Managing Large Data Centers

July 9, 2003

Arland Kunz

Agenda

- **Enterprise landscape**
  - What do we have today
- **How can you manage this today**
- **Management trends**
  - What to expect in the future

Enterprise Ecosystem

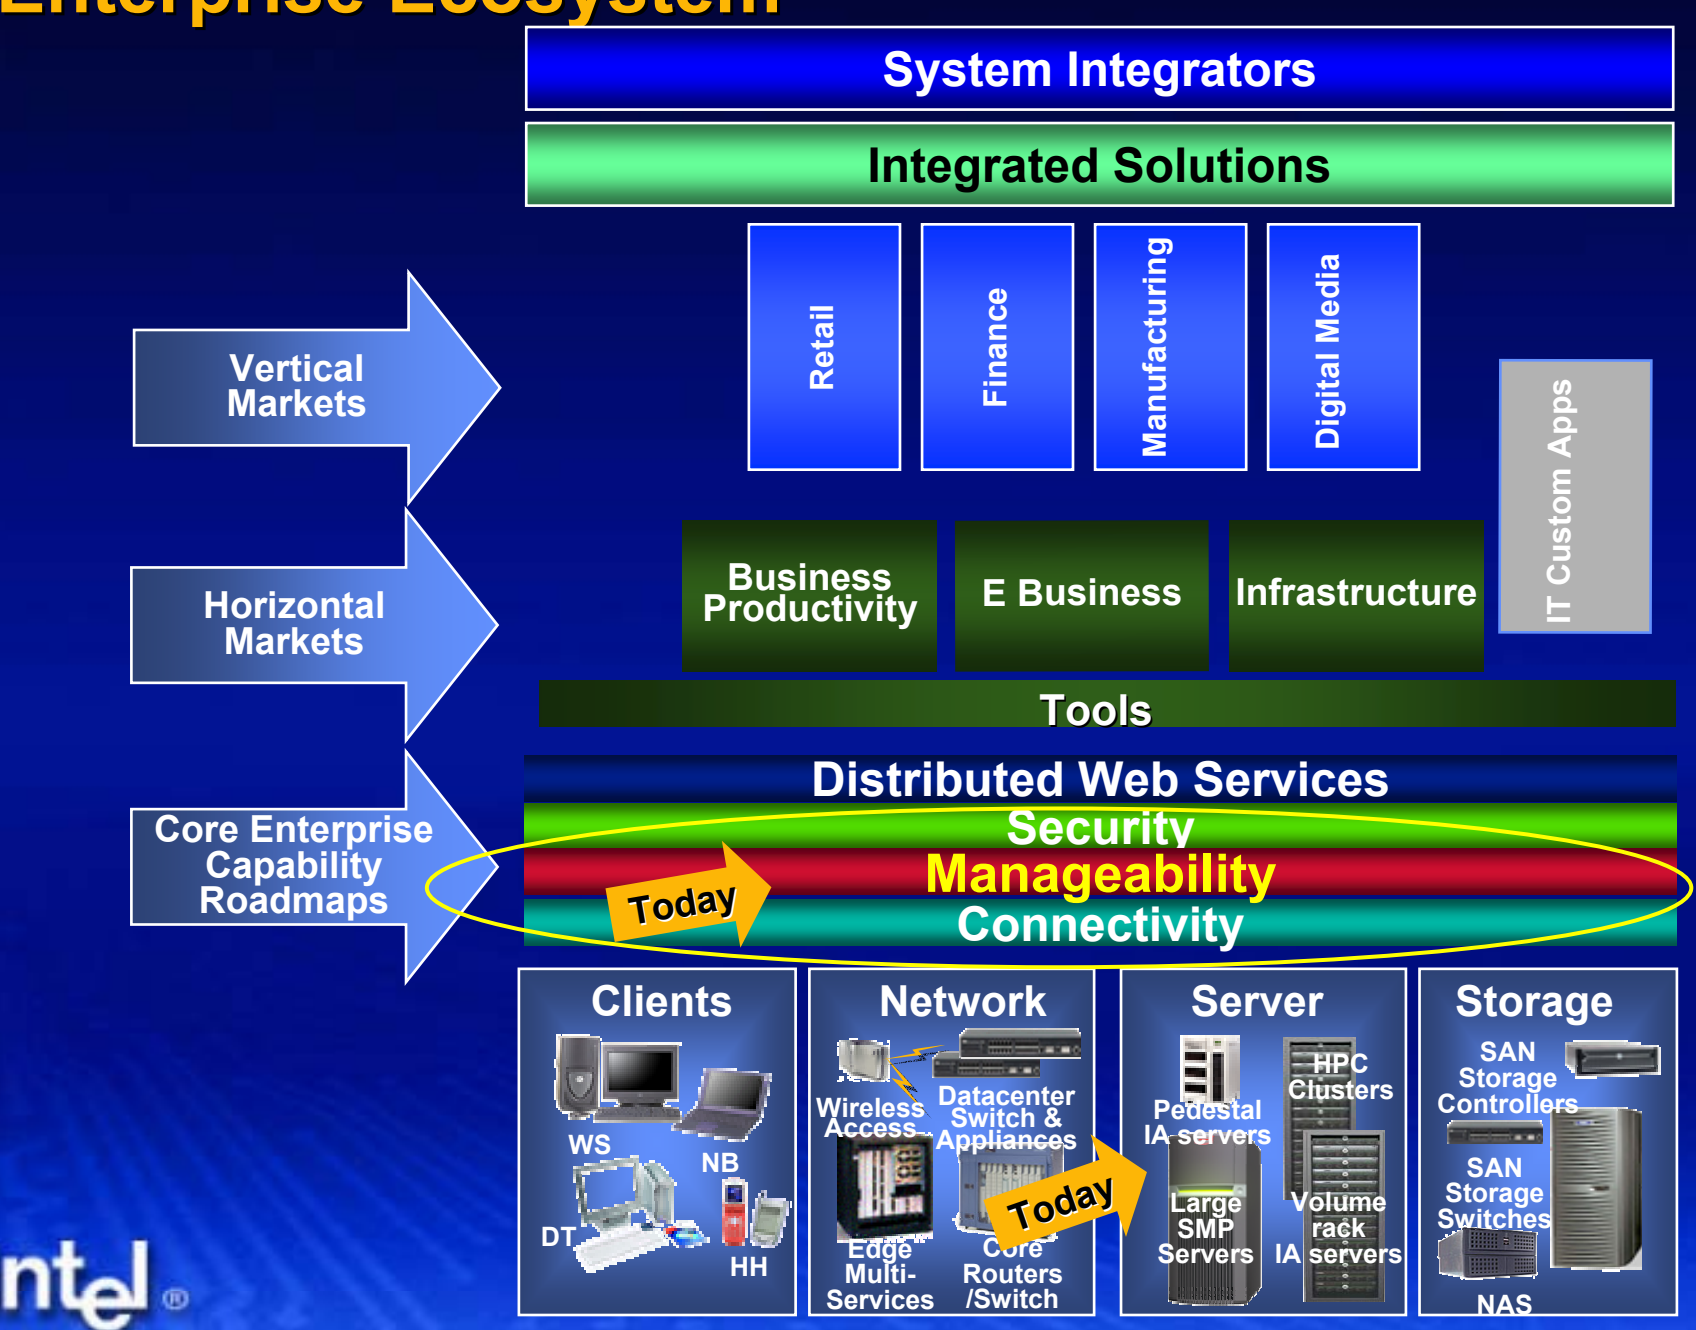

TCO Elements for Enterprise Services

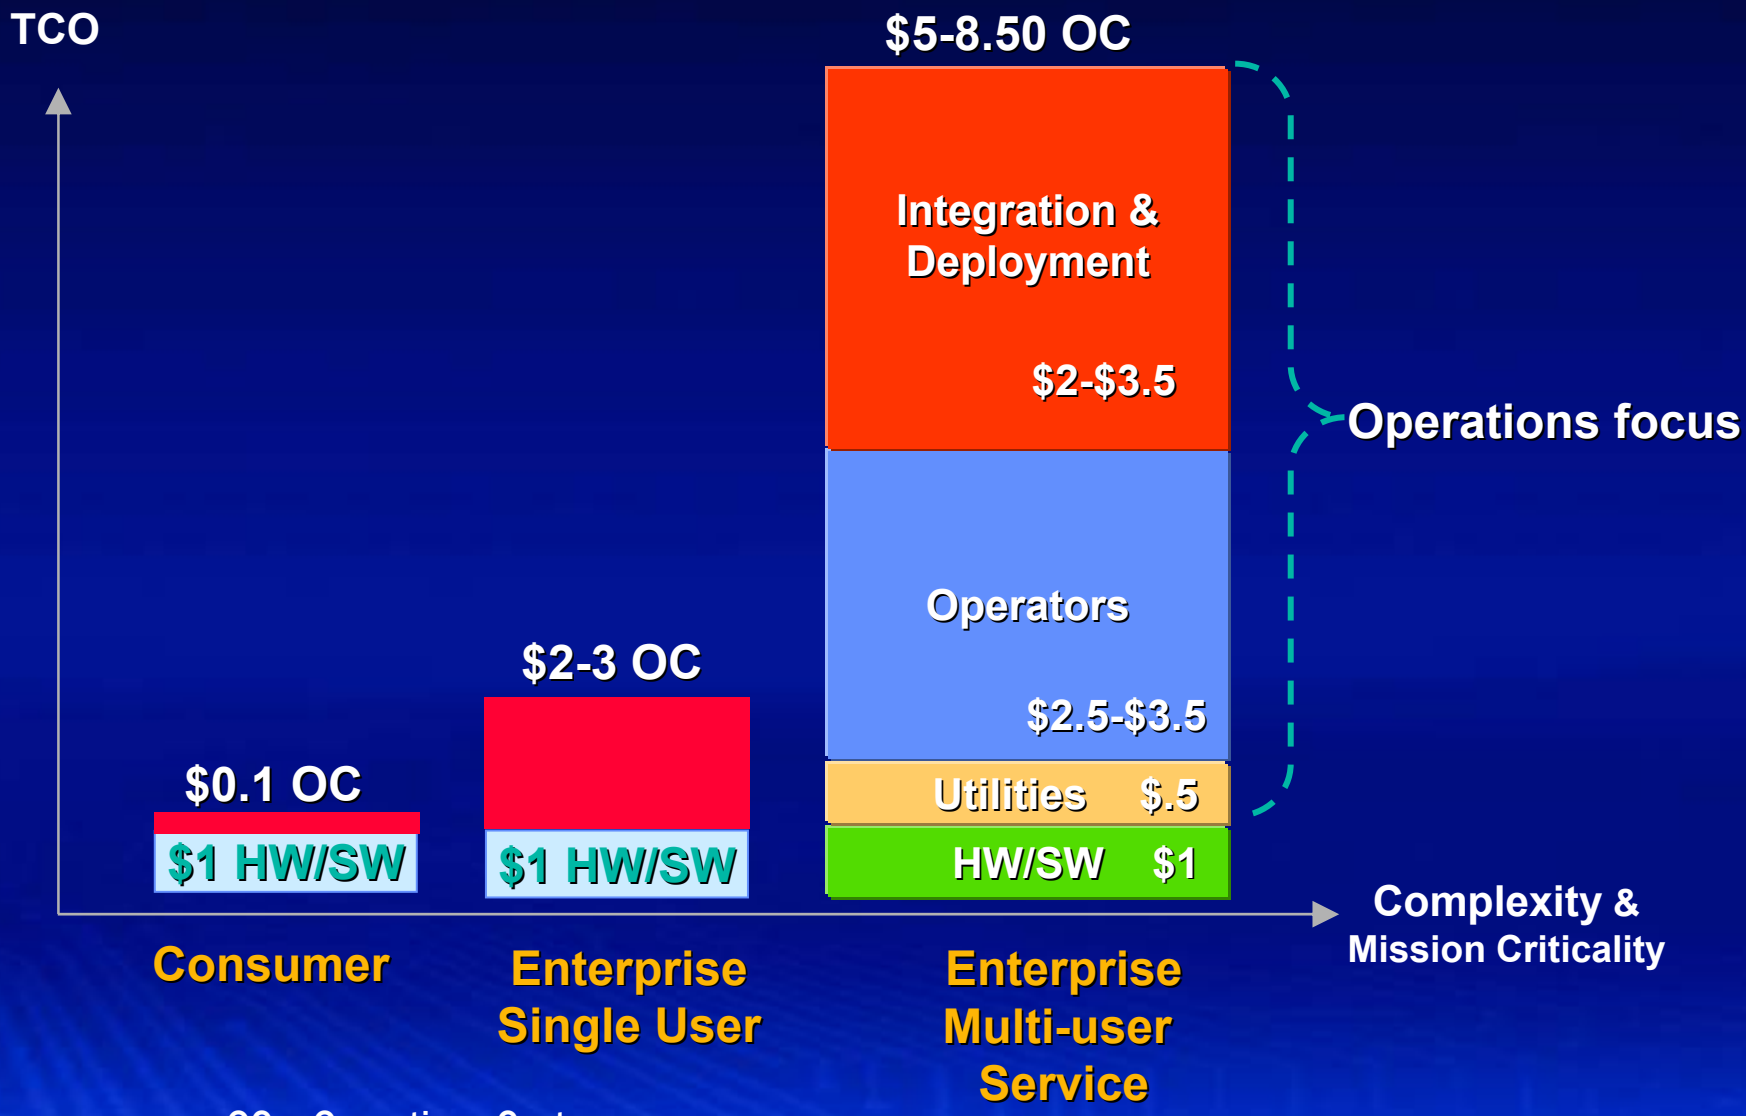

OC = Operations Cost
Source: Intel, Gartner studies

IT Manageability Life Cycle

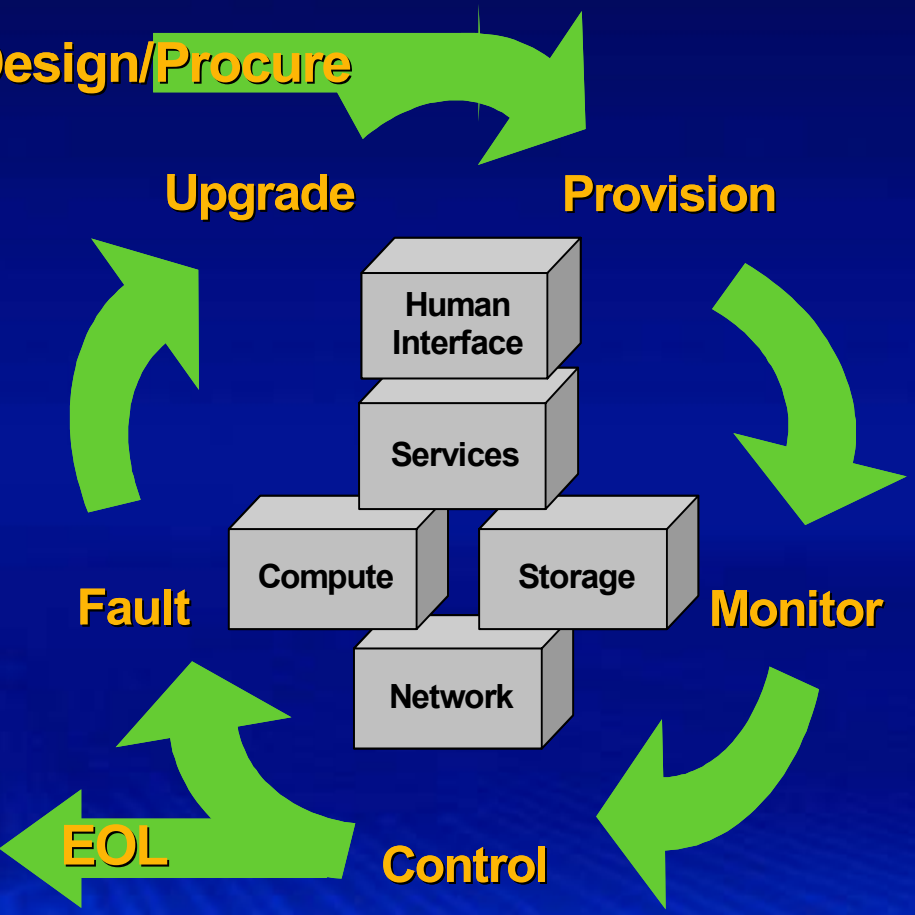

1. **Design**
   Define requirements and specify solution
2. **Procure/Deploy**
   Purchase systems as specified and place in system prior to starting operations
3. **Provision**
   - Discover and configure compute, storage and network elements to the operational state
4. **Monitor**
   - Discover, monitor, and alerting of the ongoing state, health and performance of services
5. **Control**
   - Regular and preventative maintenance and service optimizations
6. **Fault**
   - Preventing, predicting and recovering operational state from faults
7. **Upgrade**
   - Change management, version control, & system staging
8. **End of Life**
   - Remove from operation and dispose of material in appropriate manner

intel

Agenda

- **Enterprise landscape**
  - What do we have today
- **How can you manage this today**
- **Management trends**
  - What to expect in the future

Server Management Stack

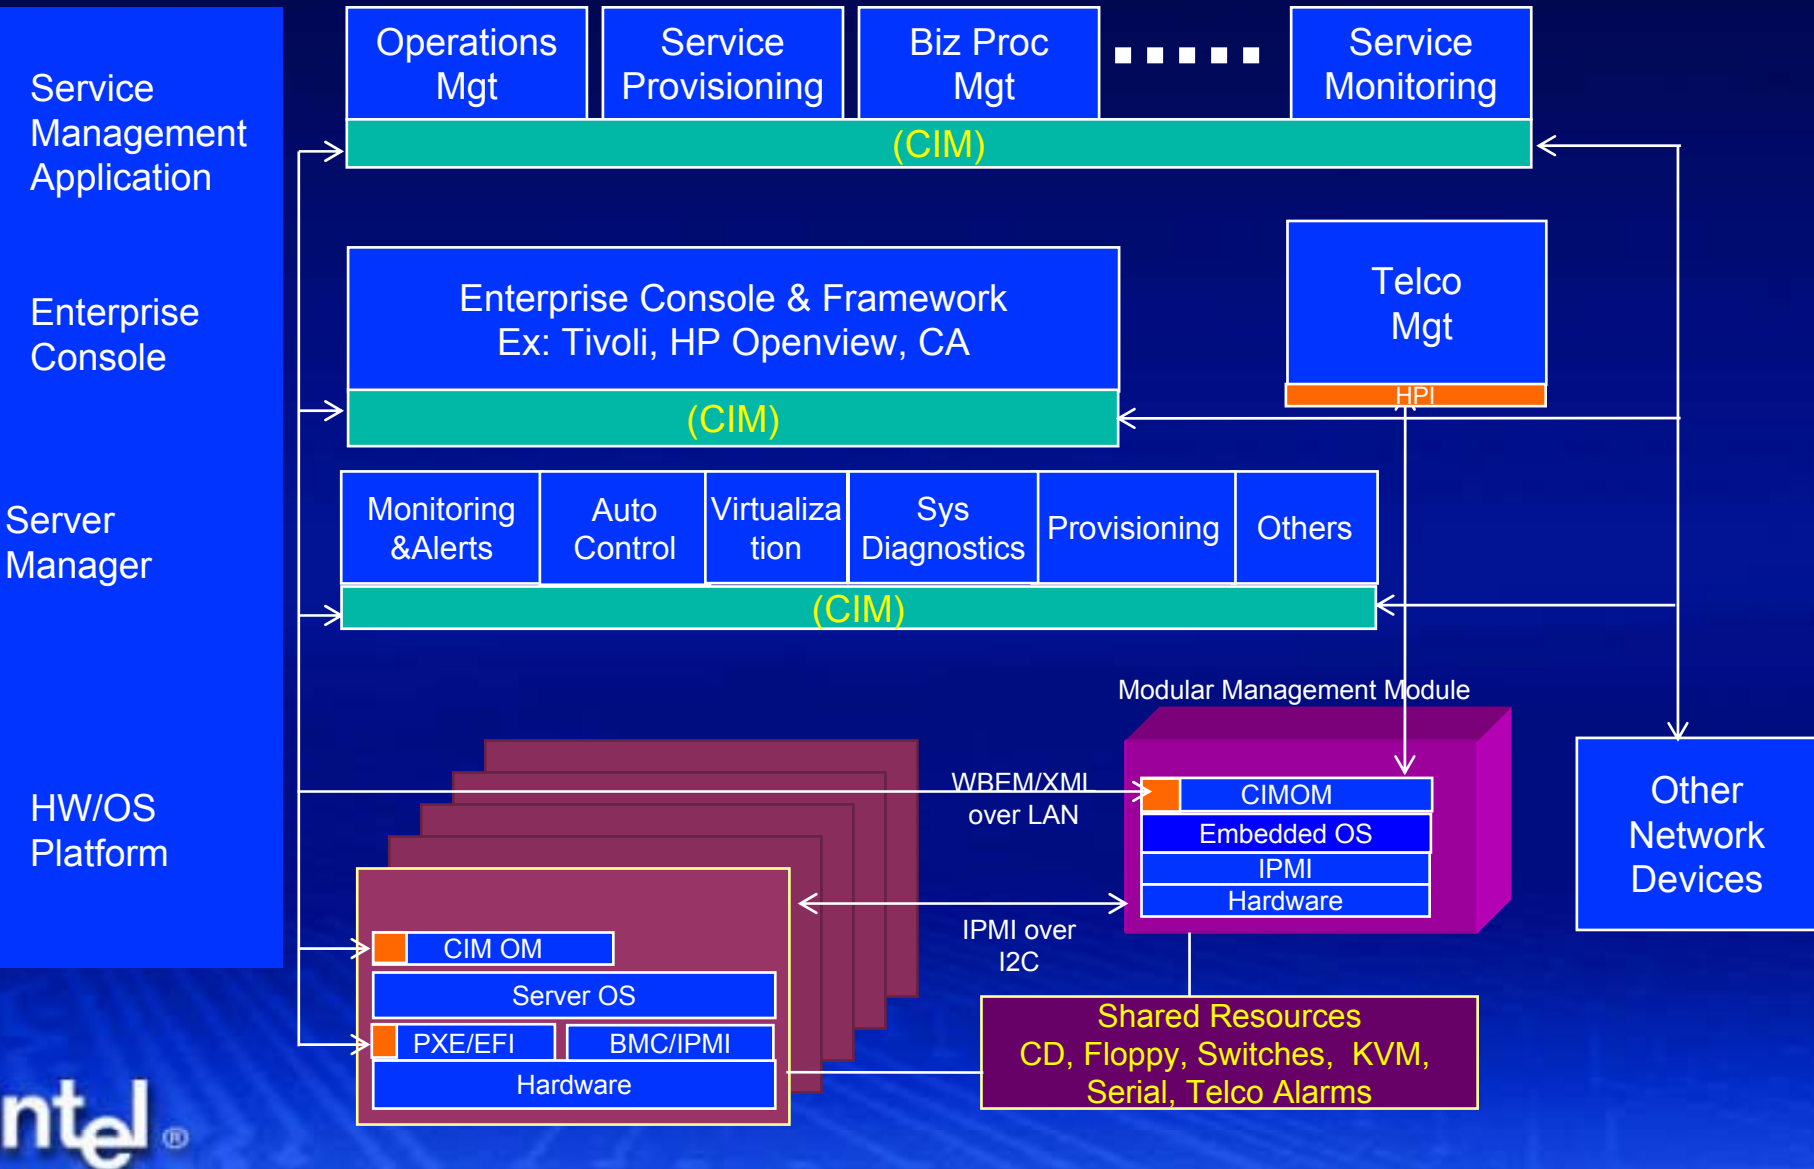

Building Blocks

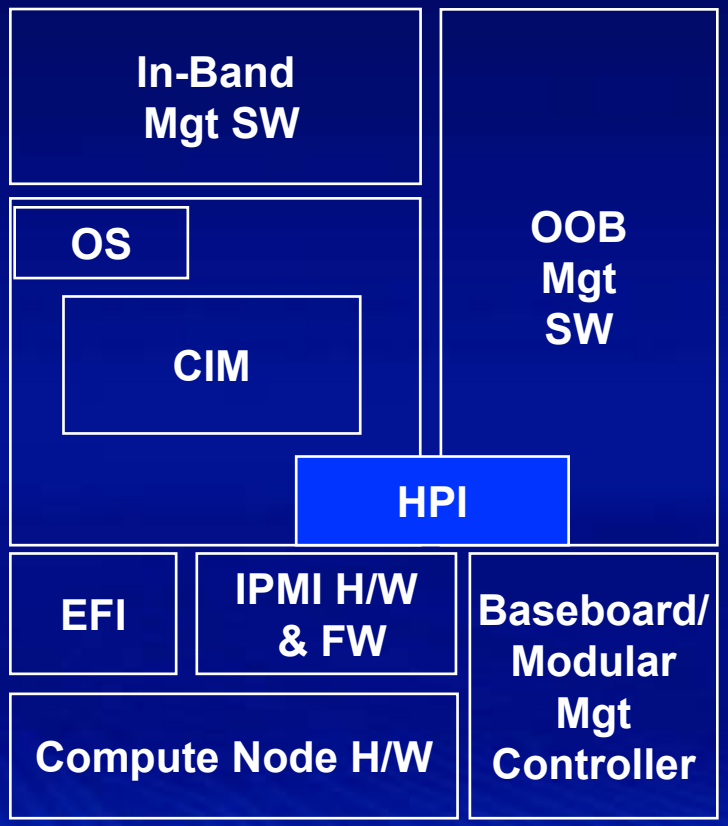

- **In-Band**
  - CIM
- **Out-of-Band**
  - LAN/Serial
- **Management IF**
  - CIM
  - IPMI
  - HPI
- **Management H/W**
  - BMC/MMC
  - EFI
  - PCI Express
  - InfiniBand* Architecture

* Other names and brands may be claimed as the property of other

Agenda

- **Enterprise landscape**
  - What do we have today
- **How can you manage this today**
- **Management trends**
  - What to expect in the future

Data Center Is Complex

Management can make it simple for the user

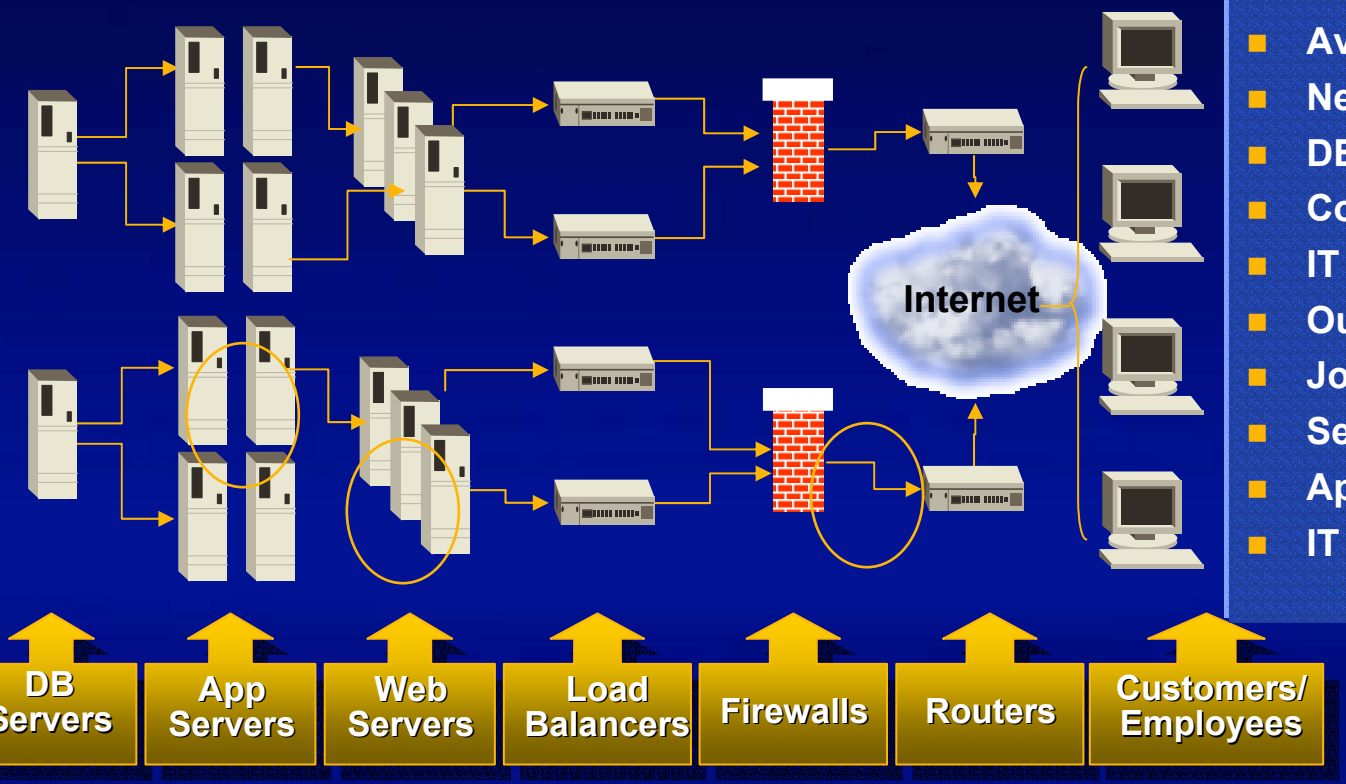

Management Sub-segments

- Availability & Performance
- Network Management
- DBMS Management
- Configuration Management
- IT Service Desk
- Output Management
- Job Scheduling
- Service Management
- Application Management
- IT Asset Management

Different Management Software needed to Manage complex Data Center

intel

Get some help from commercial Apps

Management Segment Map

Availability & Performance
- IBM
- BMC
- HP
- CA

Network Management
- HP
- IBM
- Network Assoc
- Cisco
- CA

DBMS Management
IBM
BMC
Oracle
Quest

Configuration Management
- IBM
- MSFT
- Symantec
- LANdesk

IT Service Desk
- Peregrine
- CE
- IBM
- Network Assoc

Output Management
- IBM
- Mobius
- Beta Systems
- CA
- HP
- BMC

Job Scheduling
- IBM
- CA
- BMC

Service Management
- IBM
- Concord
- Lucent
- InfoVista

Application Management
- BMC
- IBM
- CA
- NetIQ

IT Asset Management
- Peregrine
- CA
- Tally

intel

There is even more

Provisioning/Life Cycle Management

- Platspin
- Veritas
- CA
- HP

Cluster Management

- Veritas
- IBM
- Oracle
- NPCAI

Storage Management

- Veritas
- Legato
- CA
- BMC
- IBM

Write you own based on Building Blocks

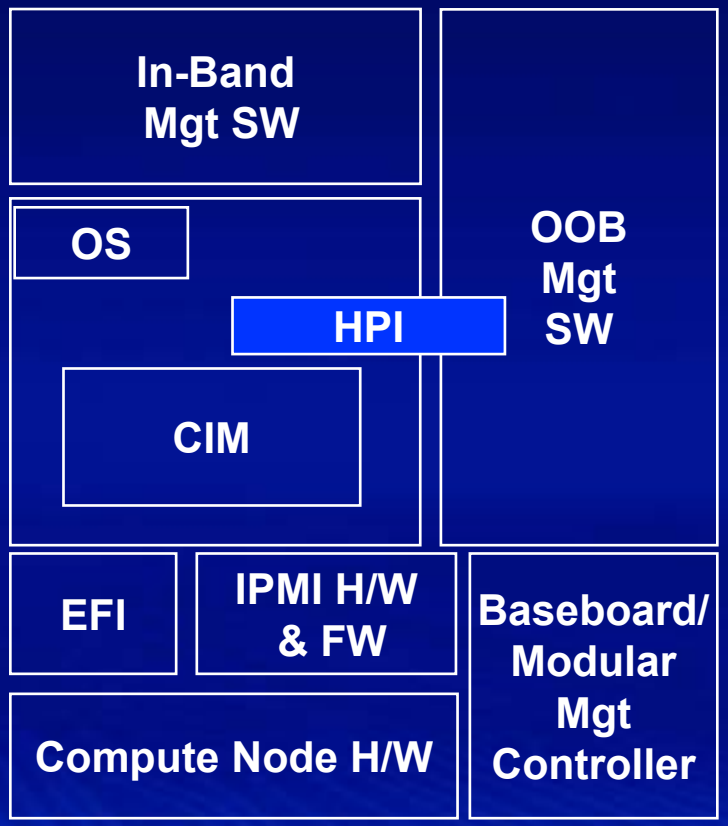

- **In-Band**
  - CIM
- **Out-of-Band**
  - LAN/Serial
- **Management IF**
  - CIM
  - IPMI
  - HPI
- **Management H/W**
  - BMC/MMC
  - EFI
  - PCI Express
  - InfiniBand* Architecture

* Other names and brands may be claimed as the property of other

Management Building Blocks - Purposes

Core building blocks simplify and integrates in-band and remote out-of-band management

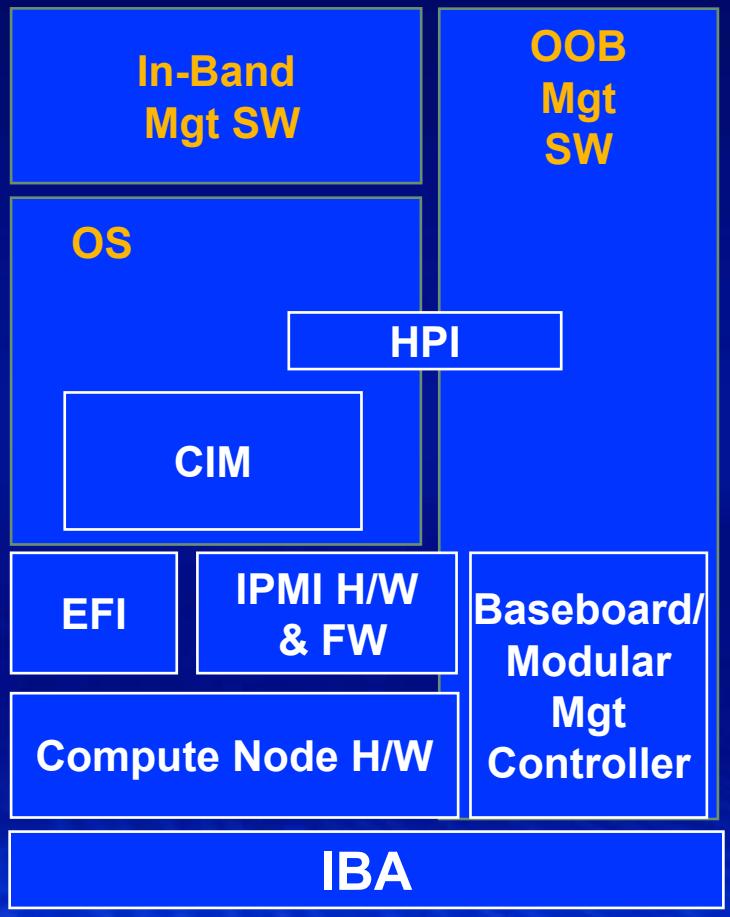

- **IPMI** is the platform instrumentation solution
- **EFI** is the preferred platform provisioning and virtualization solution
- **CIM** is the preferred in-band management management framework
- **BMC/MMC** is the central point of managing the modular server as a single unit
- **HPI** is the platform management API for Telco and non-CIM environments

InfiniBand* Architecture

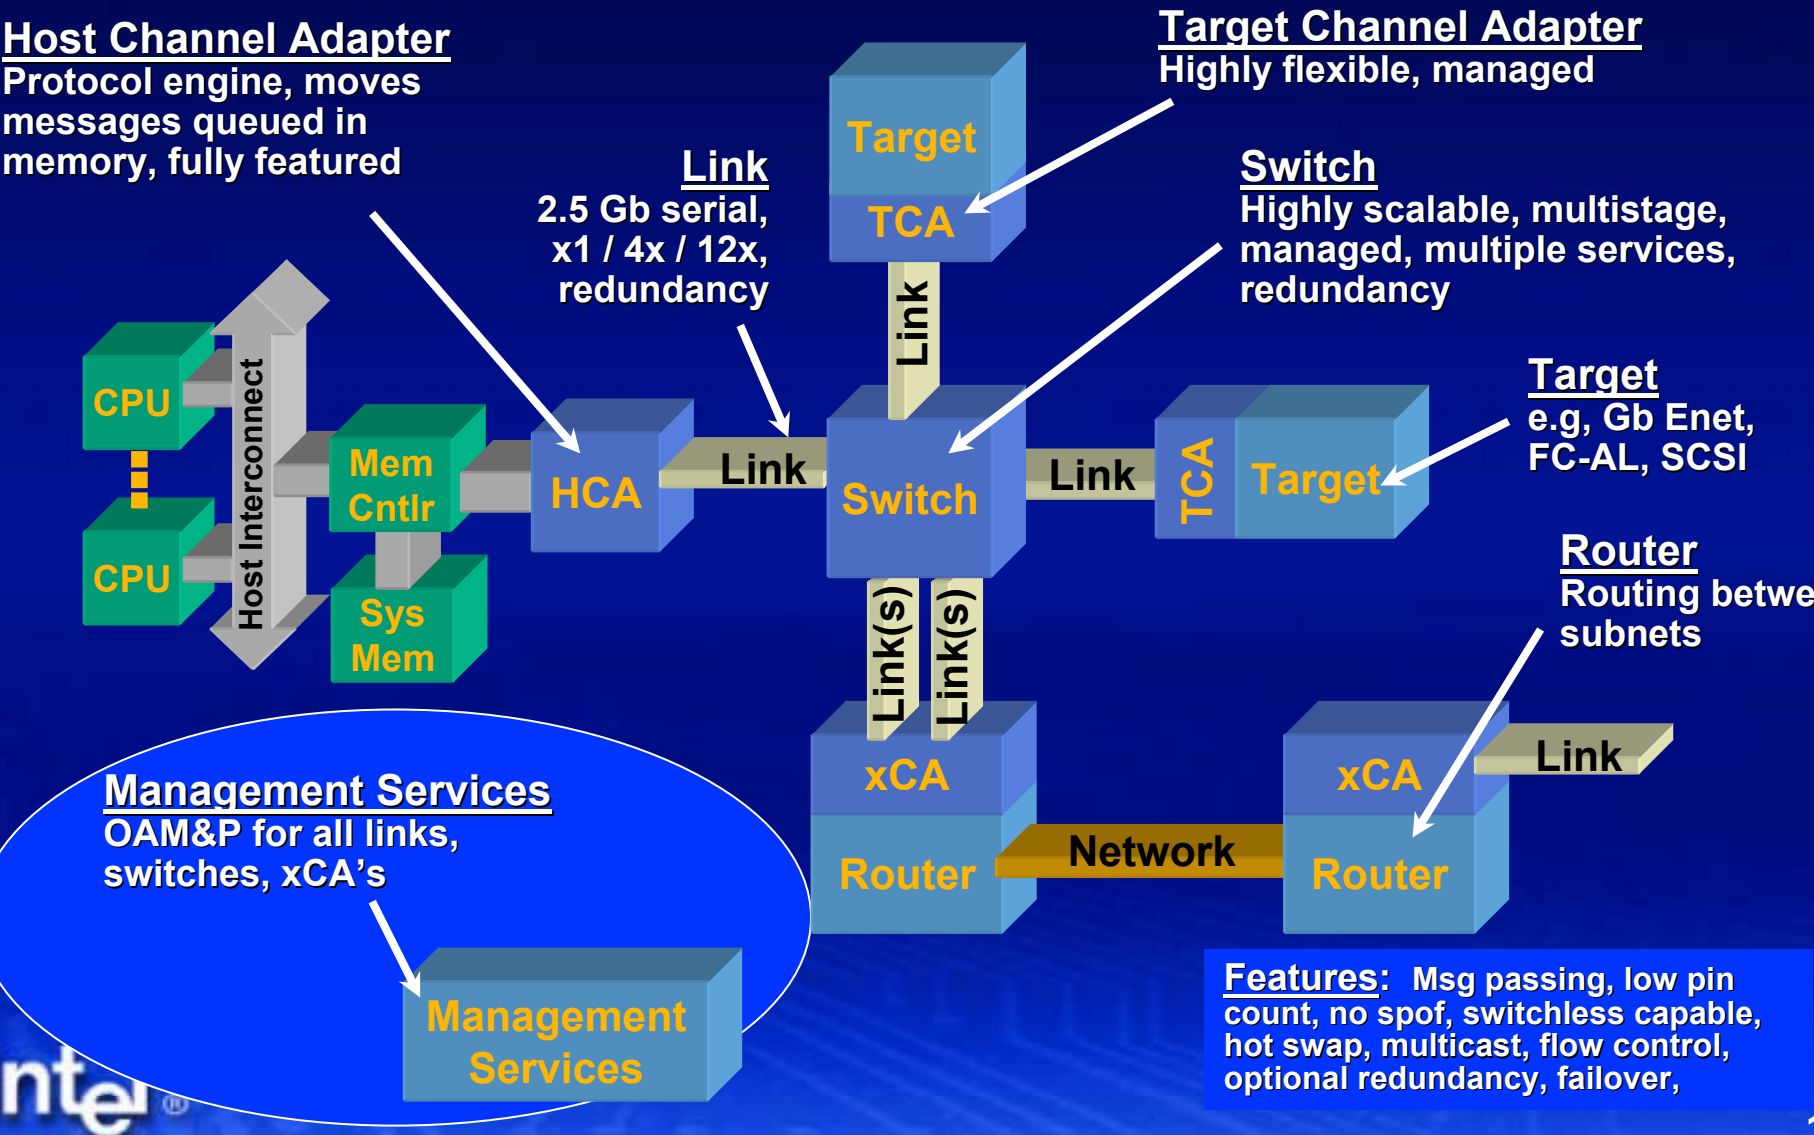

Subnet Management

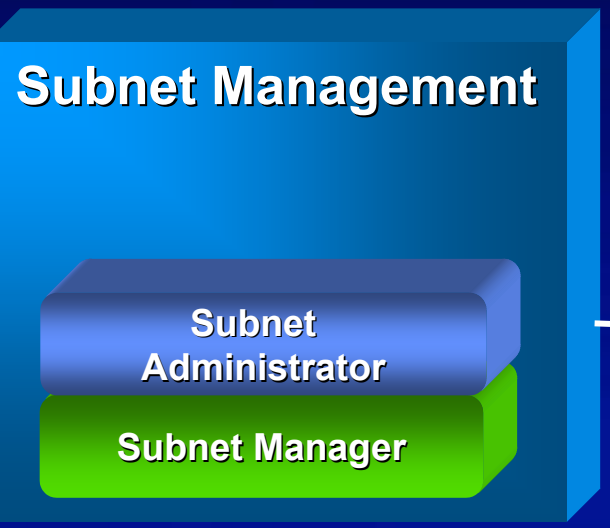

- Configures subnet paths through switches
- Discovers and configures devices as they appear on subnet
- Maintains event and multicast services

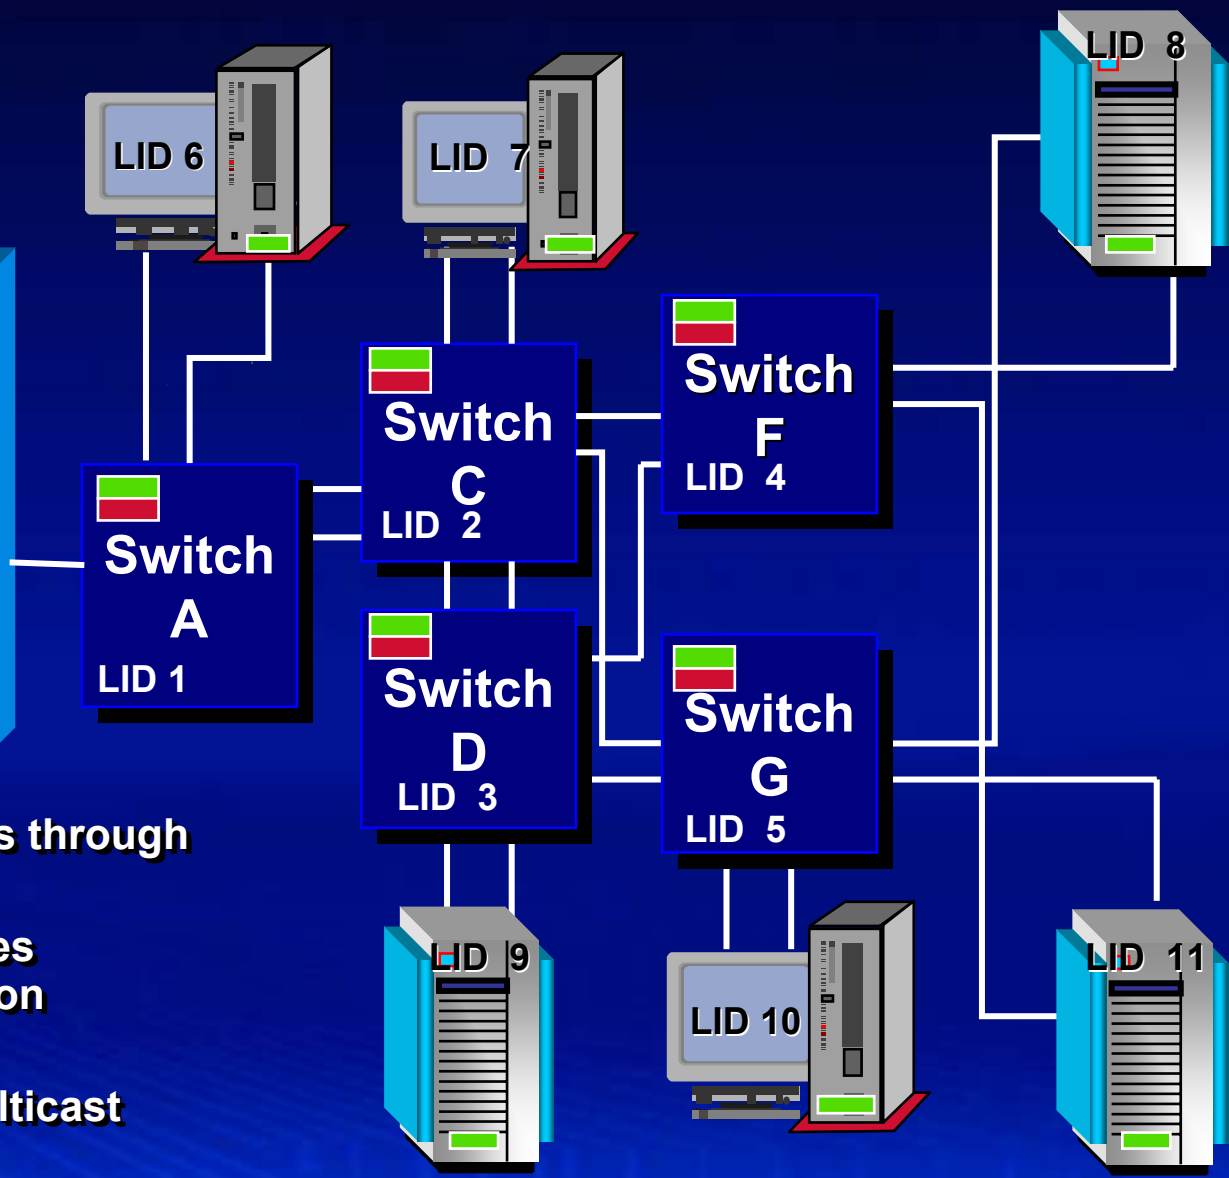

Management Model

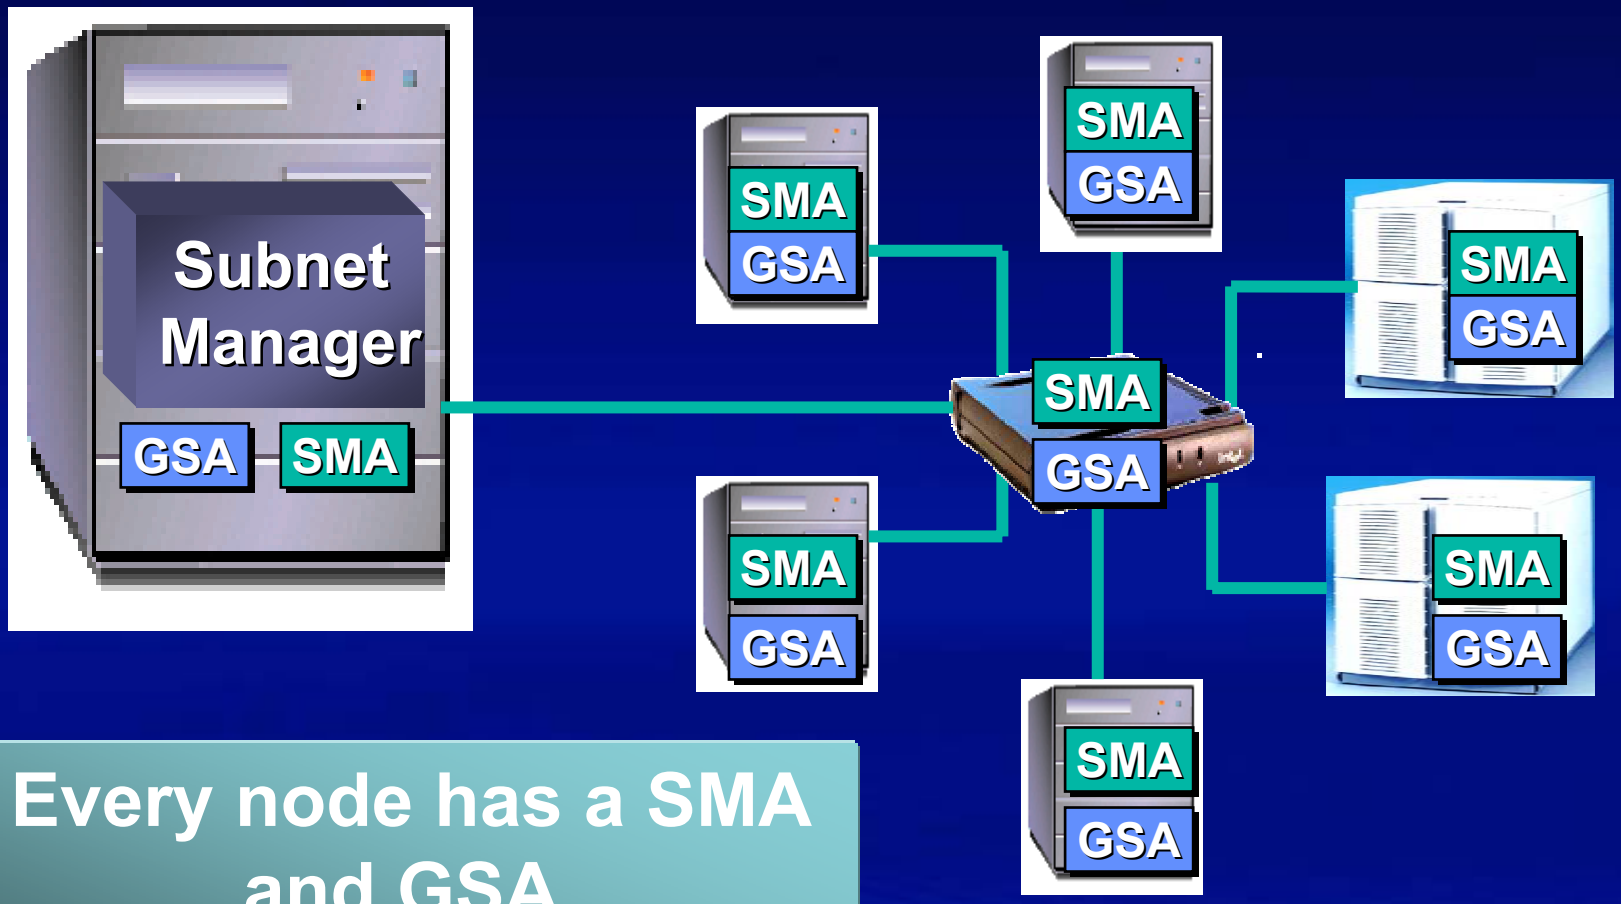

Every node has a SMA and GSA

Managers can be located on any node

SF IB Project Overview

- **SourceForge.net – Linux community openly collaborates**
  - http://sourceforge.net/projects/infiniband
  - Dual-license Open source licensing (GPL, “BSD+patent”)
- **Projects comprised of sub-projects**
  - Maintainers and developers
  - O/S distributors involved in development process
  - Sub-projects focus deliverables of interested parties
  - New sub-projects may be added at any time
- **Evolutionary development model**
  - O/S distribution occurs as development completes
  - Completion dependent on dev. Community involvement

High Level Architecture

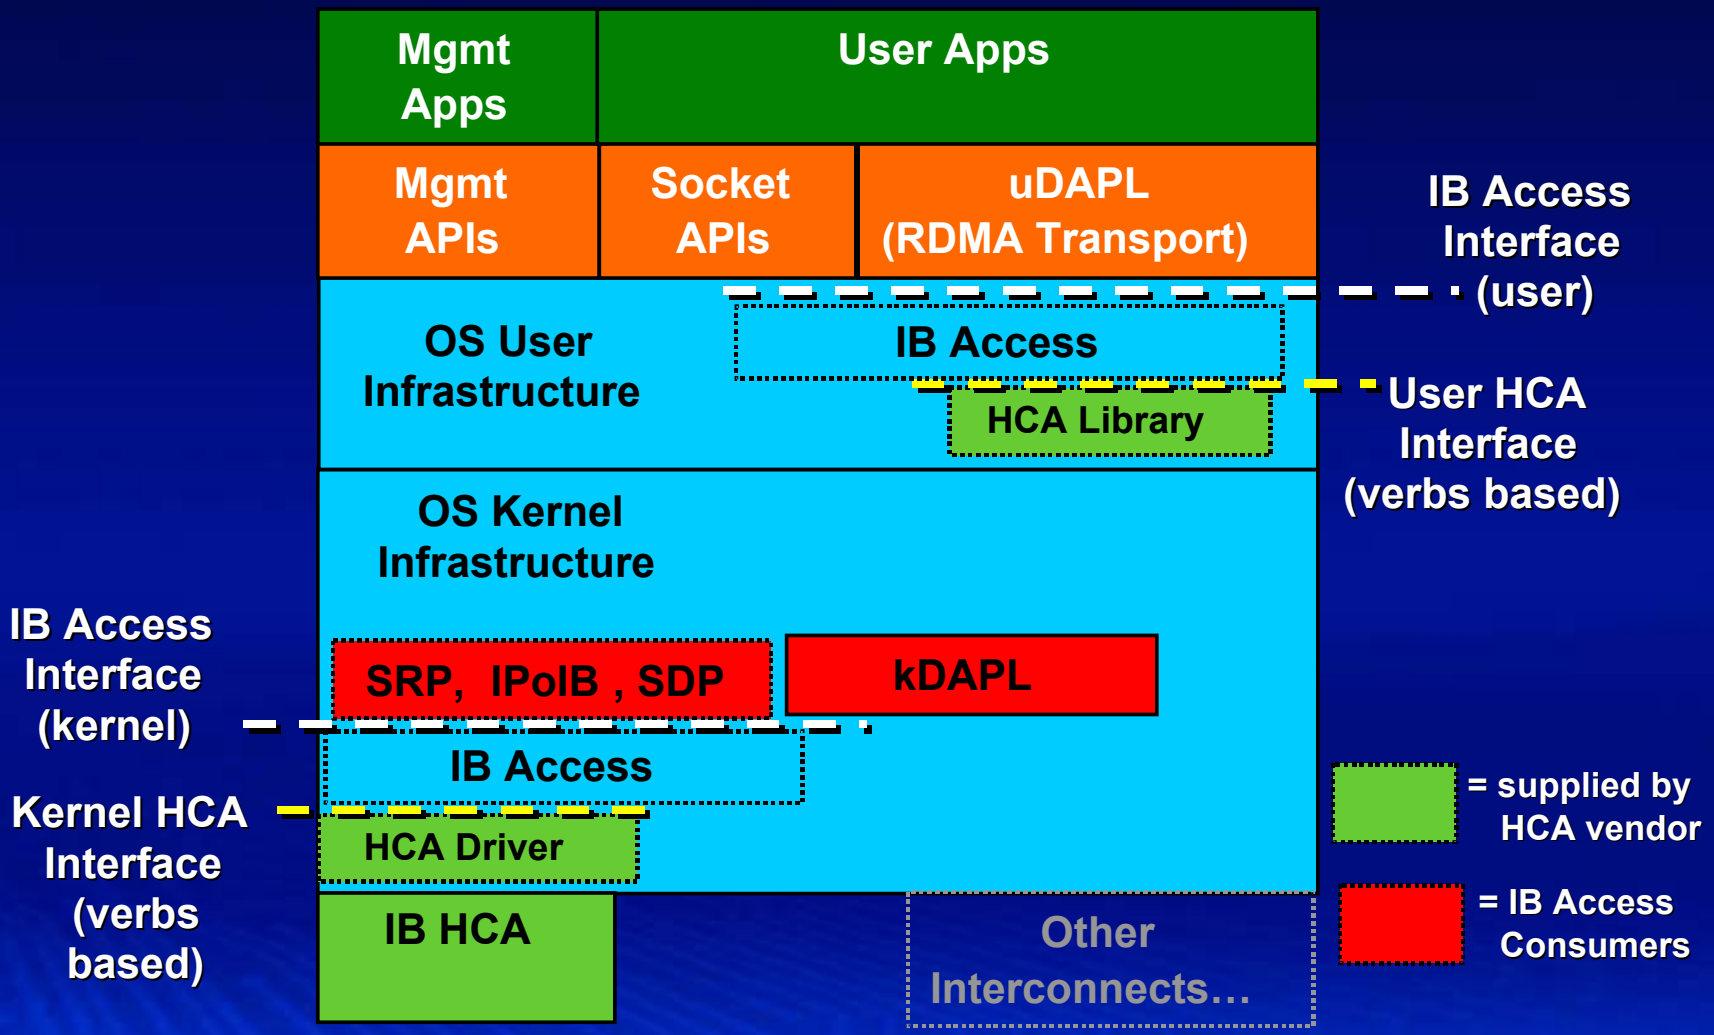

Agenda

- **Enterprise landscape**
  - What do we have today
- **How can you manage this today**
- **Management trends**
  - What to expect in the future

Enterprise Management Trends

Impact

- *Integration servers: ISVs*
- *OS/mgmnt tools interaction*

- **Enterprise End Users**
  - Demanding anytime, anywhere computing
  - Demanding transparent deployment of security, upgrades, management
- **IT Management**
  - Virtualization of compute and storage resources
  - Improve manageability and security of applications and environment
  - Greater demand for lowering TCO is driving demand for better integration tools
- **ISVs**
  - Service-centric infrastructure creating opportunity for new management app categories
    - Adaptive Provisioning/Mgmt, Service Level Mgmt, Workload Mgmt, Clustering (IDC)
  - Increased complexity and integration requirements driving demand for more robust management applications
  - Enterprise System Management revenues under pressure (Gartner)
    - Standards commoditizing “agent” capabilities
    - Less expensive distributed licenses replacing mainframe licenses ($$)
  - Both Server Consolidation and Scale out creating opportunities for System and Storage Mgmt vendors

intel

New Technology
Impacting Manageability

- EFI
  - Pre-OS management
  - Virtual machine agents
  - Diagnostics
  - Event Agents
- Tiano
  - XML interface
  - IPMI hooks
- PCI Express
  - Virtual Channels, Traffic Classes (QOS)
  - Hot-Plug/Swap and Surprise Removal
  - Enhanced Configuration and Power Management
  - Advanced error logging/reporting
- InfiniBand
  - IO devices directly visible to Mgmt
  - Boot over Fabric

New Technologies brings new Management capabilities

Manageability Direction

To drive internal activities and engagement with ISVs

- **Impact lower TCO by enabling new capabilities in management software**
  - Remote management across geographies
  - Manage services not just servers
  - Support proactive management (automated response to events)
  - Automate system provisioning
  - Automate system operations tasks
  - Simplify distributed, scaled-out systems management
  - Integrated, end-to-end management framework across HW, OS, application and network
  - Up-level focus of management to business services NOT just components (servers, OS, applications, M/W)
- **Impact:**
  - Addresses the #1 end user need: Complexity & TCO

intel

Agenda

- **Enterprise landscape**
  - What do we have today
- **How can you manage this today**
- **Management trends**
  - What to expect in the future

IT Manageability Life Cycle

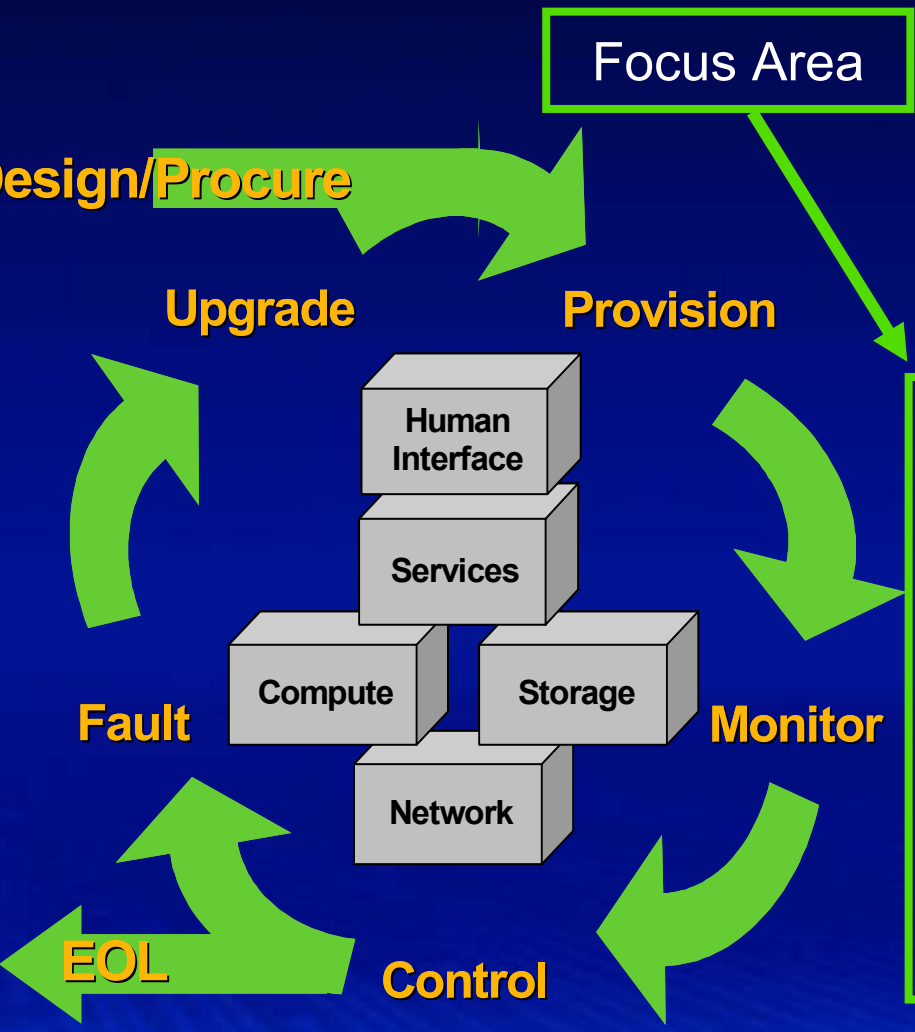

1. **Design**
   Define requirements and specify solution
2. **Procure/Deploy**
   Purchase systems as specified and place in system prior to starting operations
3. **Provision**
   - Discover and configure compute, storage and network elements to the operational state
4. **Monitor**
   - Discover, monitor, and alerting of the ongoing state, health and performance of services
5. **Control**
   - Regular and preventative maintenance and service optimizations
6. **Fault**
   - Preventing, predicting and recovering operational state from faults
7. **Upgrade**
   - Change management, version control, & system staging

End of Life

- Remove from operation and dispose of material in appropriate manner

Vision: Self-management of heterogeneous pool of virtualized IA compute, network, & storage resources in an industry-consistent manner

The "Virtualization" Landscape:
What needs to happen…

- **Single management user interface**
  - Common management across operating systems
- **OS and applications provisioning**
- **Repurpose nodes to meet changing workloads**
- **Automatic configuration**
- **Failure prediction, 'self healing' nodes**
- **Failure resolution**
- **Remote**
  - OS and Software Applications upgrades
  - flash BIOS/Firmware upgrades

Virtualization needs to start at the platform level through common management building blocks

Server Management Stack

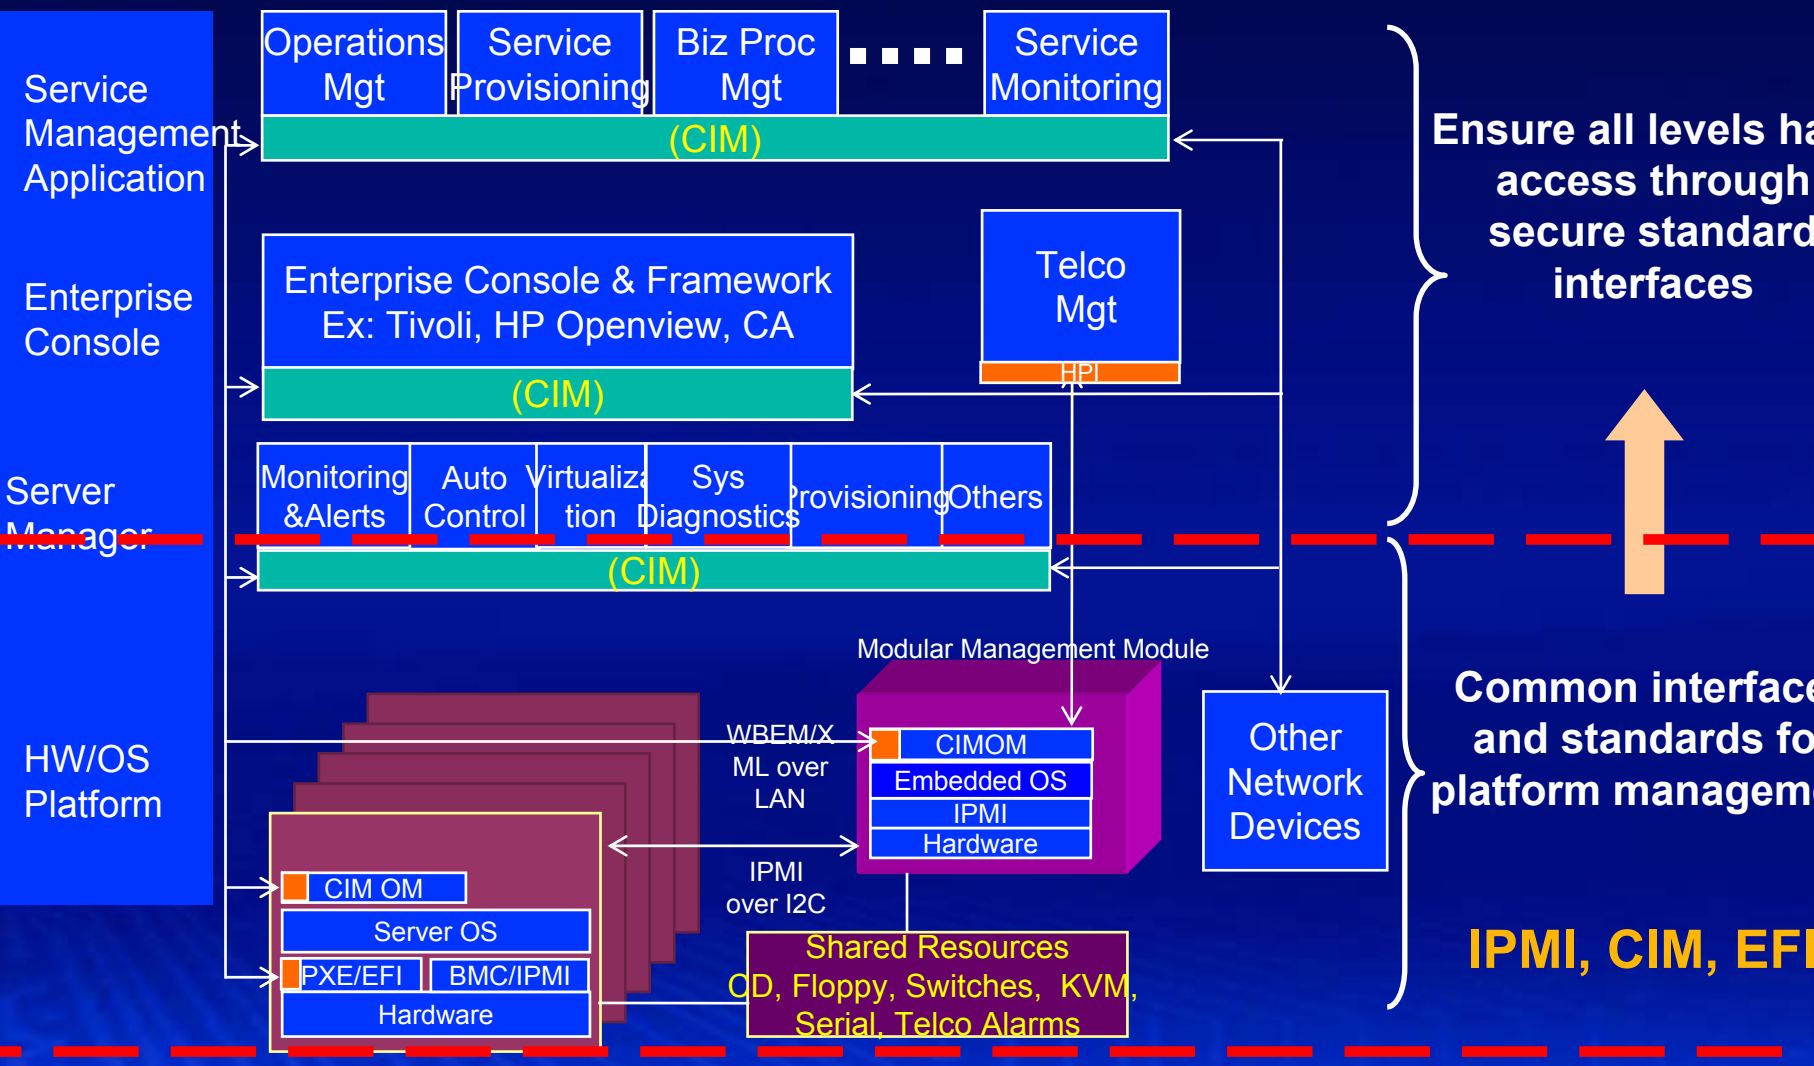

Intel's focus at the platform level

Intel's Drive towards the Virtual Data Center:

Platform Technology Leadership

CPU

- Non-volatile storage and logging
- Sensor configuration (initialization agent)
- Failover between two LAN ports
- SMbus (for sensors and IPMB channel)

Chipset

- Memory mirroring
- PCI Express Hot Swap
- Intel® x4 Single Device Data Correction (x4 SDDC)

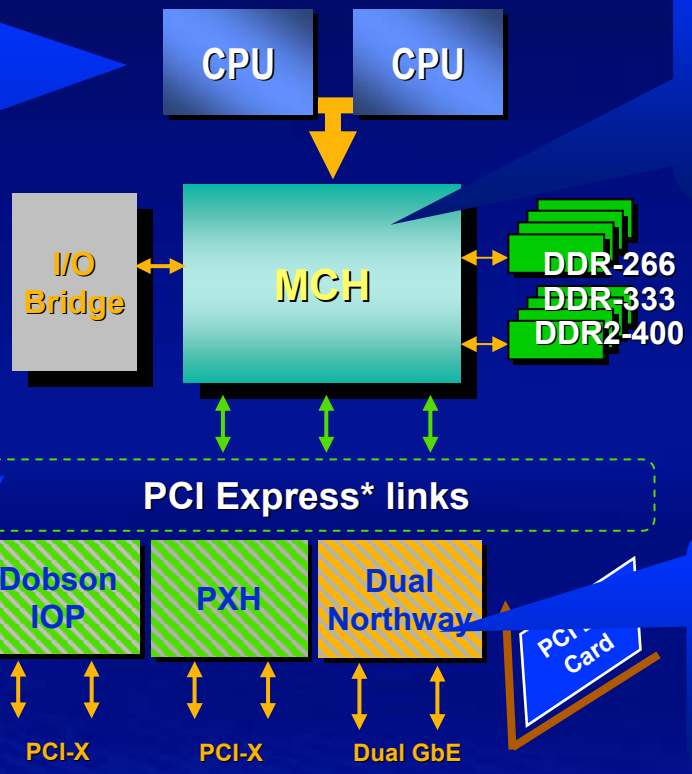

PCI Express

Extended error handling
Hot plug/ Hot Swap
Extended configuration data (on the cards)
Several more error registers built into architecture

LAN Controller

- ASF 1.0/2.0
- TCO bypass
- IPMI 1.5
- Baseline BMC

Significant management capability, all managed through IPMI, is built into Intel platforms in 2004

Summary

- **Data Center are complex**
  - Labor largest % of data center expense
- **Commercial apps available in many segments**
- **Write own apps**
  - Use industry Standards
    - CIM
    - IPMI
    - HPI
  - Take advantage of New Technology
    - EFI
    - Tiano
    - PCI Express
    - InfiniBand Architecture
- **Industry trends**
  - Virtualization of resources
  - Automate as much as possible
  - More manageability built in

Backup

Working Propositions

Manageability is critical to bring down TCO

- Market analysis studies show this as a key issue
- Dynamic provisioning, “Service-centric computing”, and Security mgmt must be easier (i.e., approaching automatic) to cause IT managers to upgrade/expand systems and software
- Large companies (Schwab, Morgan Stanley, etc.) are asking us for help
- Sun using mgmt features to market against IA
- Microsoft recognizes the problem and is working to address it
- **Messages:**
  - **Intel’s enterprise ramp strategy is gated by the need for robust management capabilities, standards**
  - **IPF success is gated by availability of antivirus, system management apps, utilities**

Manageability is critical to modular computing

- **Manageability enables Enterprise Modular Computing**
  - **Availability - Provisioning**
  - **Application management - Security**
  - **Scalability - Performance**
- **End goal: Single management tool for system, cluster, OS, applications**

Management Fit