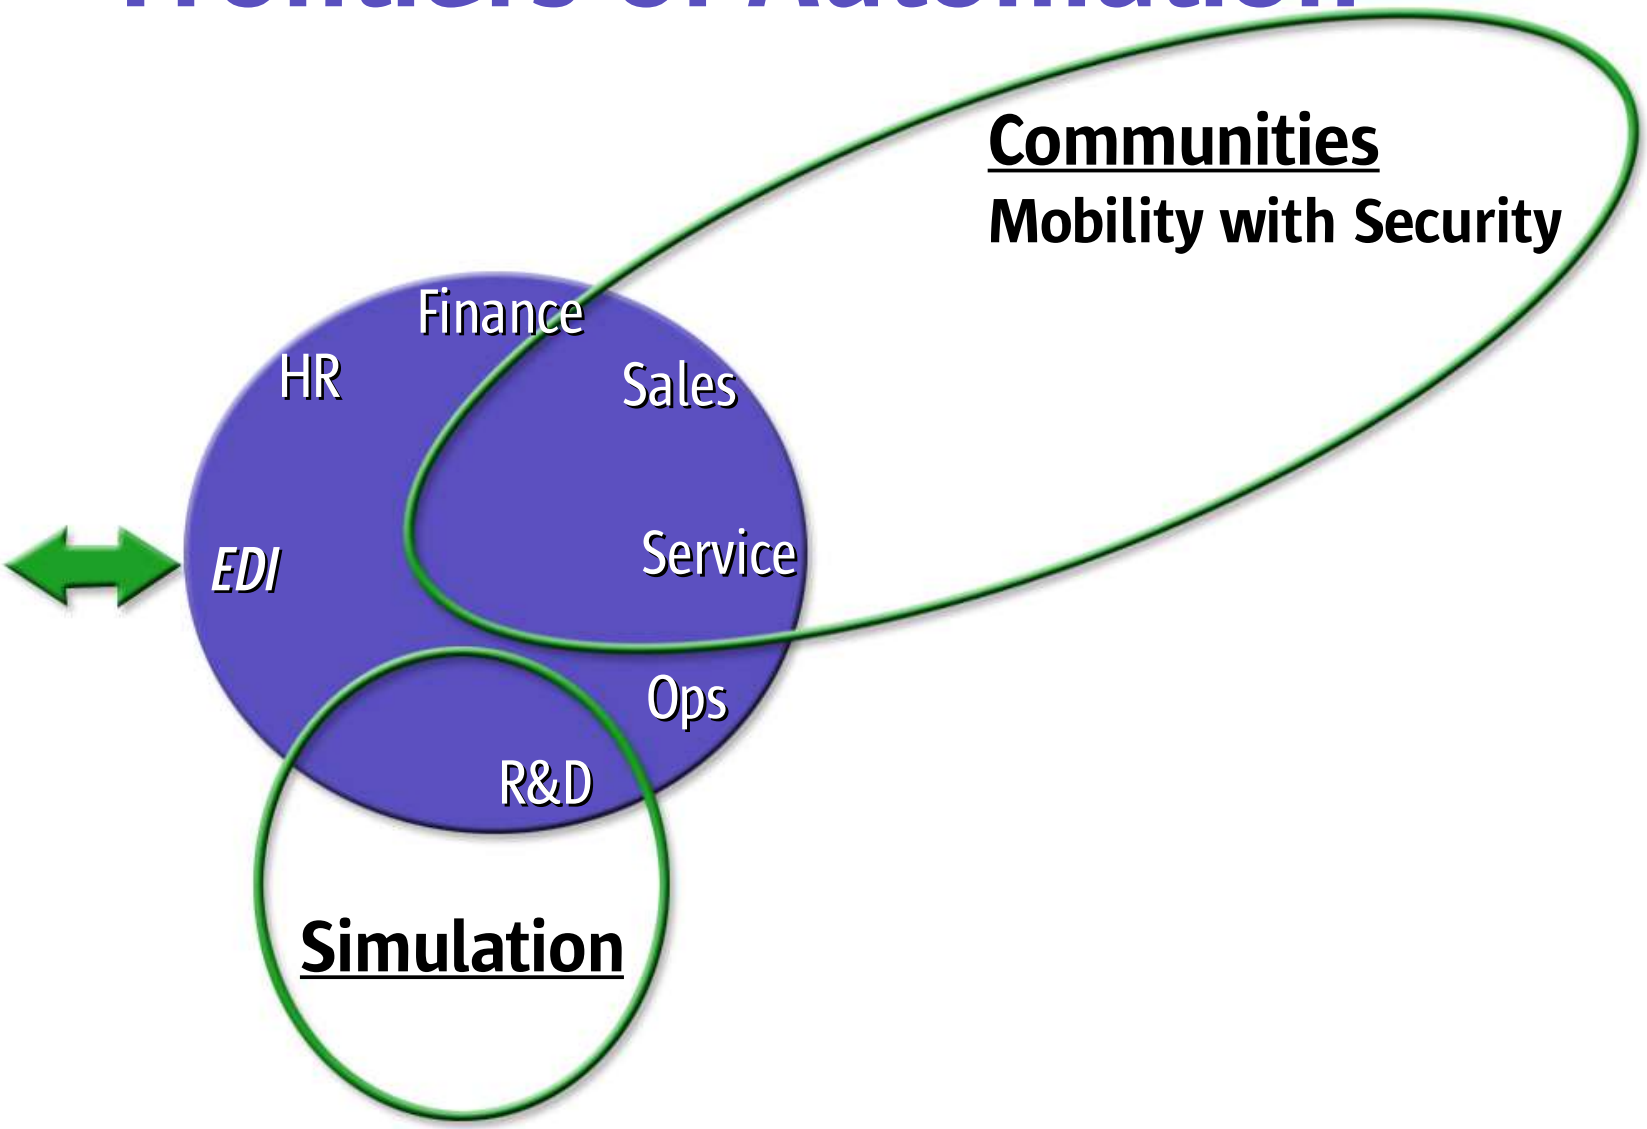

Frontiers of Automation

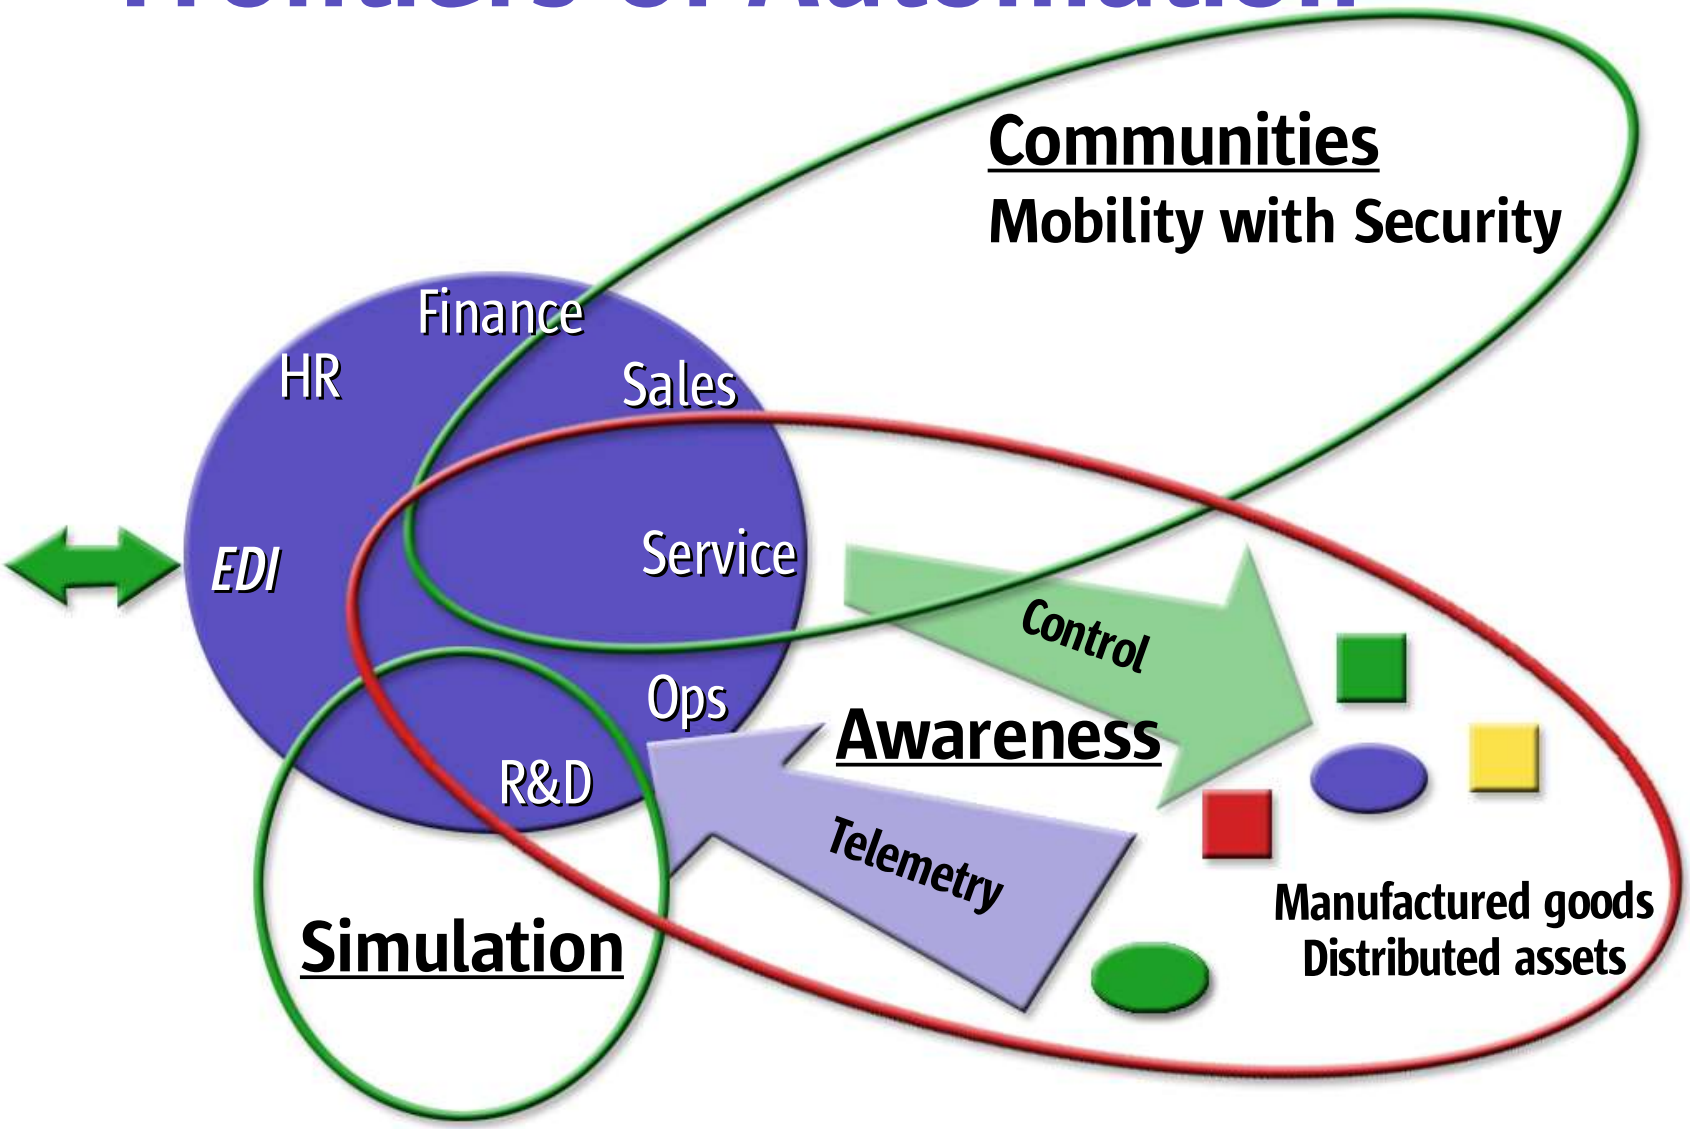

Virtualization of Resources

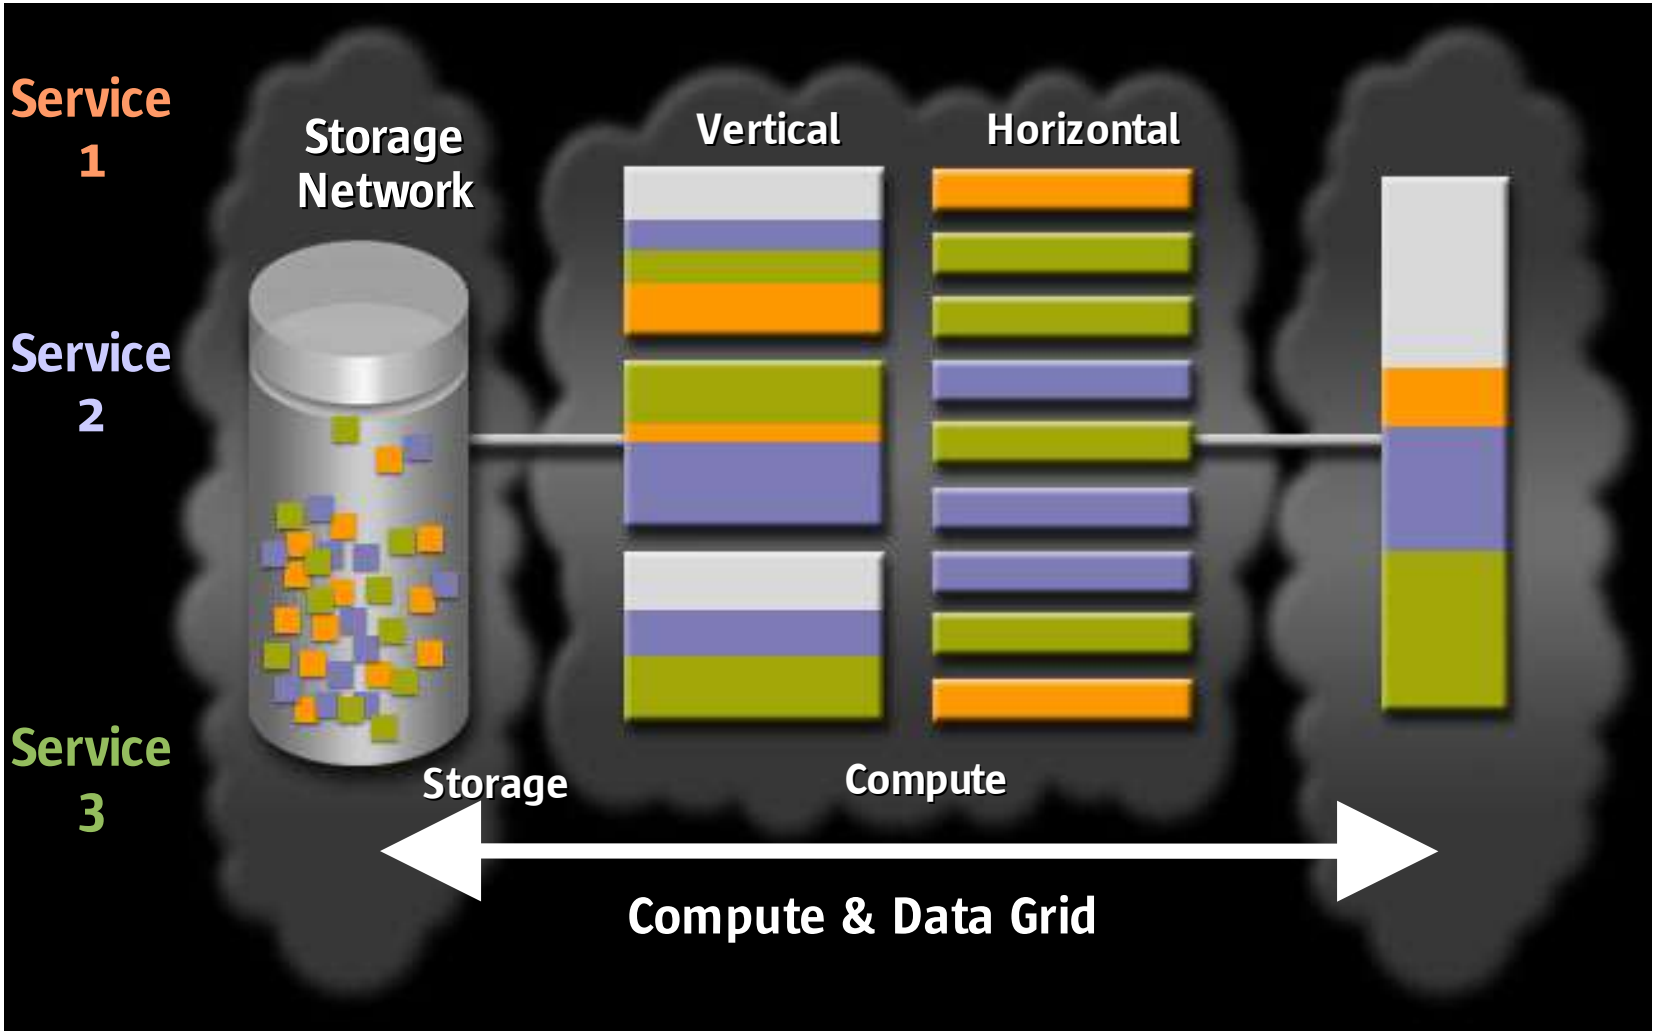

Services sharing resources

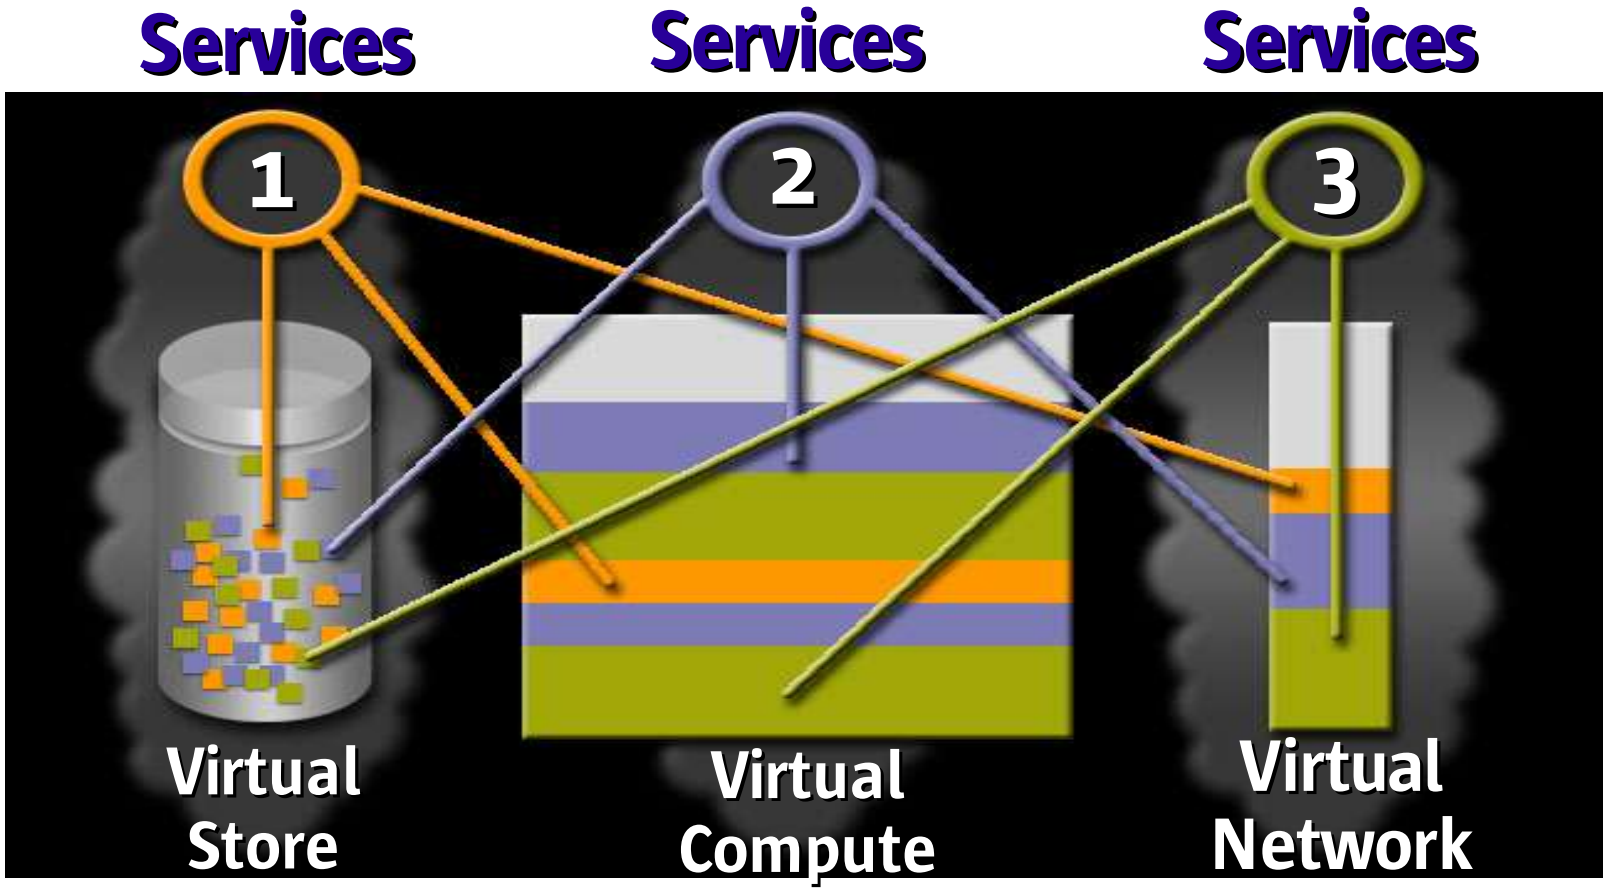

N1: managing services, not servers

Two Views of the Stack

DEVELOPERS

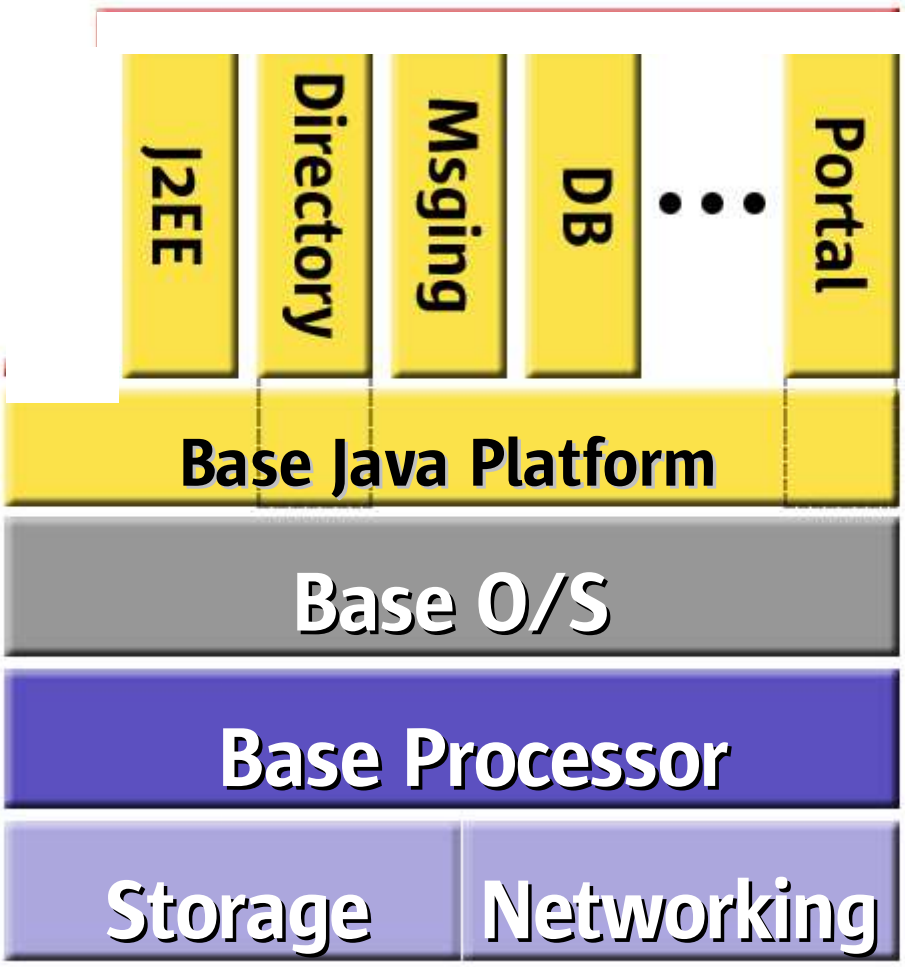

OPERATORS

Software Components Are Networked Graphs

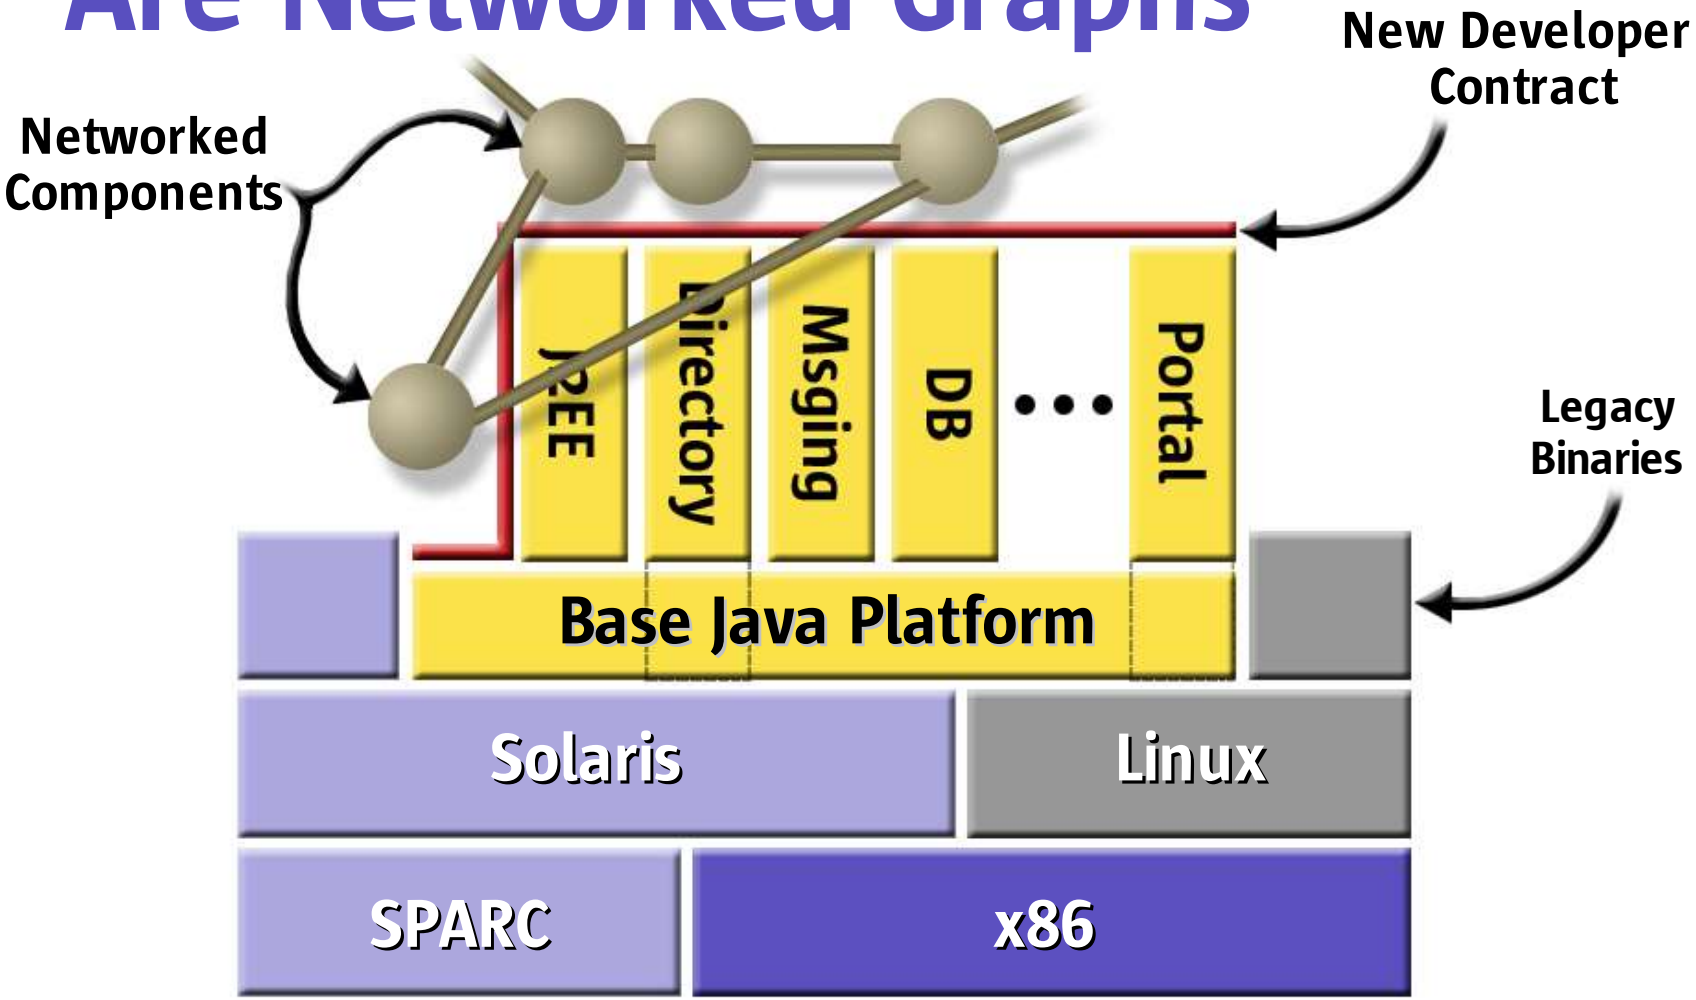

Orion: The Release Vehicle

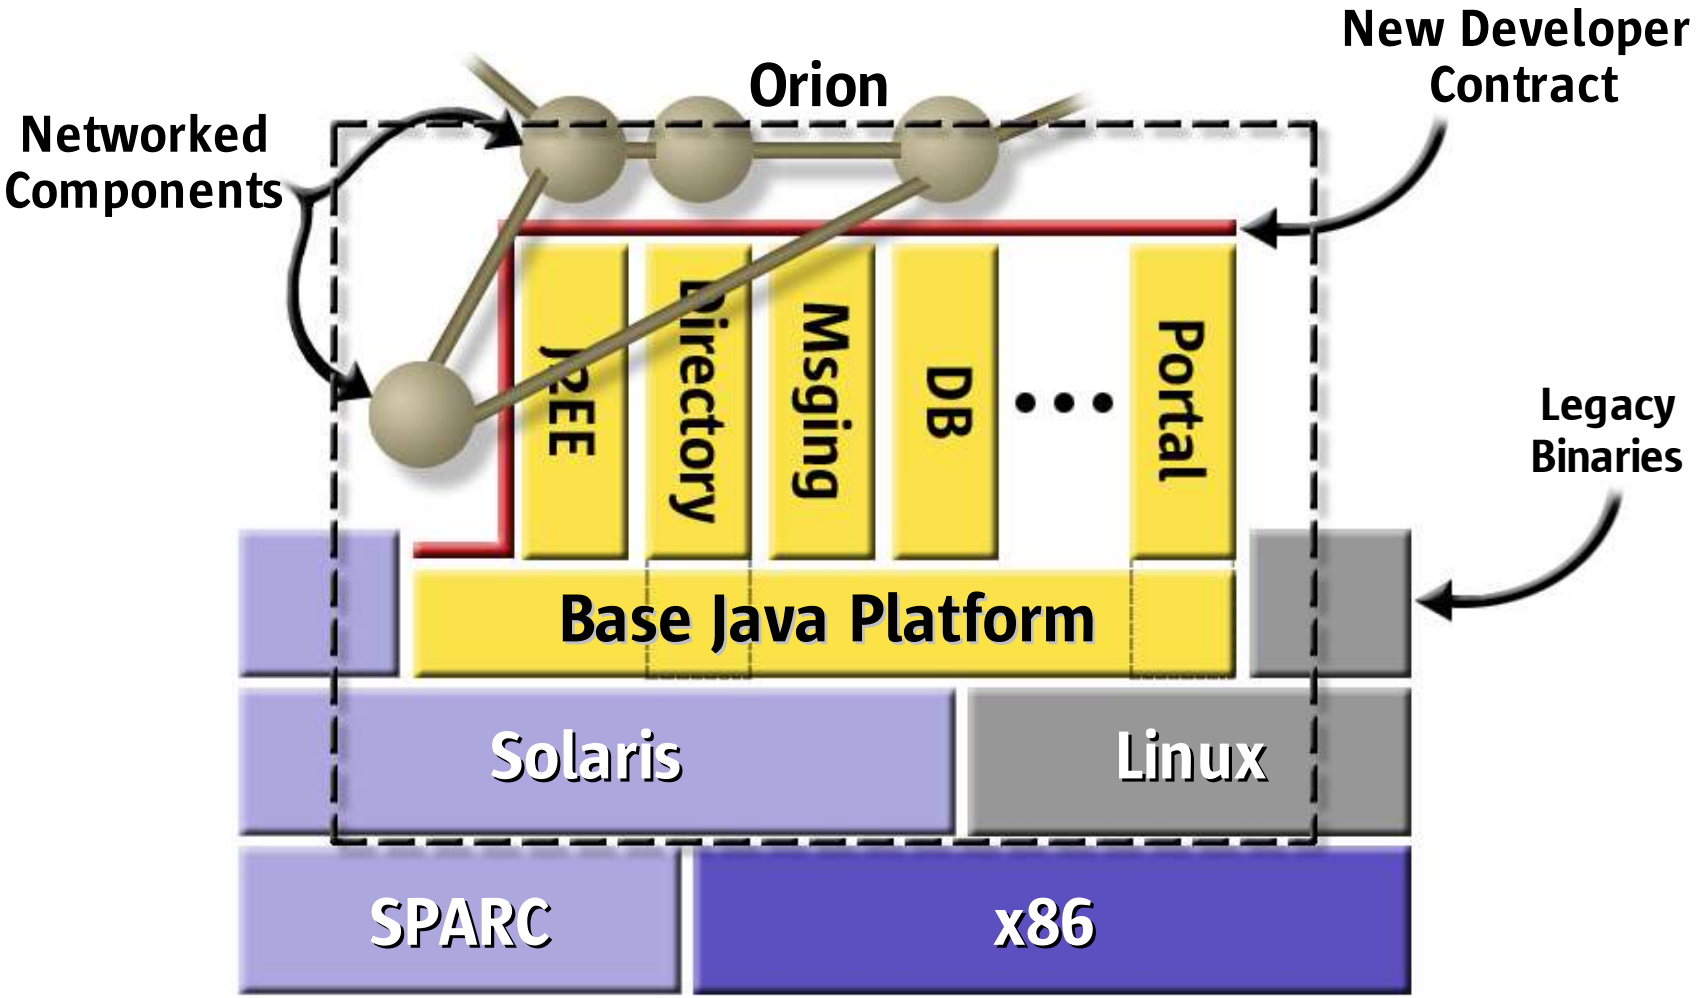

Two Views of the Stack

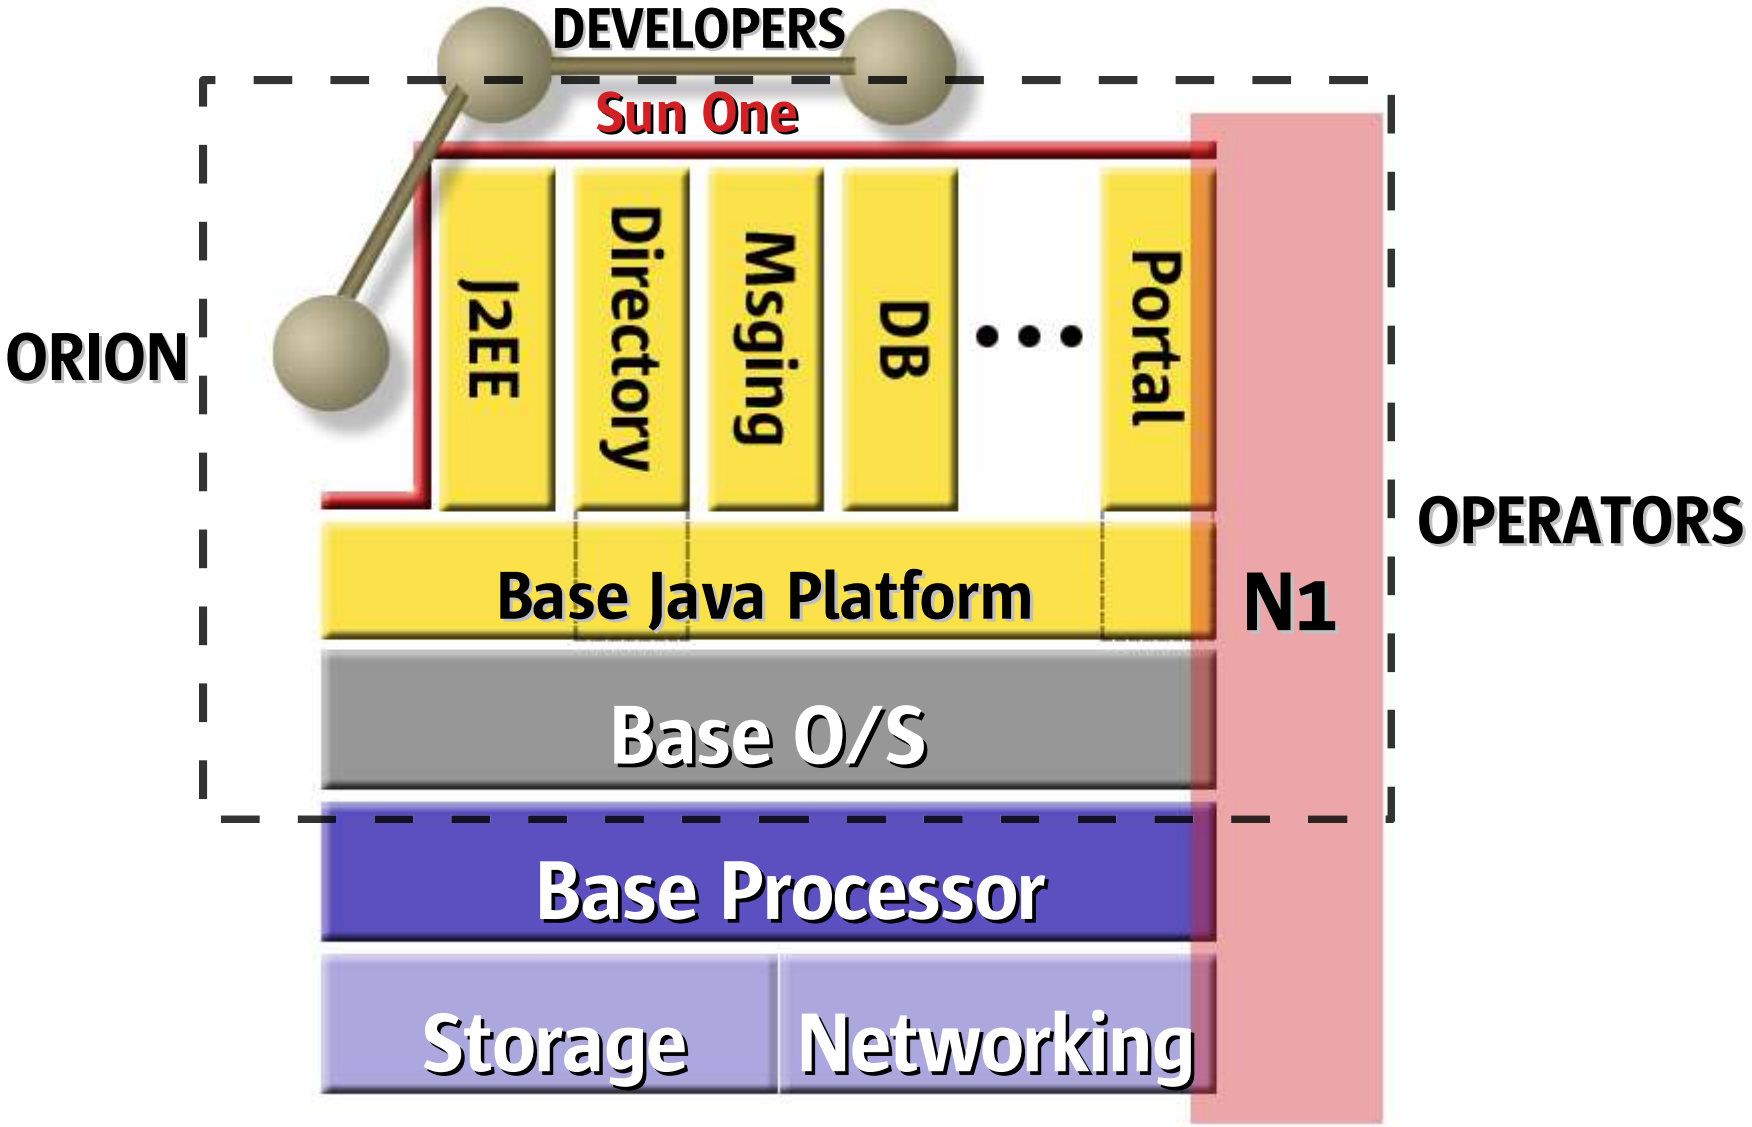

Server and Storage Elements are NOT "Commoditization" of Existing Systems

- The commoditization is essentially the end-game of stagnant architecture
- Current architectures are solving the old SW problems
  - Processor-centric, not network-centric
  - Not nearly efficient enough
- Must and will enable much higher rates of innovation

p.s. Keep the Conventional Wisdom coming!

Confluence of Two Megatrends

- Software application components have incredible amounts of “thread-level” parallelism
  - Each user of a website causes many threads to be created
  - The numbers of users served per second is the key metric (otherwise delays per user)
  - [This is the sweetspot for SMPs]
- Moore’s Law has progressed to the point where we can integrate a substantial multithreaded Server-on-a-Chip

Servers will become…

- Single boards (or single board + memory expansion board)
- Low latency, IP networking with RDMA is the only I/O
- 1 to 4 processor sites within a "sea of memory"
- Kernel O/S sees as a 32-way to 1000-way SMP
- From the point of view of the N1, stateless components

The Microsystems Era

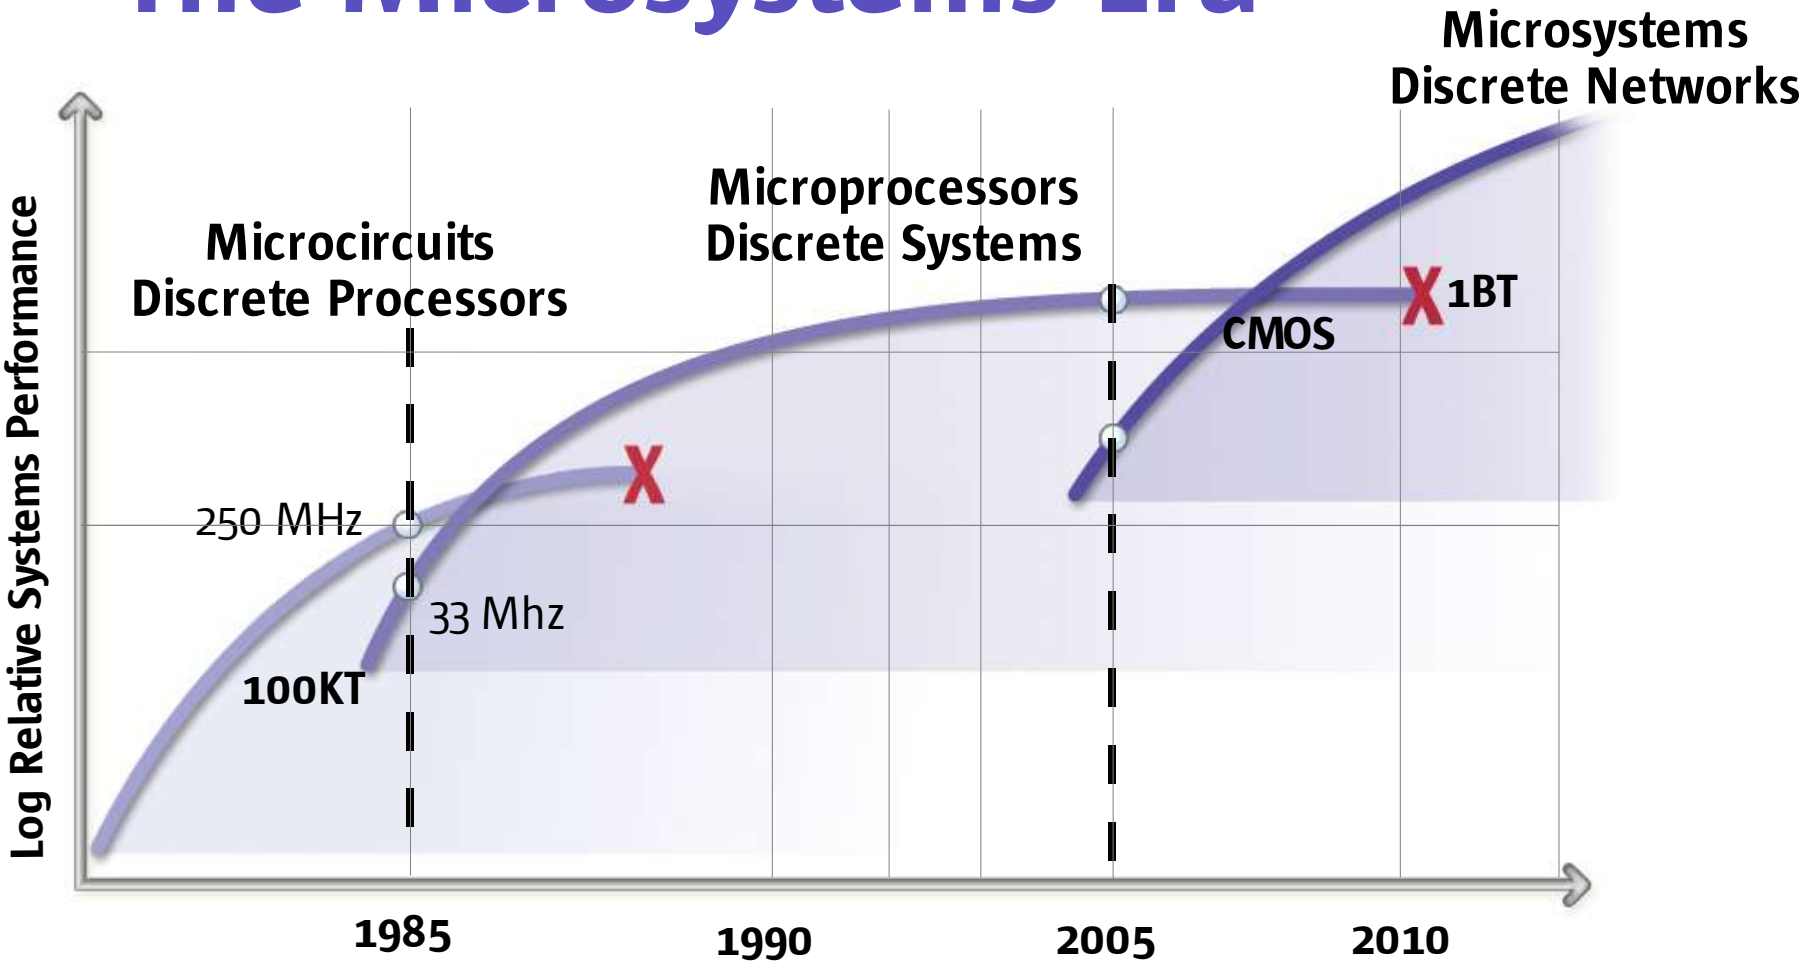

Crossing the Chasm…

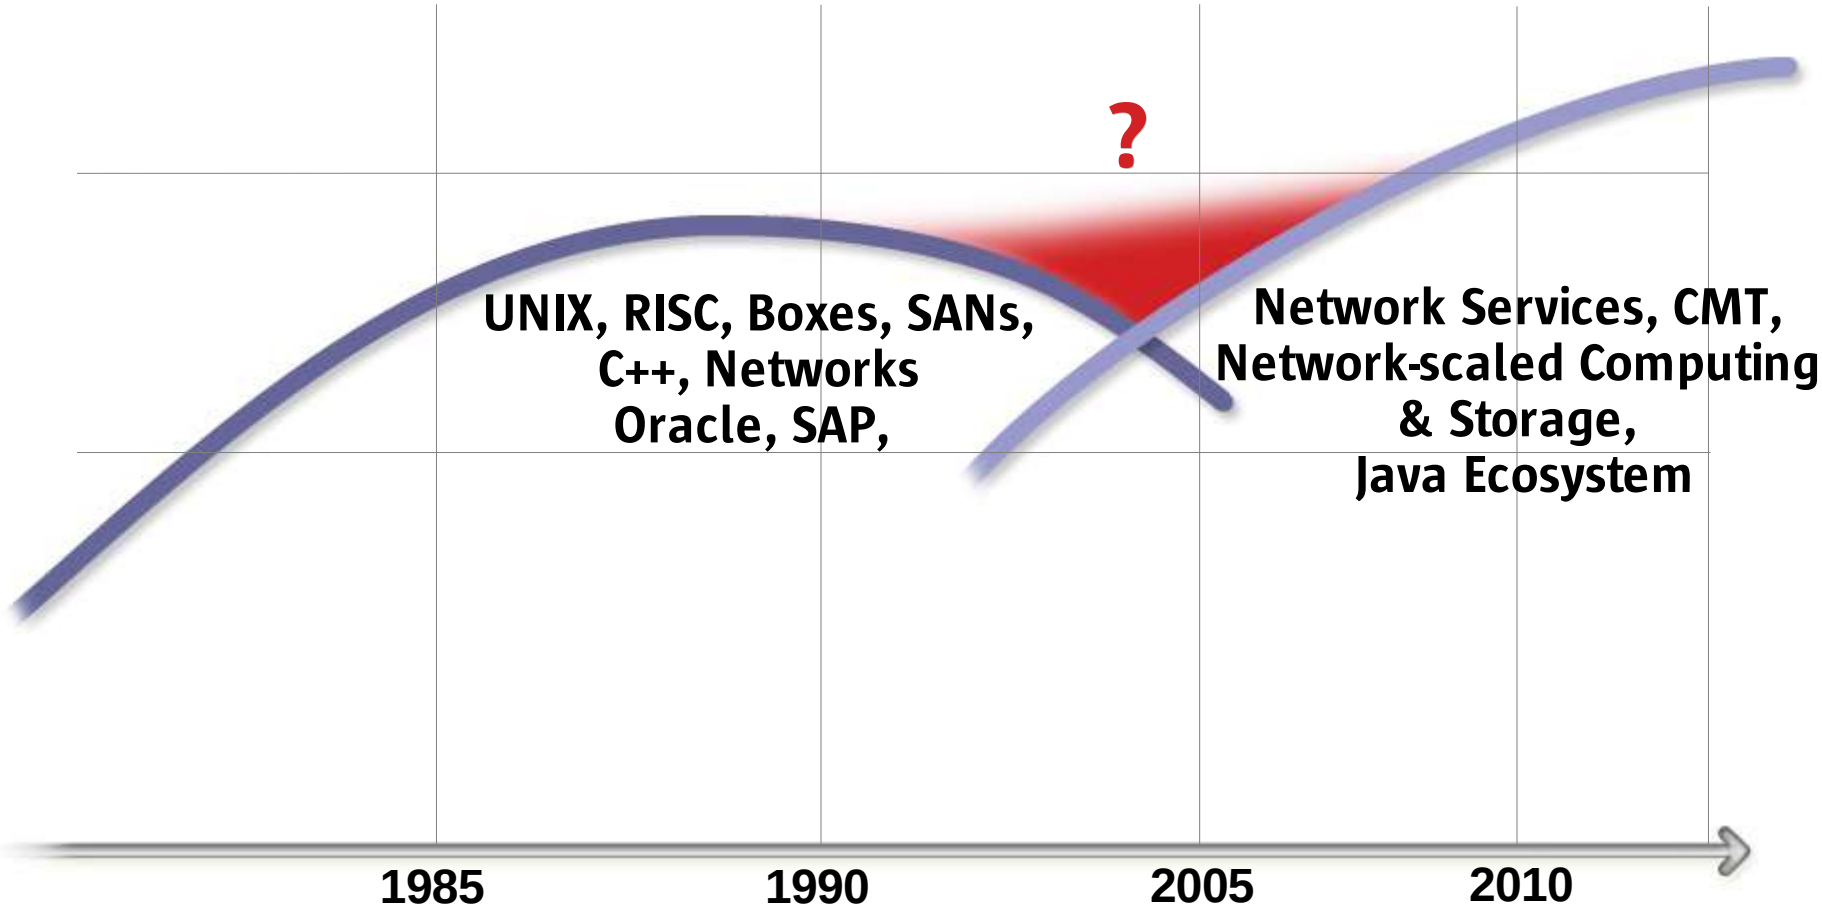

Three Technology Pillars

1. Developers: **SunONE/Orion**
2. Operators: **N1**
3. Systems: **CMT/Microsystems**

Conclusion

- Software is changing, from single user applications to mega-user services
- Systems are changing to support these new services
  - The network is (the backplane for) the computer
  - Macro-virtualization is enabling technology
  - Multithreading reinvents processors
- Both of these align to the strike zone of Sun's technologic competencies

Bottom Line

- Still in the first phases of networking
  - Expect the number of networked things to grow by another factor of 1M
  - New value from automating still-manual processes
- Expect radical transformation of SW and systems
  - Systems, in particular, are at a watershed
  - Mask complexity through people or eliminate through engineering

Sun's HPTC Strategy: "All Grid, All the Time"

Data Grids

Compute Grids

Graphics Grids

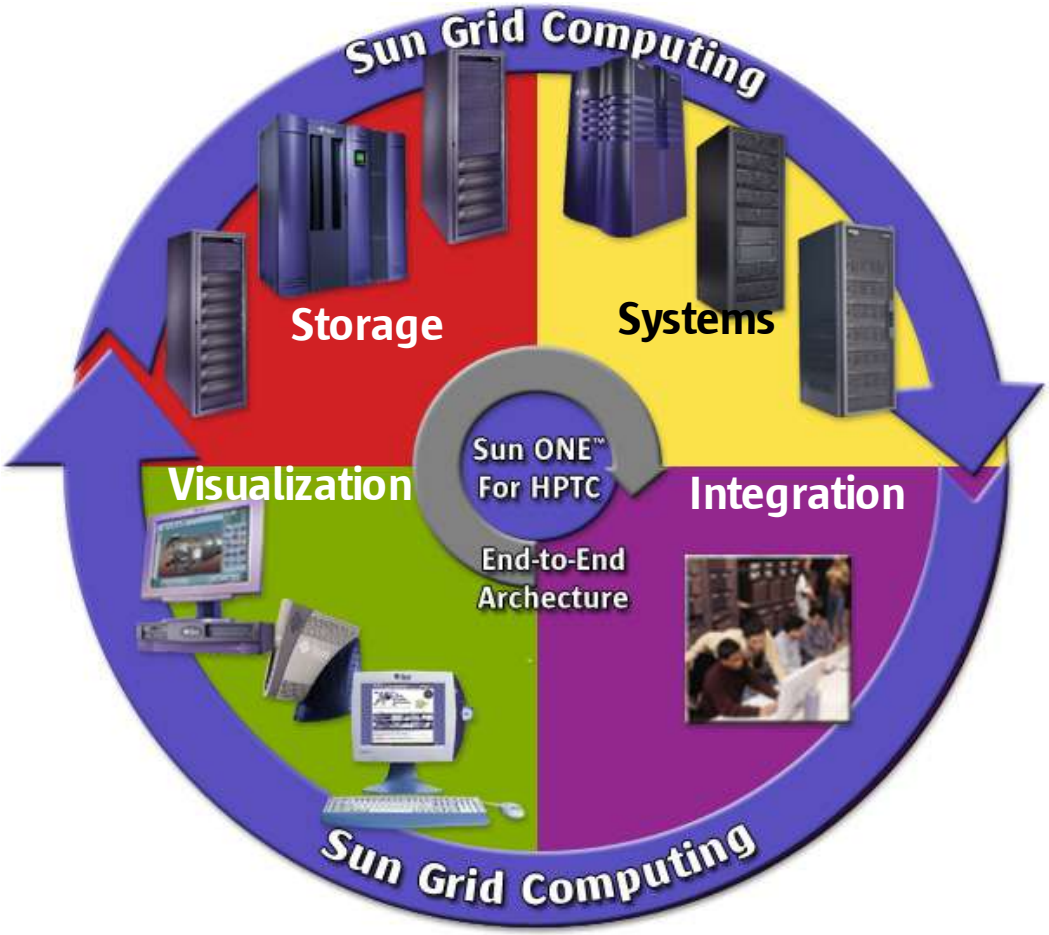

What is HPTC?

Commercial and HPTC Characteristics

- Commercial
  - Transaction oriented
  - High volume of computationally light and rapid operations
  - High volume of small data packets
- Workload types
  - Web page access
  - E-mail
  - Database access / update
  - Banking transactions
- Service expectation
  - Completion in seconds
  - Guaranteed response time regardless of number of requests
  - Always on availability

- High Performance & Technical
  - Job or project oriented
  - Small volume of computationally intense or lengthy operations
  - Small volume of large complex models, algorithms, & data sets
- Workload types
  - Crash Testing
  - Ocean and Climate Modeling
  - Genomics and Proteomics
  - Portfolio and Risk Analysis
- Service expectation
  - Completion in hours or days
  - Guaranteed quality of results
  - Results always available within required timeframe

HPTC Strategy

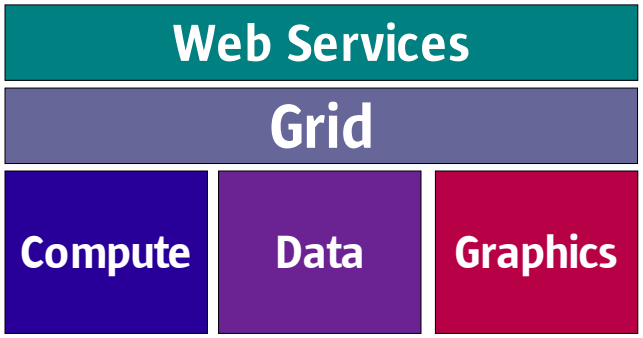

- Industry Trend: Grid
  - "All Grid, All the Time"
  - Grid product family
- Seeds: High-End
  - HPCS, iHEC
  - Sun Labs
- GTM: Solutions
  - Intersection of solutions and Grid
- Technology: HPTC Web Services
  - GGF, OGSA, ISVs, Web Svcs

HPTC Focus

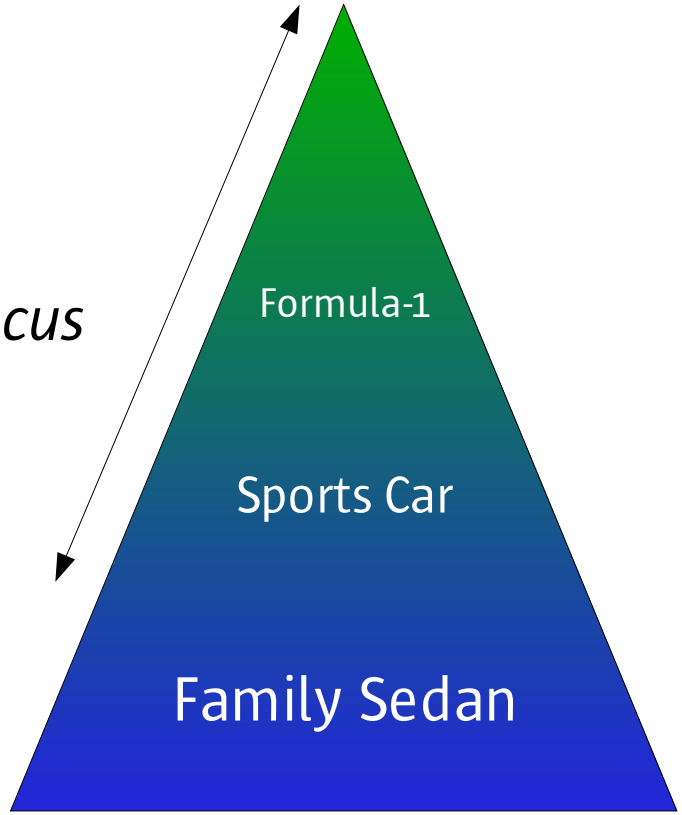

The Network is the Computer

Evolution in HPTC

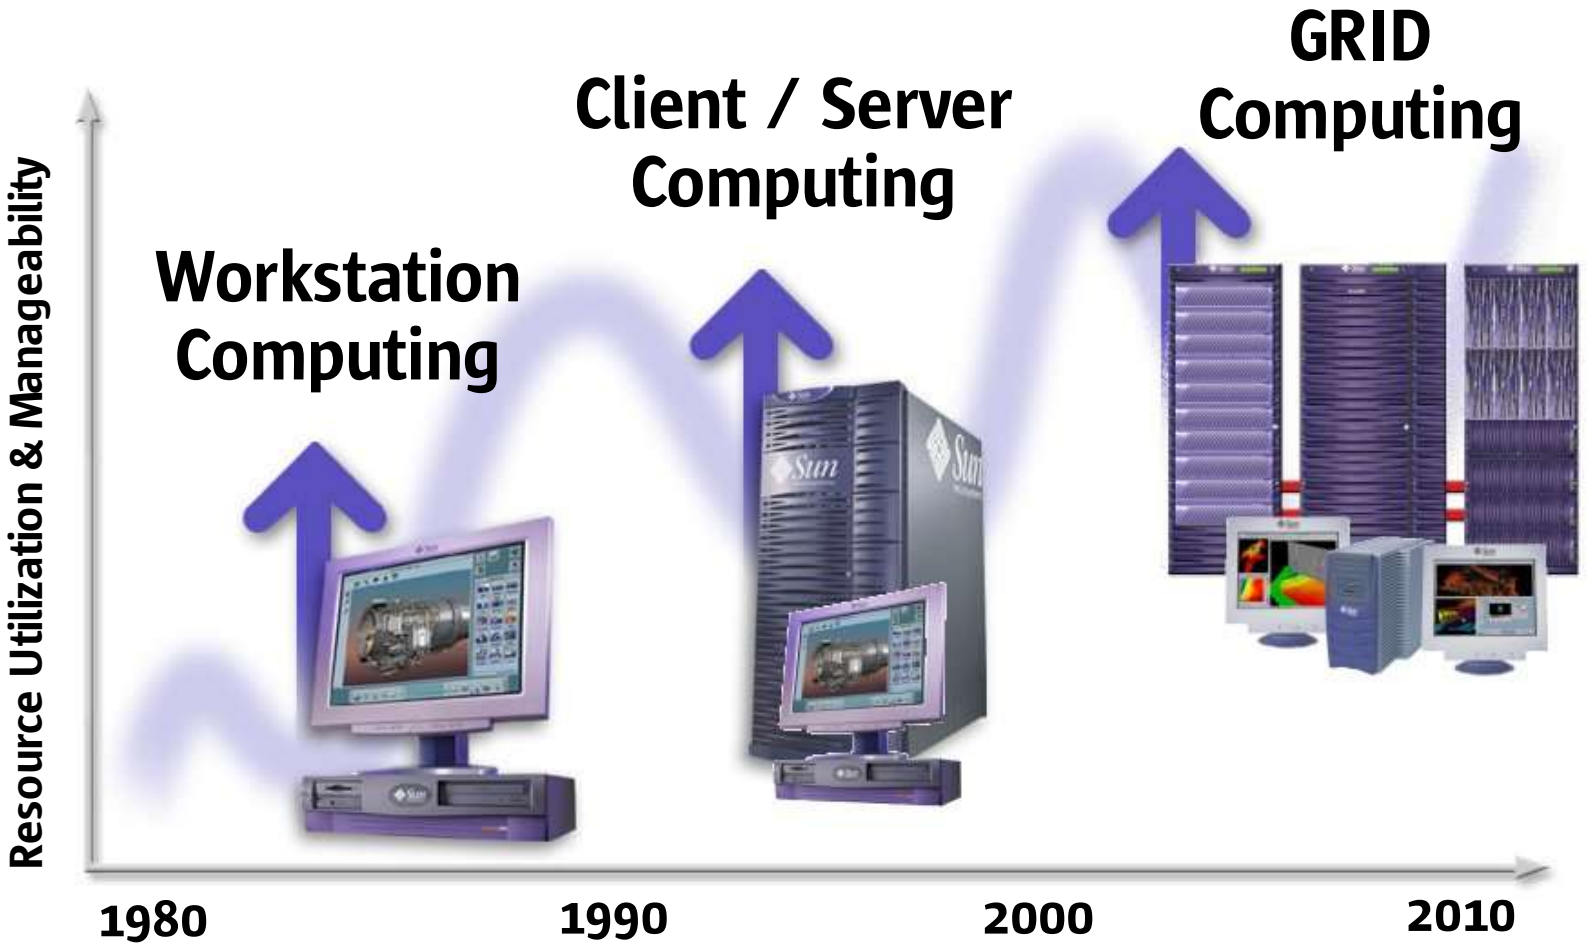

Grid Computing Components

Large SMPs

Storage

Small SMPs
Thin Servers
and Blades

Visualization Systems
Workstations
and Thin Clients

More than just clusters

More than just connected storage, compute and display systems

Data Grid: HPC SAN

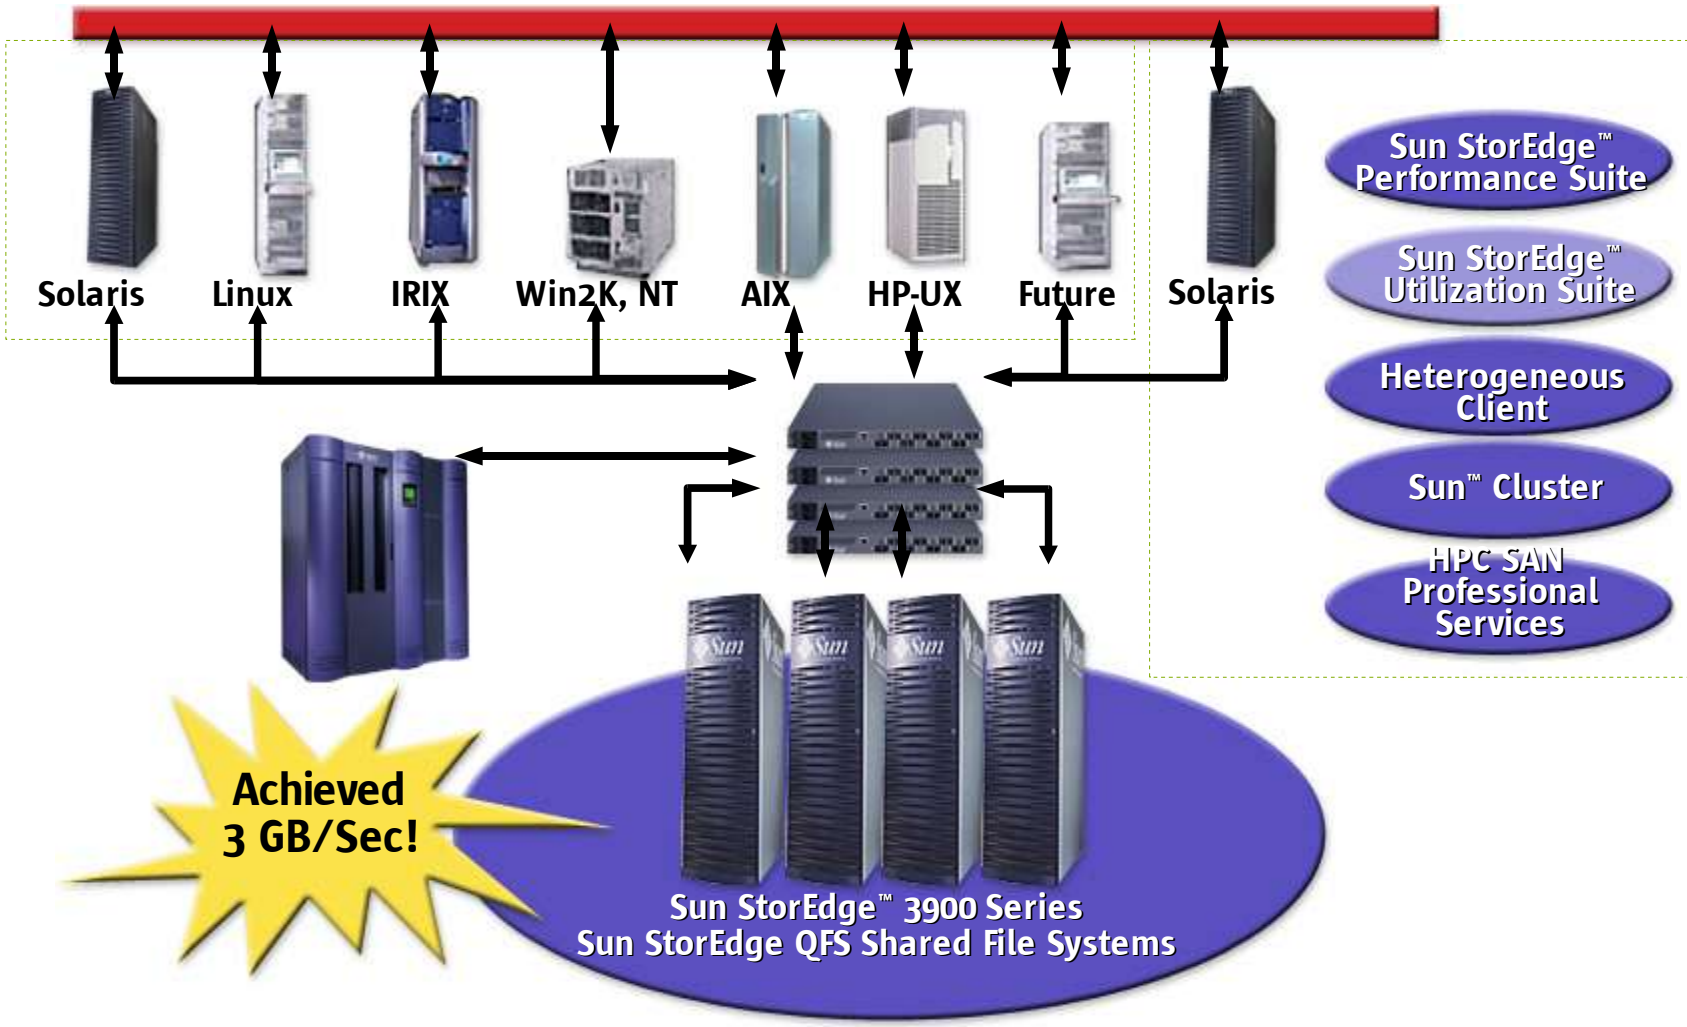

Innovation Pays:

DARPA HPCS Phase II Selection July 8th 2003

- Sun received $49.7 million to continue work on a revolutionary sustained peta-scale system concept
- DARPA HPCS Mission:
- Provide economically viable systems for the national security and industrial communities in the latter part of this decade that:
  - Improve the computational efficiency and performance of critical national security applications
  - Reduce cost and time of developing HPCS application solutions
  - Insulate research and operational HPCS application software from system specifics
  - Deliver improved reliability to HPCS users and reduce risk of malicious activities

Utility Computing

- DMdata (Denmark)
  - Leader in providing Utility Computing services
  - Available, reliable, scalable infrastructure
  - Server-based outsourcing offering has opened up opportunities in finance, transportation, the public sector, healthcare and agriculture; “25 to 30% lower cost”
  - Pay-as-you-go pricing
  - Fast, flexible allocation of resources
  - Supported by Sun's expertise