IBM Systems & Technology Group

Helping your customer understand what system is the Right Fit.

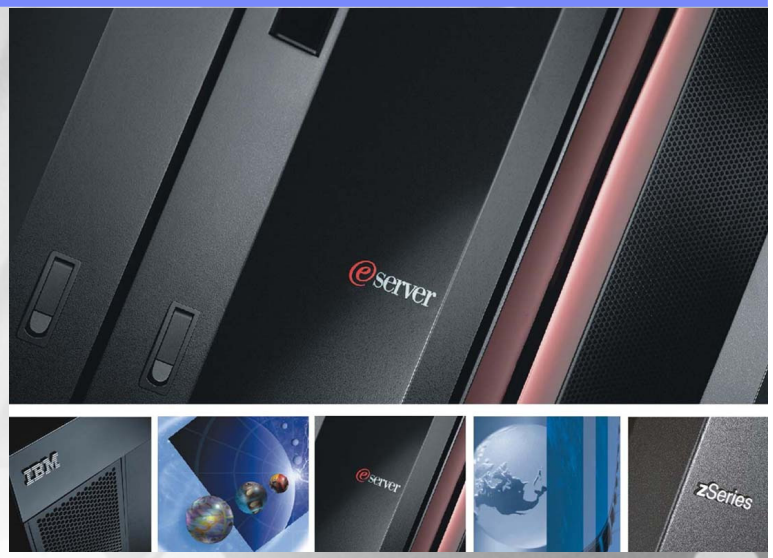

Irby Bilsky
WebSphere z Sales Specialist

ON DEMAND BUSINESS™

Agenda

- **Right fit for workloads-The basic elements**
  - Non-Functional requirements
  - Architectural Considerations/Linguistics etc.
  - TCO
  - Conclusions
- **Questions**

How do companies select a platform for their applications?

- **First question is**
  - "Will it run there?"
- **Second question is**
  - "How much does the hardware cost?"
- **Done!**
- **But……Is that all we should be thinking about?**

What did we miss? Non-Functional requirements

- **Shouldn't they have asked some questions about:**
  - Availability? Backup? Site Disaster Recovery?
  - Space? Power? Cooling?
  - Operations? Monitoring? Server management?
  - Integration?

- **So, what are these non-functional requirements?**

Non-Functional Requirement Methodology

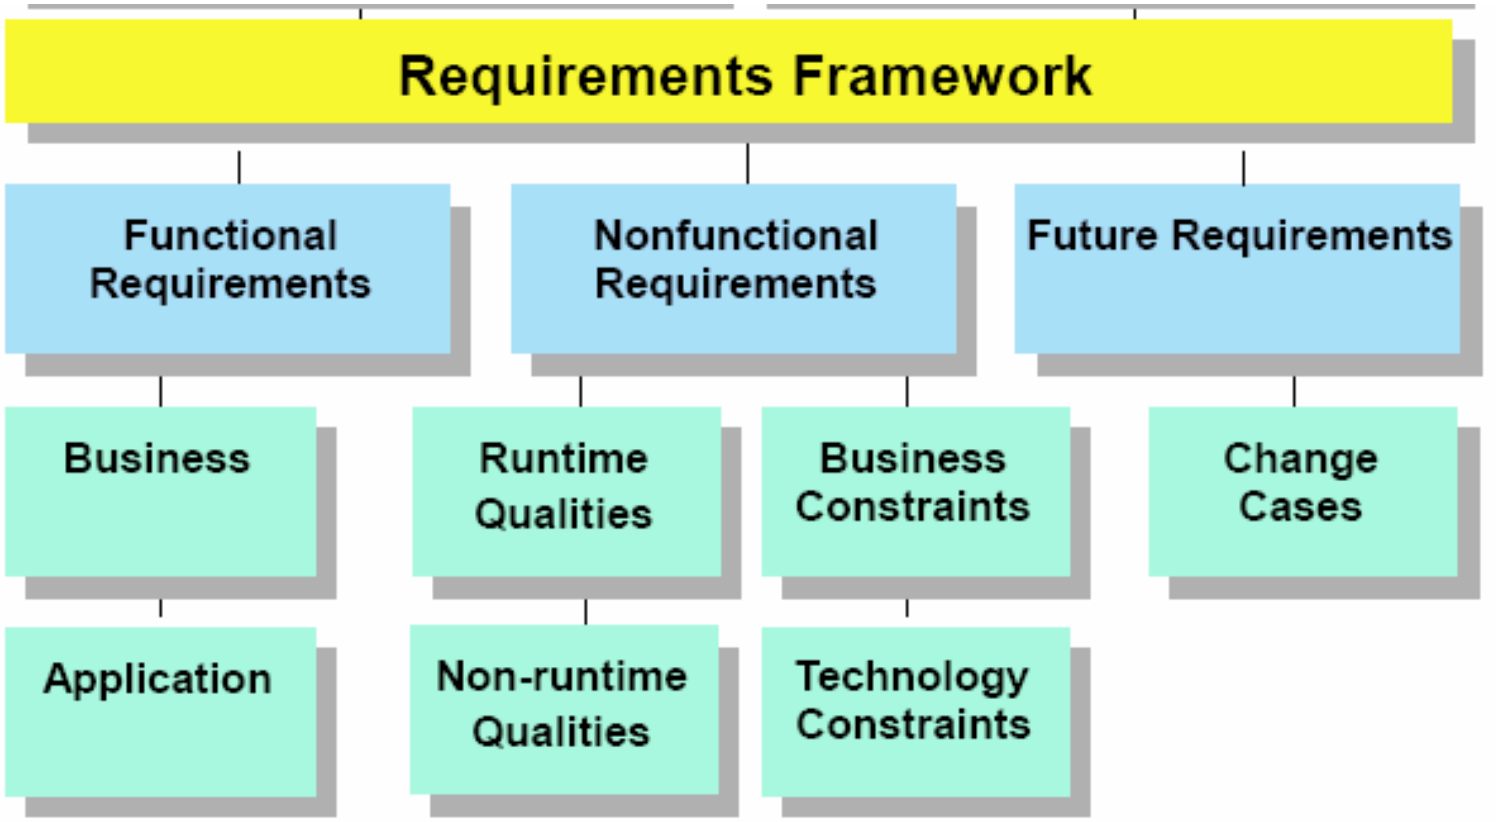

Non-Functional Requirement Methodology

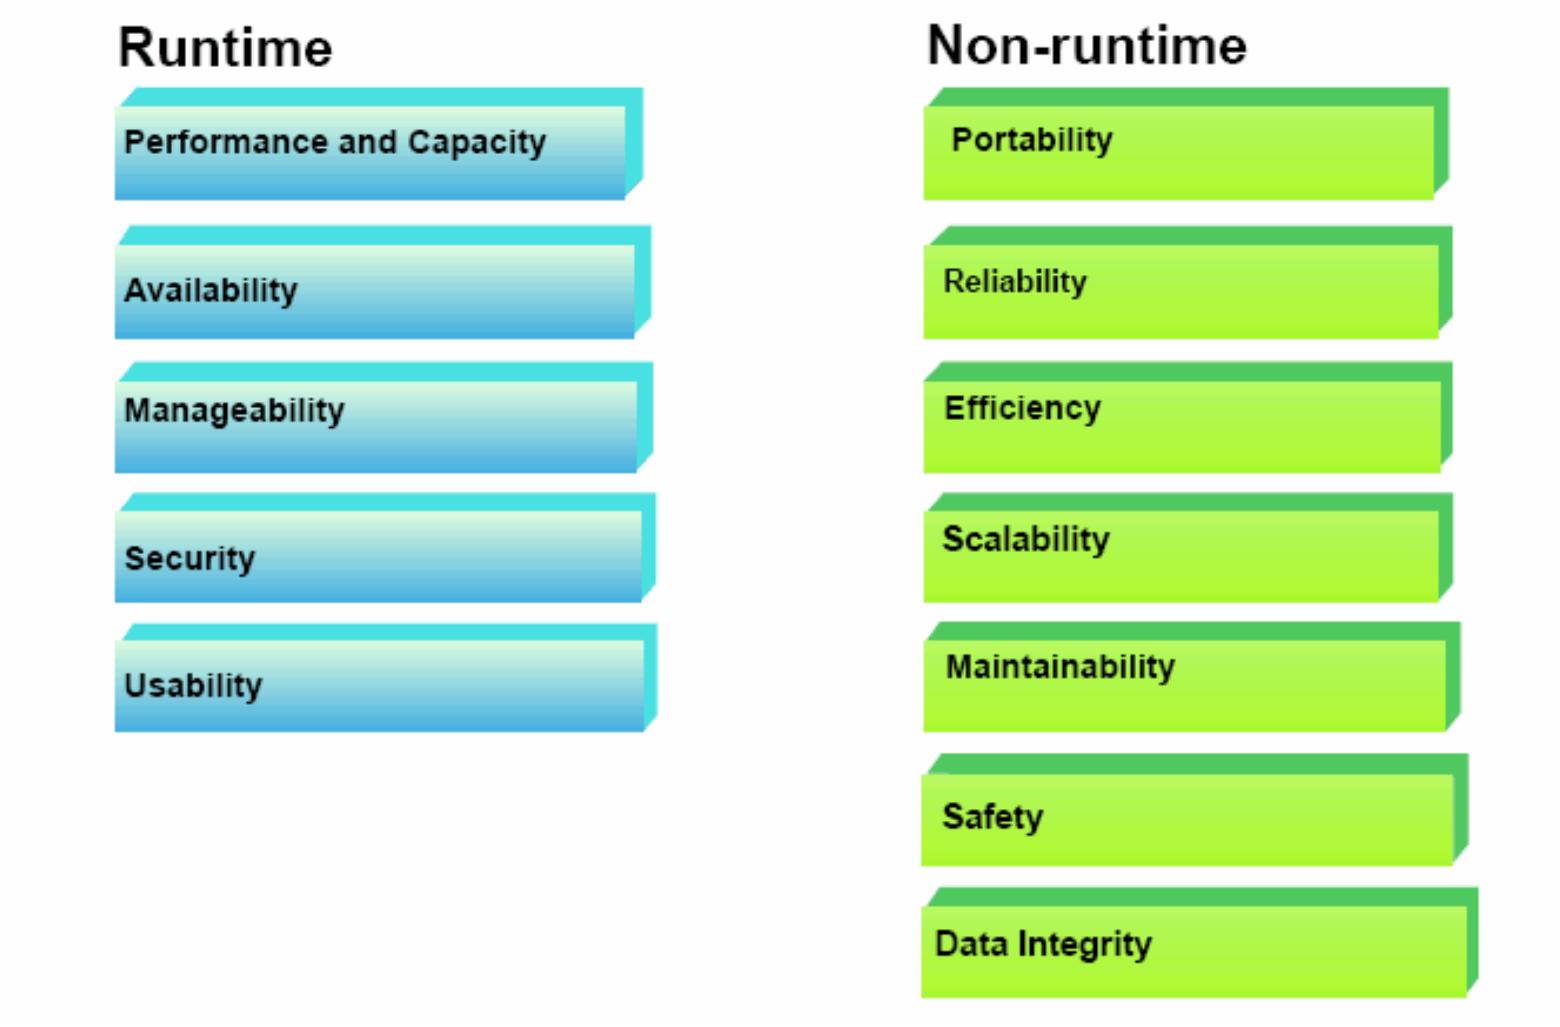

Non-Functional Requirement Methodology

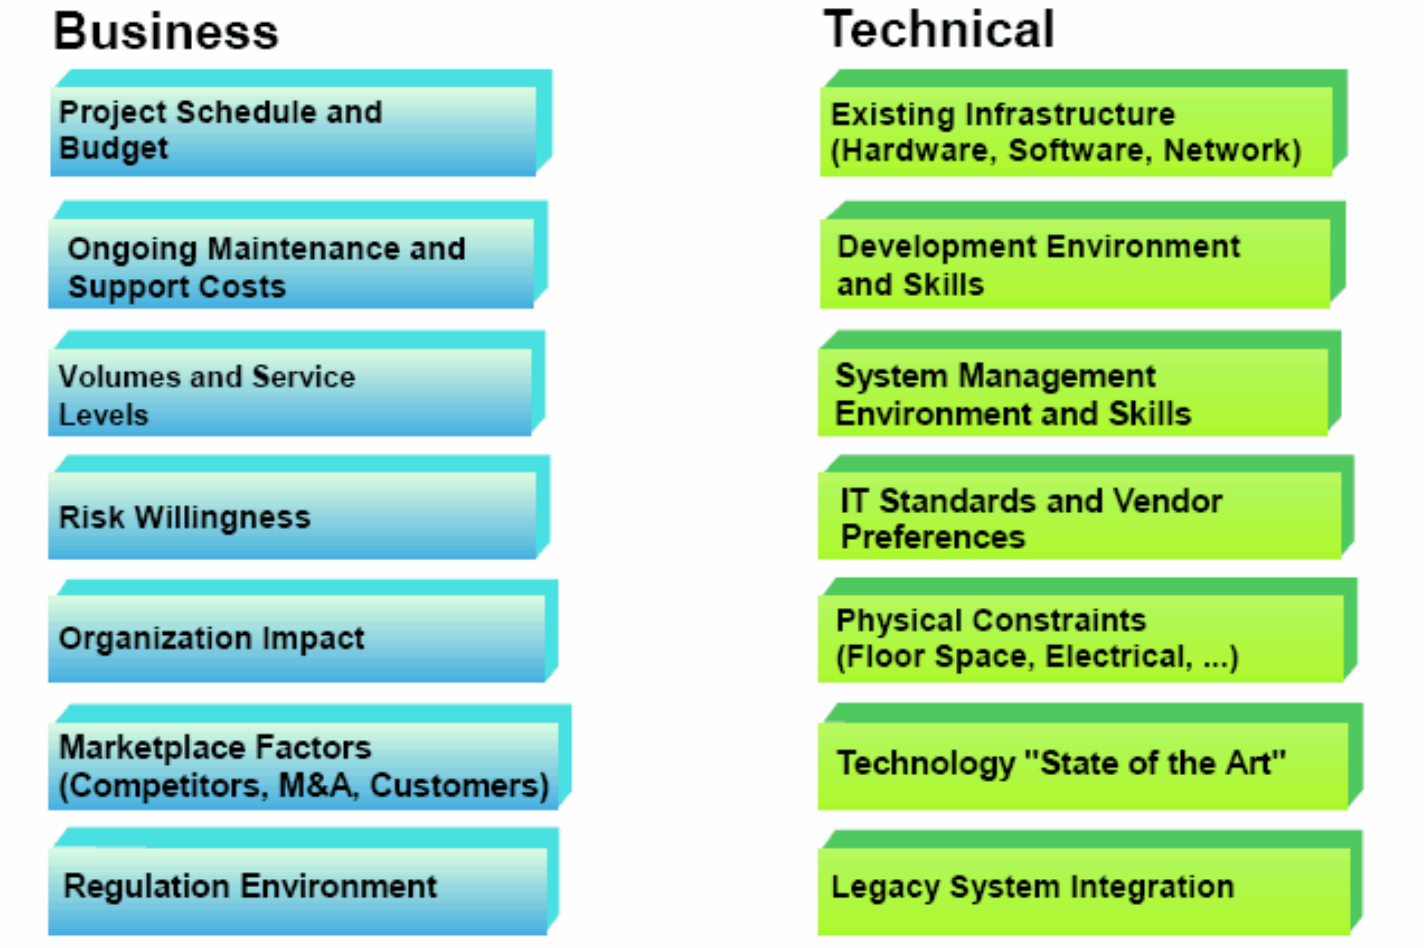

Don’t all platforms allow you to take these issues into consideration?

Good question……Let’s see if we can find a good answer!

The right ‘tool’…All of these tools can move a person from one place to another…real fast….

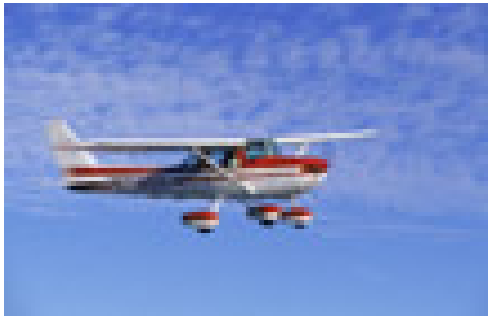

But…which is the right tool… to move 1 person? 100 people? 400 people?

Platform Heritage plays a significant role…….

S/360 (aka System z)

- **Initial design point**
  - Multiple users
  - Mixed workloads
  - Chip balance – I/O

Large costs = need mixed workloads and maximum utilisation

RISC/Unix

- **Initial Design point**
  - Workstation (Single user)
  - Compute intensive
  - Chip balance - compute

Low cost = single user

WinTel

- **Initial Design Point**
  - Personal computer (single user)
  - GUI interface
  - Chip balance – compute but also include GUI, multi-media

Platform Go To Market Strategies

- **WinTel follows an "Industry Standard"/"Low Hardware Cost" design point**
  - Requires high volume
  - Limits high value functions
- **RISC (Unix) follows a High Volume design point**
  - Generally leads to lower cost
  - Limits high value functions
- **System z follows a High Value design point**
  - Tries to deliver function not available elsewhere

Simplified Computer Architecture

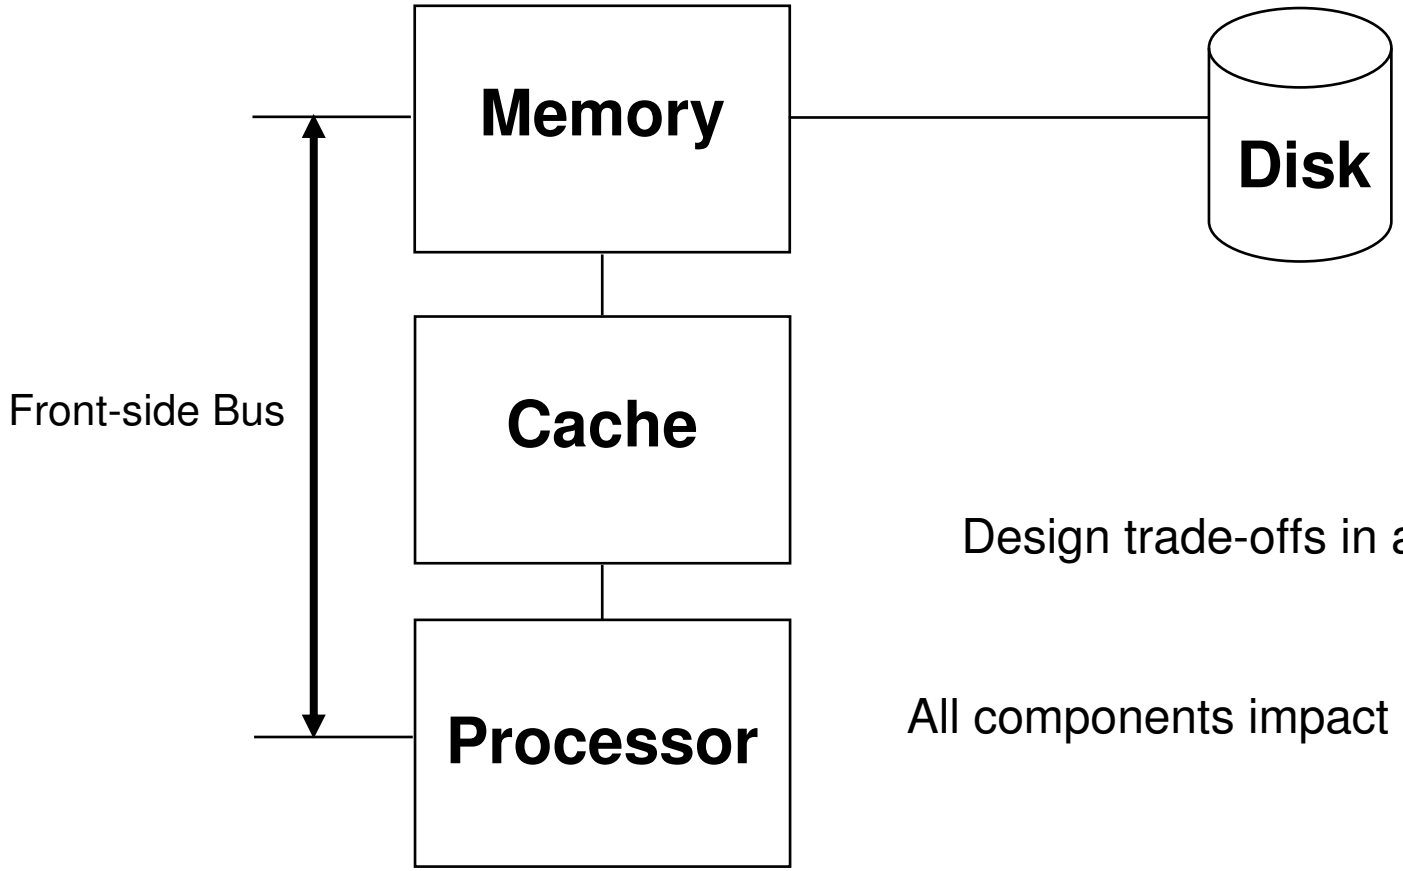

Design trade-offs in all aspects

All components impact performance

Design Philosophies determine speed…

$$ExecutionTime = CycleTime \times \frac{Cycles}{Instruction} \times PathLength$$

- **Intel designed to minimize cycle time**
  - Historically, rapid advances in clock speed
  - Recent change in strategy – Twin Castle initiative
  - Dual cores, quad cores
- **RISC designed to minimize Cycles Per Instruction**
  - RISC
  - Superscalar
  - SMT
- **System z designed to minimize path length**
  - Avoid Cache misses and reloads
  - Handle Page Faults in parallel
  - Add more…..

What determines system capacity and throughput?

There's more to performance than just processing power

- **Single system capacity is determined by:**
  - Processor speed
  - Memory hierarchy
  - I/O structure

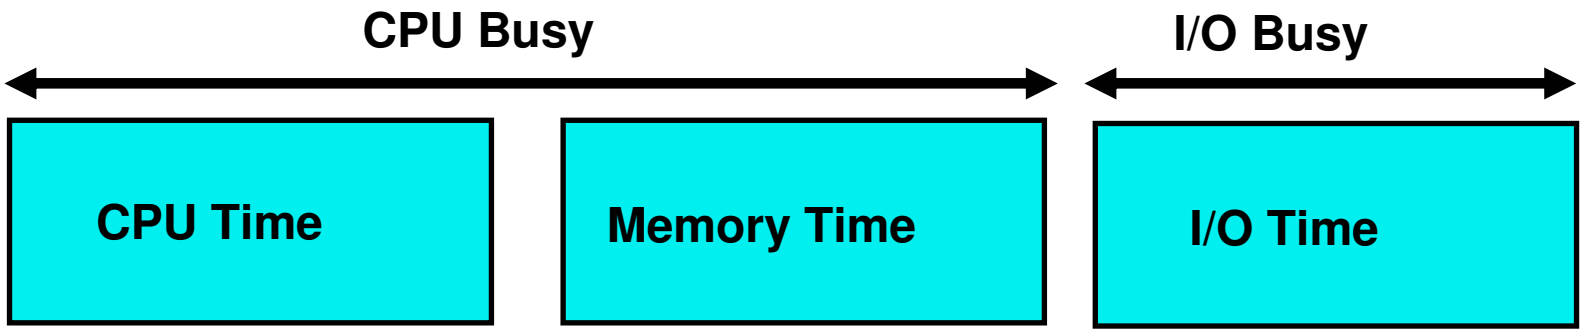

Processor, memory, and I/O times vary greatly by application and by machine type

Relative single system capacity

There's more to performance than just processing power

Workloads that do not require "balanced" computing, and rely solely on processor power will most likely perform better on Intel/RISC architectures

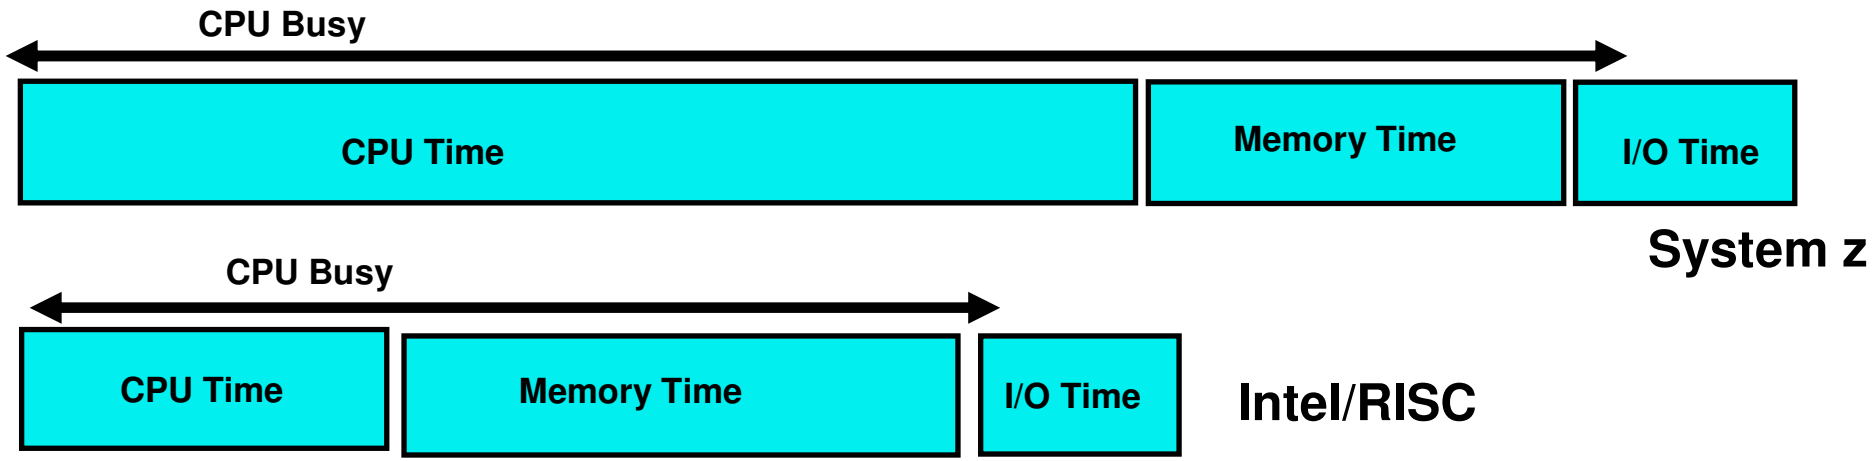

Processor- intensive workloads like SPECint, Deep Computing, and Graphic Rendering perform poorly on zSeries servers.

Relative single system capacity

There's more to performance than just processing power

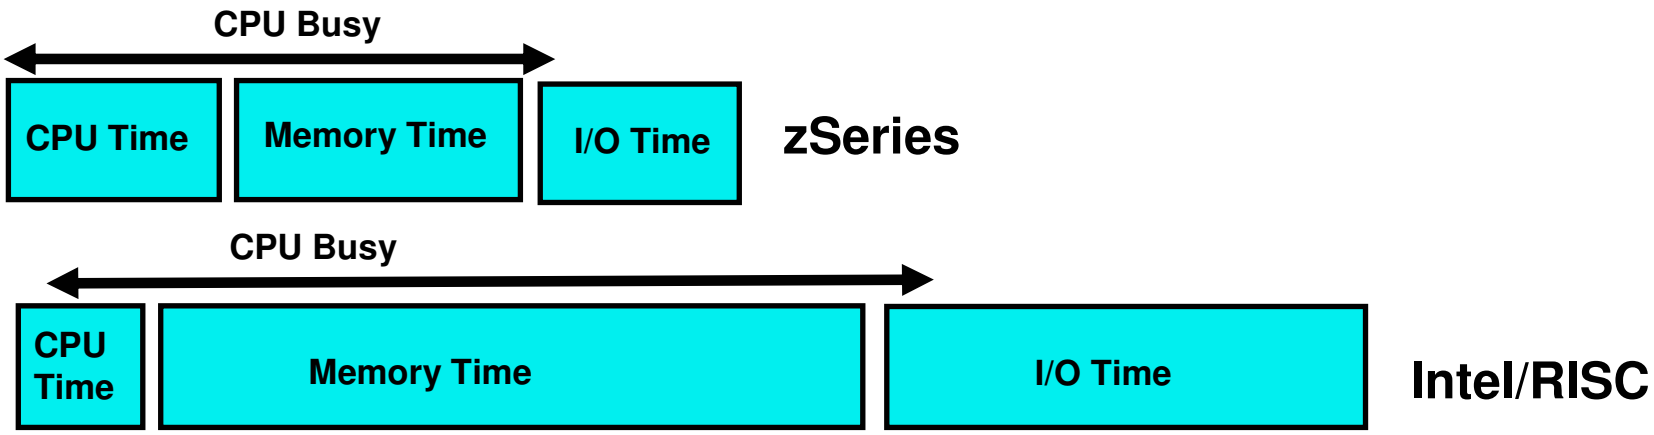

Data intensive workloads like large databases, transaction processing, object-oriented code and context switching typically run better on zSeries servers.

Comparison of System z and Intel/Risc

- **Many different ways to compare**
  - Design philosophy
  - Chip and Cache structures
  - Capacity and utilization
  - Dispatchers, Workload Managers, Mixed Workload effects
  - Partitioning
  - I/O
  - RAS features

Other Design Differences…too much to cover here

- **I/O Design**
  - Windows/Unix-Managed –vs- I/O Subsystem
- **Operating System Design**
  - Windows/Unix on top of hardware –vs- Integrated Design
- **Security**
  - Snap-on –vs- Integrated Design
- **RAS (Reliability, Availability, Security)**
- **Linguistics**
  - Work Load Manager (WLM), Virtualization (snap-on –vs- everything virtualized), LPAR (Logical Partitioning) machine level –vs-extreme granularity, etc., etc., etc.

Comparison of 32-way Machines

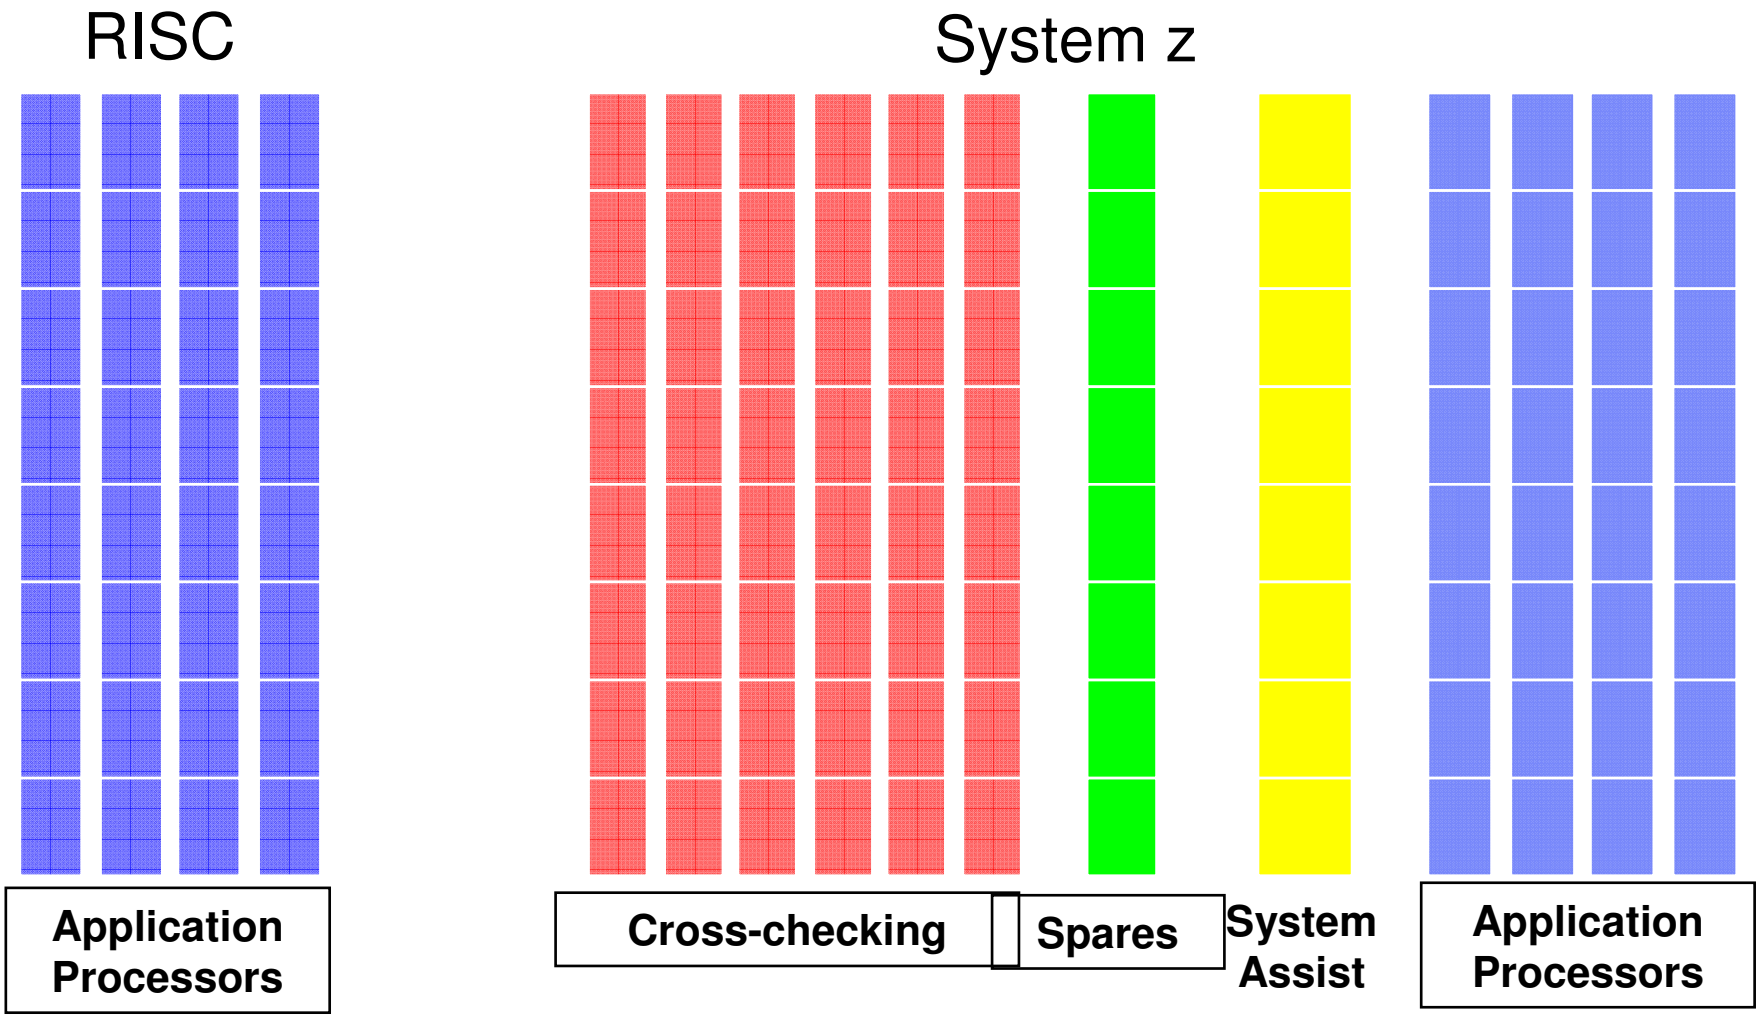

Summary of Positioning (Overly Simplified)

- **z/OS**
  - Highest Qualities of Service
  - Highest integration among applications
    - Critical for “chatty” applications
  - Highest utilization
- **zLinux -- Server farm in a box**
  - Best Virtualization
  - Significant cost savings in appropriate situations
  - Proximity to z/OS data
- **RISC**
  - Highest performance
  - Best of Breed Unix in QoS
  - Best in “Good Enough” computing
  - Very good Virtualization
- **Intel**
  - Fastest CPUs
  - Lowest hardware cost
  - Largest selection of software
  - Requires the most headroom (choose to run at low utilization)

Use the Right Technology for the Right Job

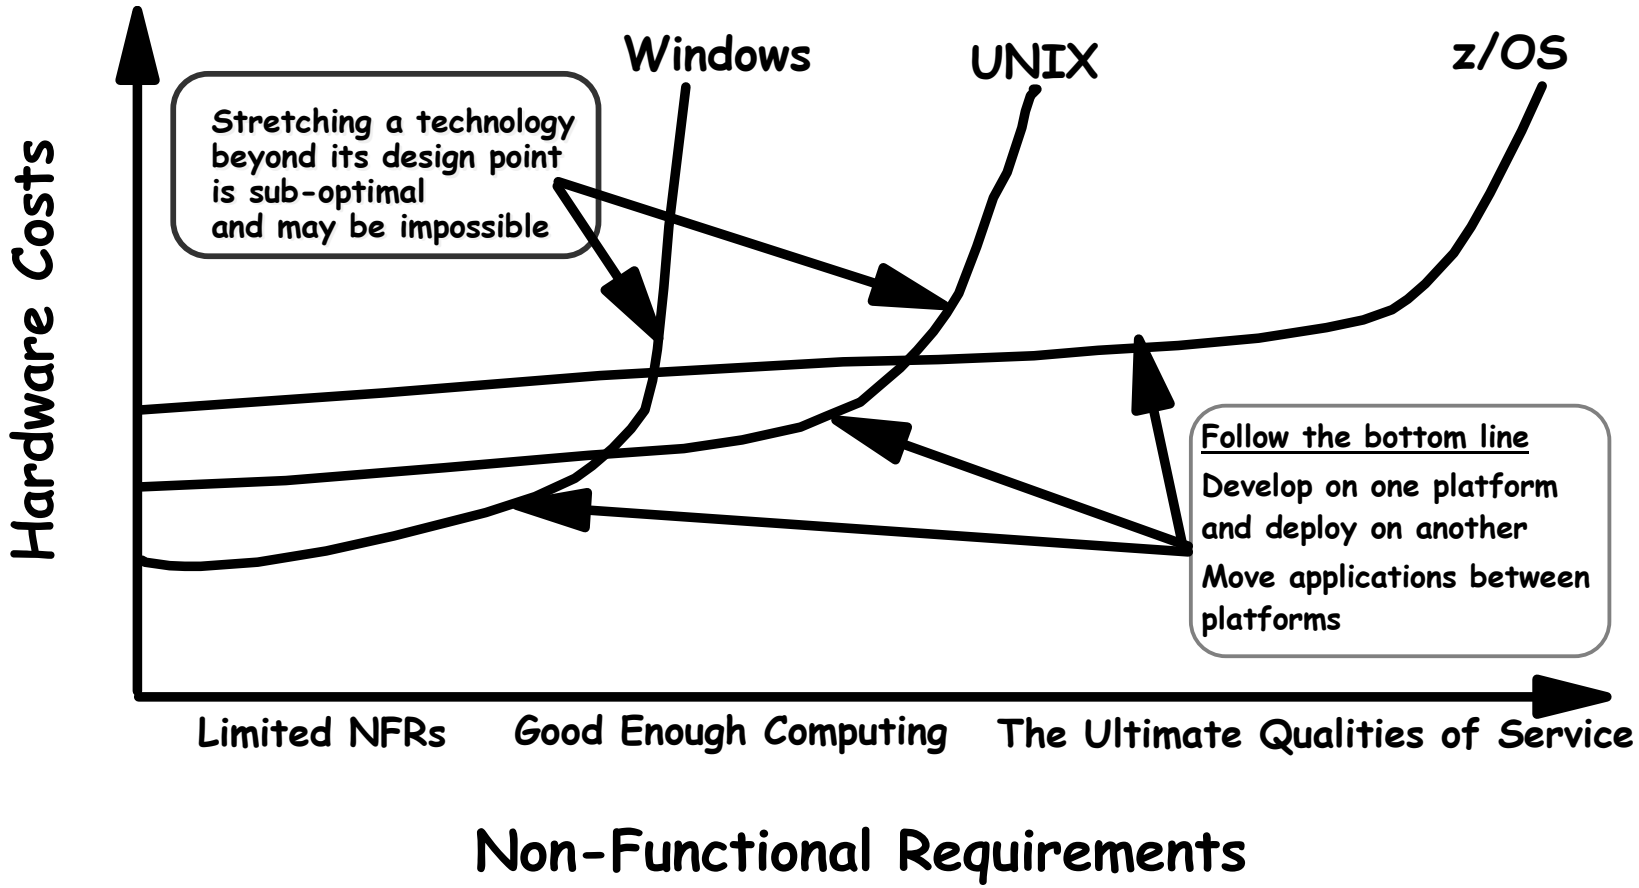

Total Cost of Ownership

TCO is more than just the price of hardware and software

Main Components of TCO

- Hardware
- Software
- People
- **Other (environmentals…)**

… and the profile is changing

People expense is now the dominant component

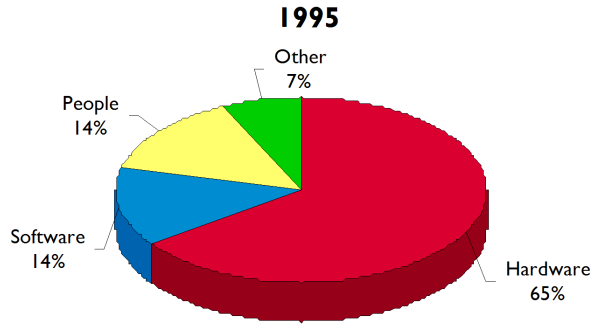

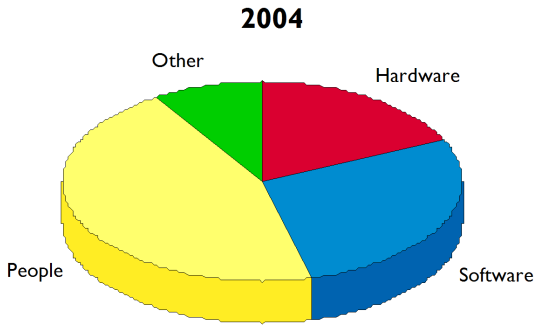

A full range of TCO factors considerations – *often ignored*

- ***Availability***
  - High availability
  - Hours of operation
- ***Backup / Restore / Site Recovery***
  - Backup
  - Disaster Scenario
  - Restore
  - Effort for Complete Site Recovery
  - SAN effort
- ***Infrastructure Cost***
  - Space
  - Power
  - Network Infrastructure
  - Storage Infrastructure
- ***Additional development and implementation***
  - Investment for one platform – reproduction for others
- ***Controlling and Accounting***
  - Analyzing the systems
  - Cost
- ***Operations Effort***
  - Monitoring, Operating
  - Problem Determination
  - Server Management Tools
  - Integrated Server Management – Enterprise Wide
- ***Security***
  - Authentication / Authorization
  - User Administration
  - Data Security
  - Server and OS Security
  - RACF vs. other solutions
- ***Deployment and Support***
  - System Programming
    - Keeping consistent OS and SW Level
    - Database Effort
  - Middleware
    - SW Maintenance
    - SW Distribution (across firewall)
  - Application
    - Technology Upgrade
    - System Release change without interrupts
- ***Operating Concept***
  - Development of an operating procedure
  - Feasibility of the developed procedure
  - Automation
- ***Resource Utilization and Performance***
  - Mixed Workload / Batch
  - Resource Sharing
    - shared nothing vs. shared everything
  - Parallel Sysplex vs. Other Concepts
  - Response Time
  - Performance Management
  - Peak handling / scalability
- ***Integration***
  - Integrated Functionality vs. Functionality to be implemented (possibly with 3rd party tools)
  - Balanced System
  - Integration of / into Standards
- ***Further Availability Aspects***
  - Planned outages
  - Unplanned outages
  - Automated Take Over
  - Uninterrupted Take Over (especially for DB)
  - Workload Management across physical borders
  - Business continuity
  - Availability effects for other applications / projects
  - End User Service
  - End User Productivity
  - Virtualization
- ***Skills and Resources***
  - Personnel Education
  - Availability of Resources

Key Points – Mainframe Costs

- The cost of running incremental workload on the mainframe goes down as the total workload grows
  - ✔Labor costs hold steady as workload grows
  - ✔IBM pricing policies designed to favor the addition of more workload
  - ✔Lower software costs per transaction as workload grows
  - ✔Special hardware pricing for new workload types
  - ✔Lower electrical and air conditioning consumption than server farms
- Consolidation opportunities accelerate the benefit
- Customers have learned (are learning) that mainframes running high workloads are the most cost efficient platform

Arcati

Predicted average cost
per end user in 2010

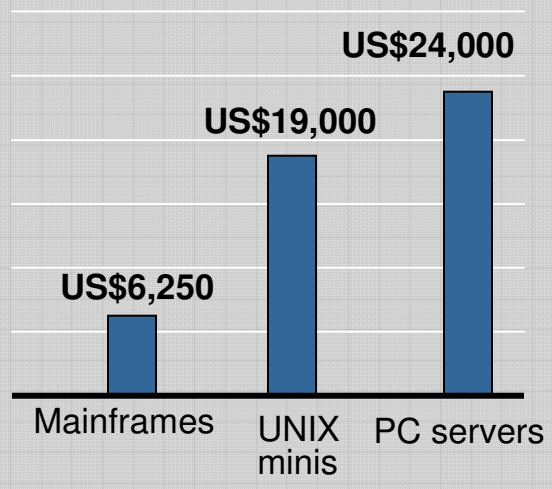

Five-year costs for hardware, software and maintenance

Arcati Research 2005.
"The Dinosaur Myth 2004 Update."

Owing to the nature of individual contracts, some details of this pricing discussion may be at variance with specific instances

Taking only TCA (Total Cost of Acquisition) into consideration (Financial institution running WebSphere)

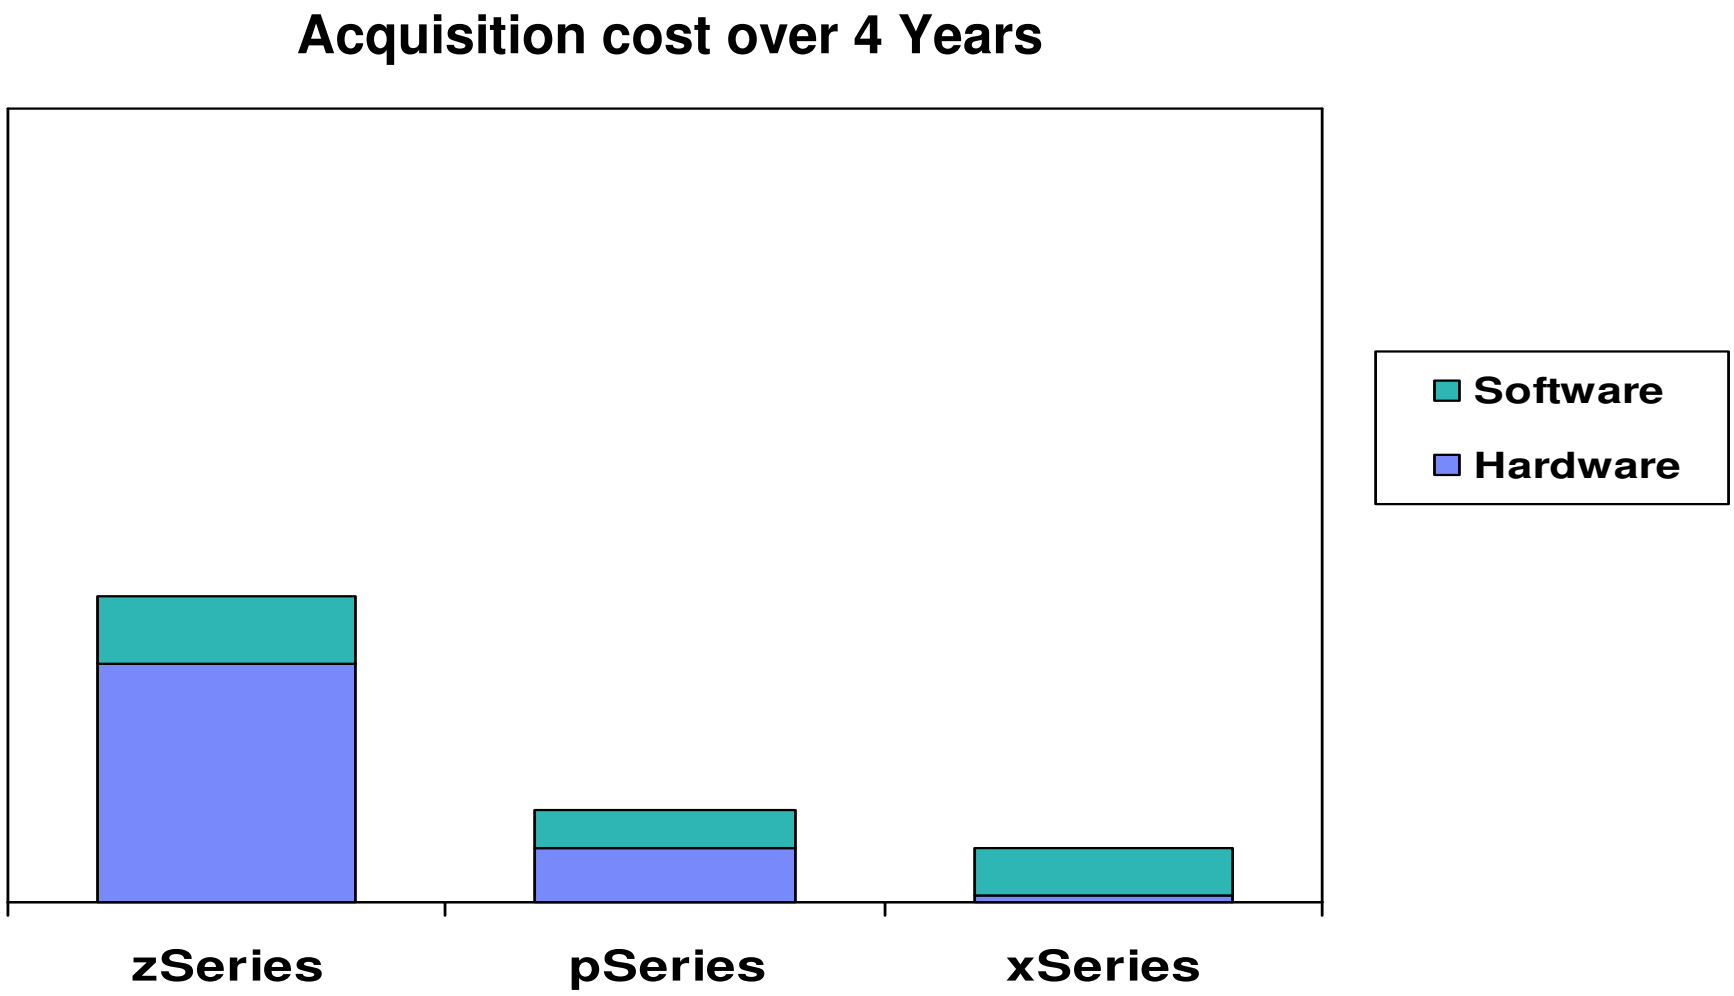

German Financial Institution TCO study

*Distributed servers have higher service, monitoring and support costs **and** cost more to develop and implement*

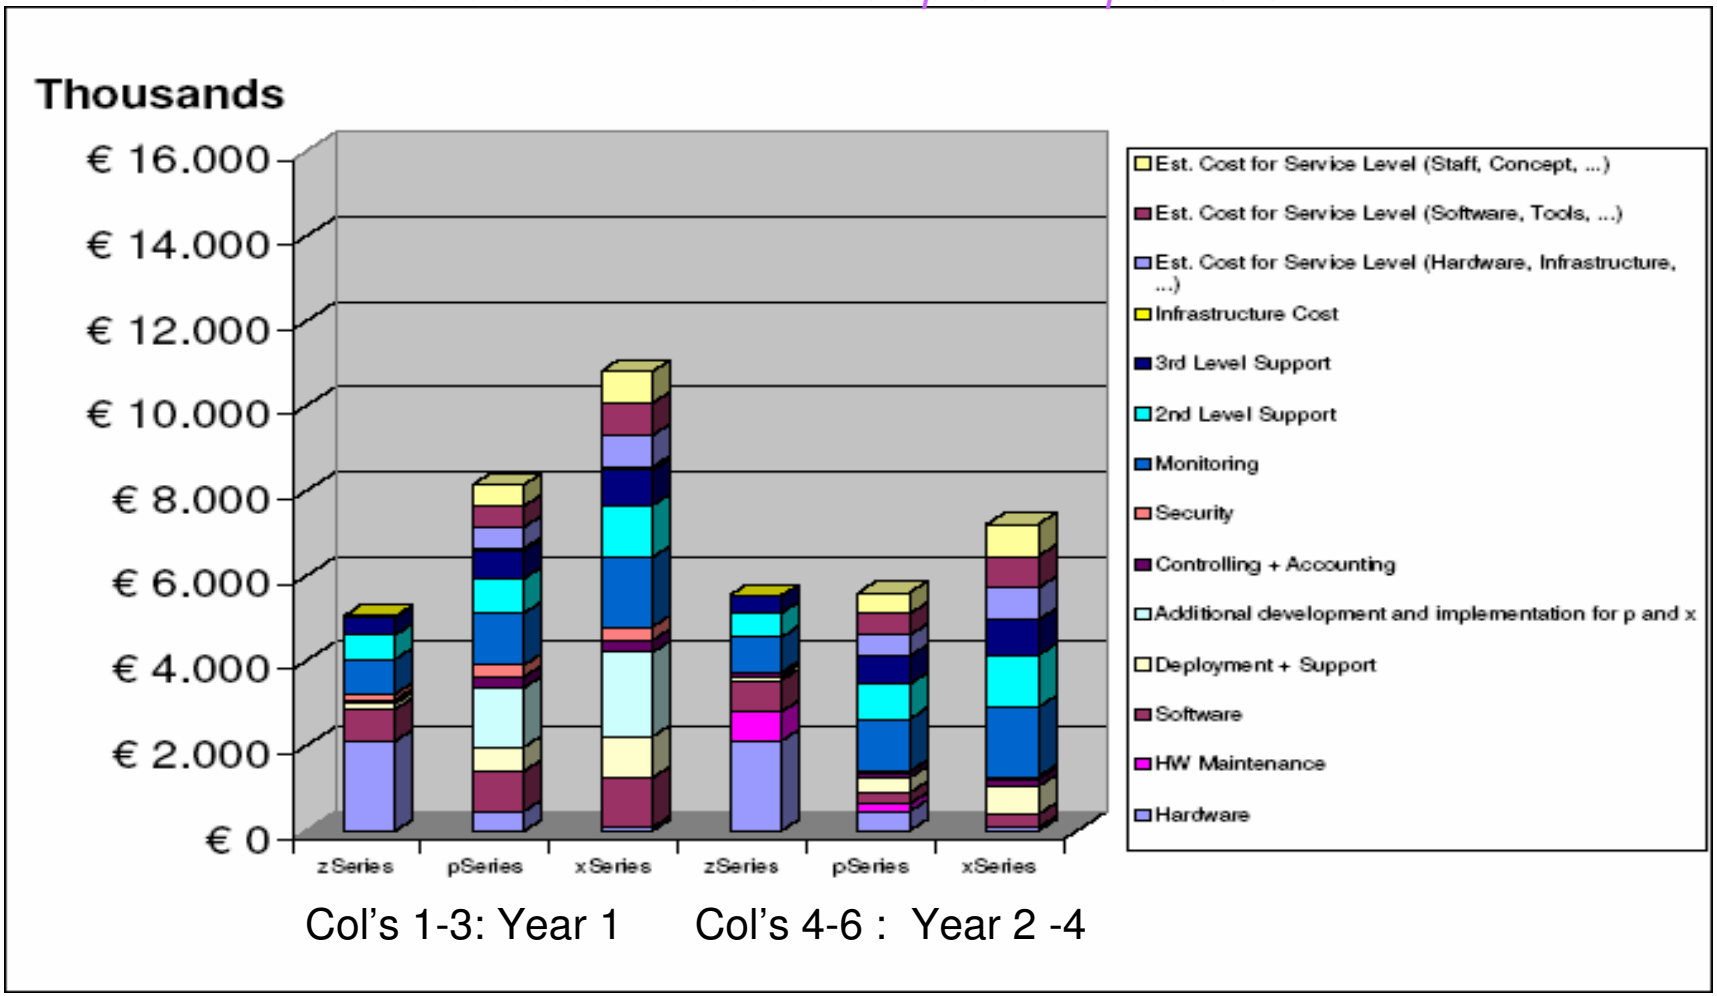

Col's 1-3: Year 1 Col's 4-6 : Year 2 -4

Customer Example: Distributed SUN Server Solution – perception… isn't always reality!

- **US Finance customer thought they only had 24 UNIX servers**
  - But these were just the PRODUCTION servers
  - In addition they had 49 servers for Development, Test and Disaster Recovery

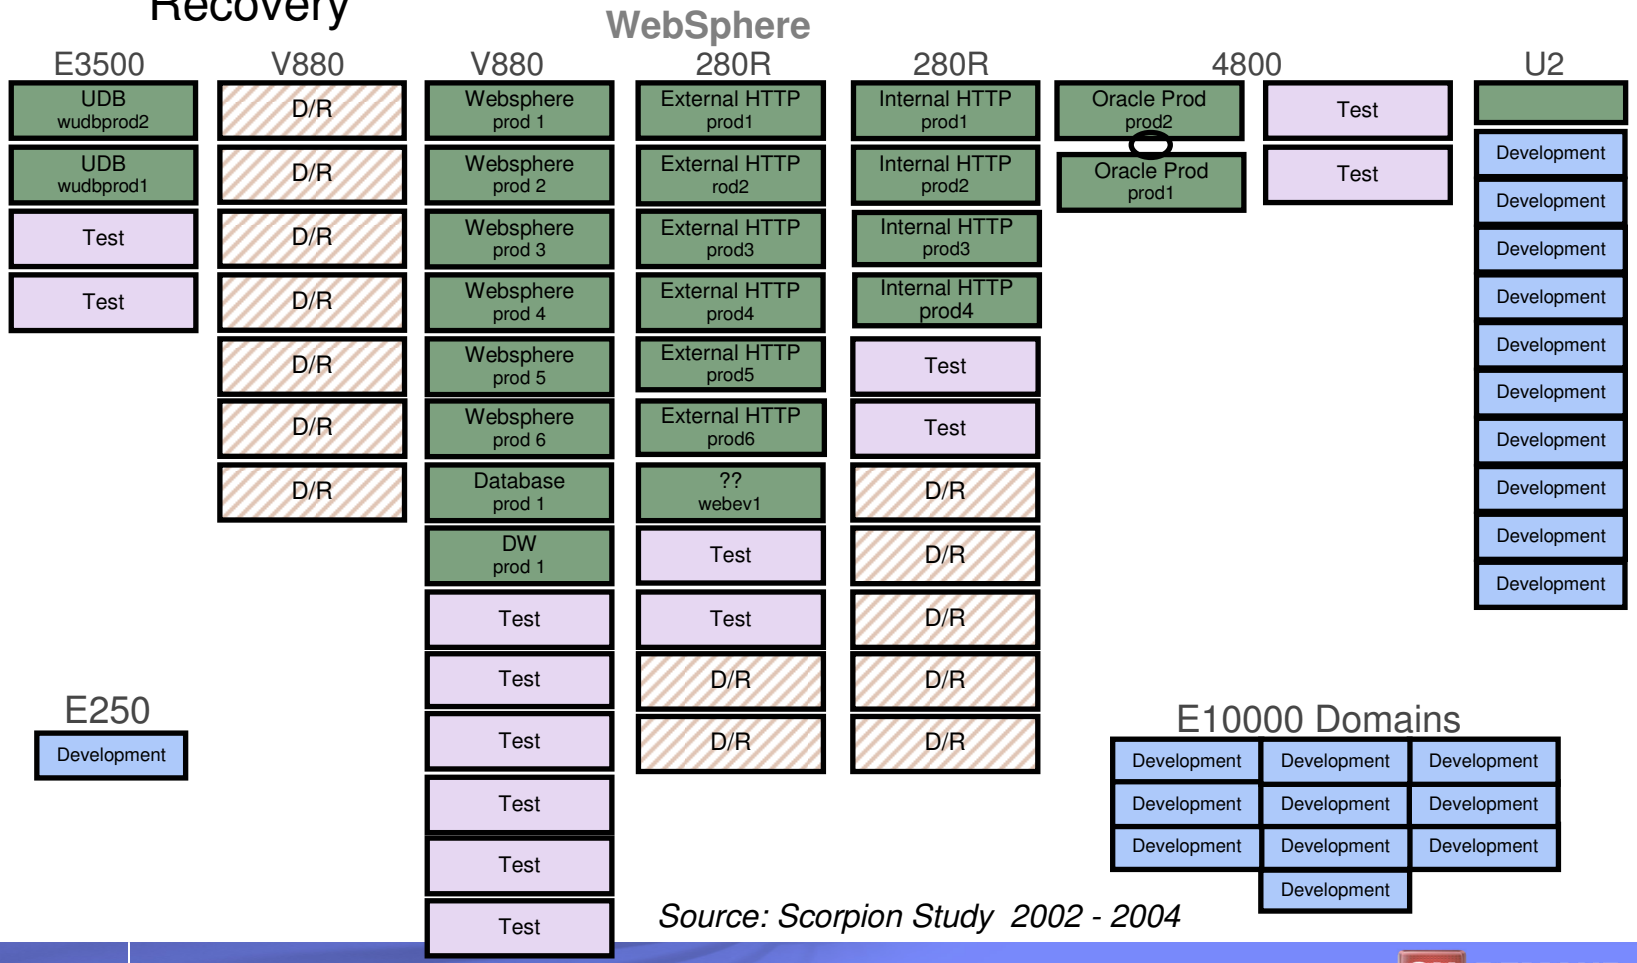

Source: Scorpion Study 2002 - 2004

People costs are often *hidden* in distributed implementations

- **In a recent typical study, a customer thought they only had 24 UNIX servers**
  - But these were just the PRODUCTION servers
  - In addition they had 49 servers for Development, Test and Disaster Recovery
- **They needed 14 people to support these servers and $7M software**
  - Running at only 20% utilization
  - Each server cost $20K per annum to support
- **A comparable zSeries implementation would have required just 20 servers**
  - Requiring 5 people to support
  - Using $6M software (over 3 years)

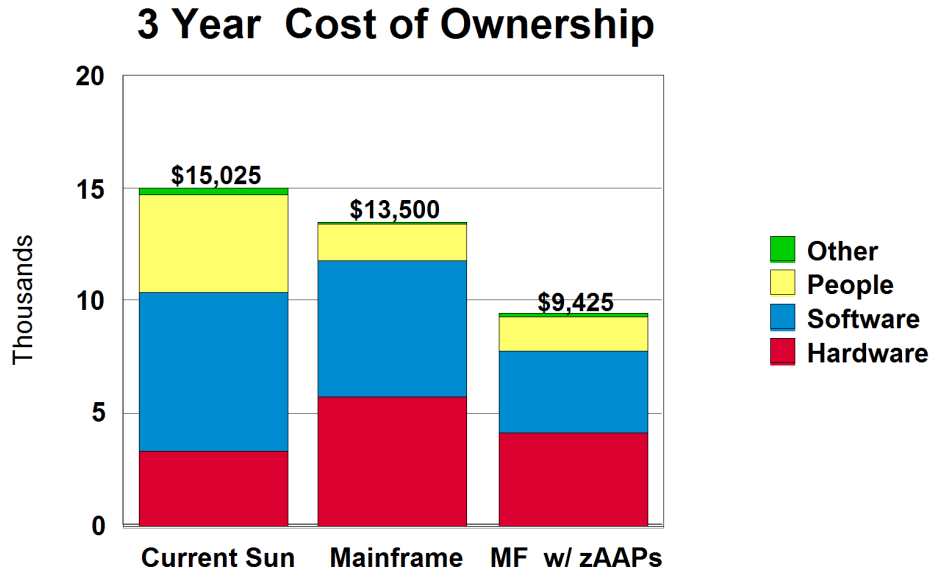

- They thought the Solaris environment was 1/5 the cost of the mainframe…
  …but in fact the **zSeries TCO was 37% less**

Conclusions…..

- **Each technology has its place**
  - They are optimized for different workloads and environments
- **There is no single "right answer"**
  - The same application at different customers may have different answers
- **We need to help customers identify what is important to them**
  - When we do, the right answer will usually be evident


Thank
You