Fit for Purpose(a workload based view)

Joe Temple – jliitemp@us.ibm.com
Session 1063

SHARE in Boston

Agenda

- Key Issues
  - Infrastructures that are Connected but not integrated, dynamic or aligned
  - Cost Trends
  - Technology constraint
  - Workload Diversity
- STG Strategy
  - Dynamic Infrastructure
  - Fit for Purpose
  - Workload Optimization
  - A structural model for “Smart IT”
- Fit for Purpose highlights and key ideas
- Workload Optimization and platform positioning

Islands of Computing

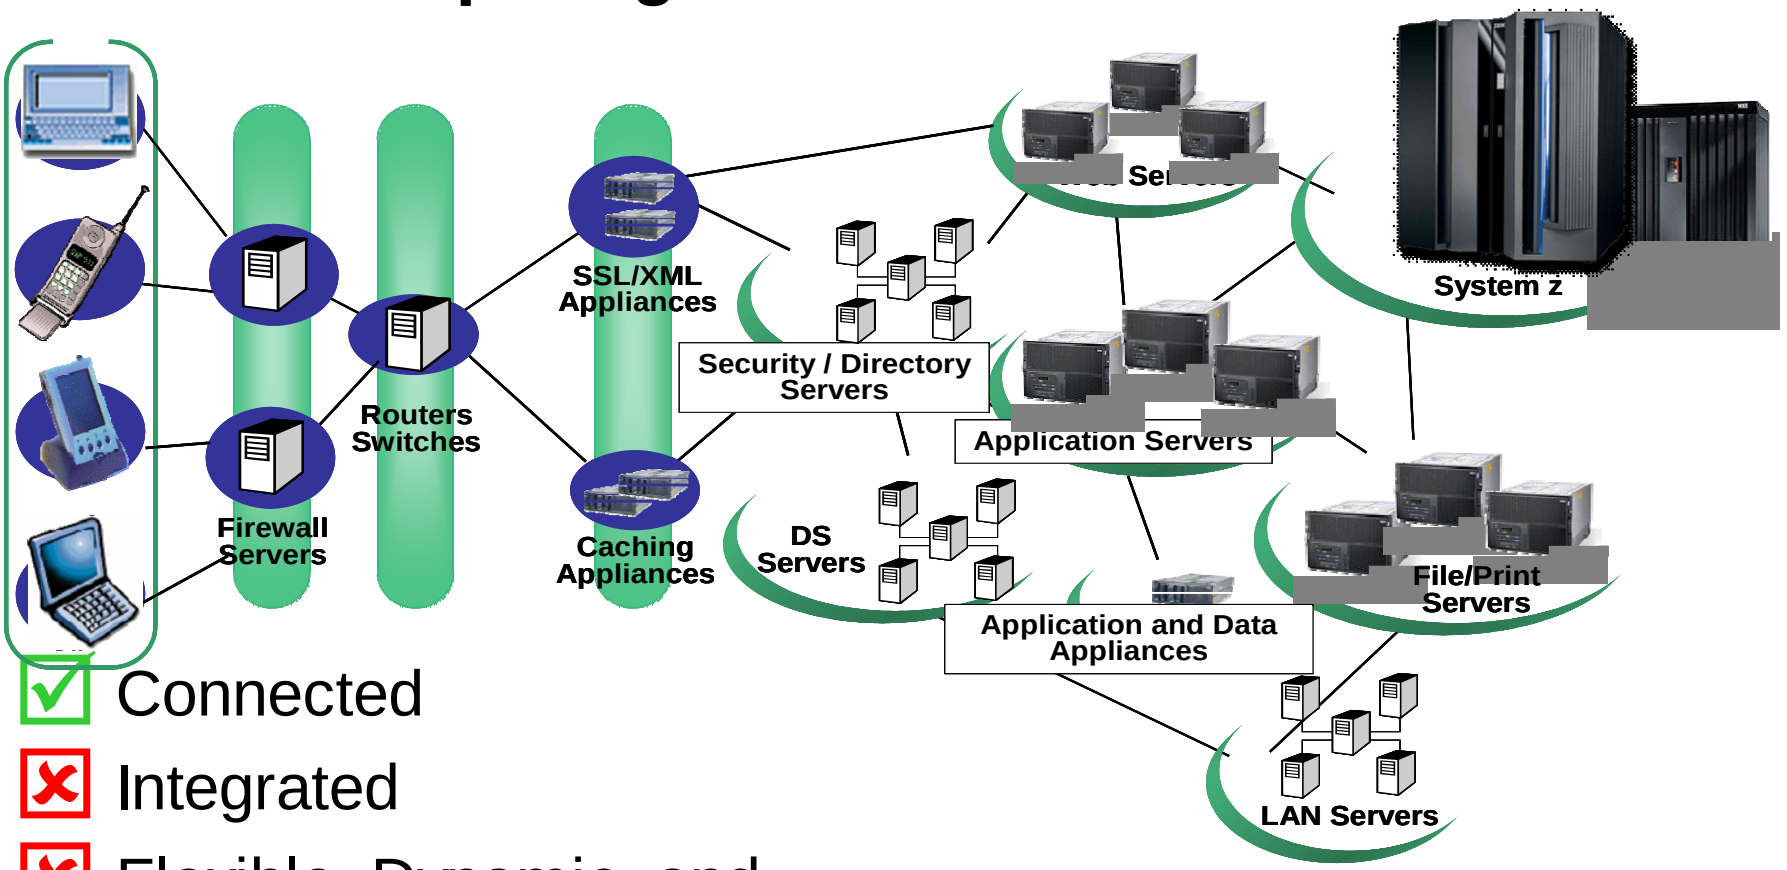

- ☑ Connected
- ☒ Integrated
- ☒ Flexible, Dynamic, and Responsive
- ☒ Aligned with Business Objectives

IDC – Shifting Management Requirements

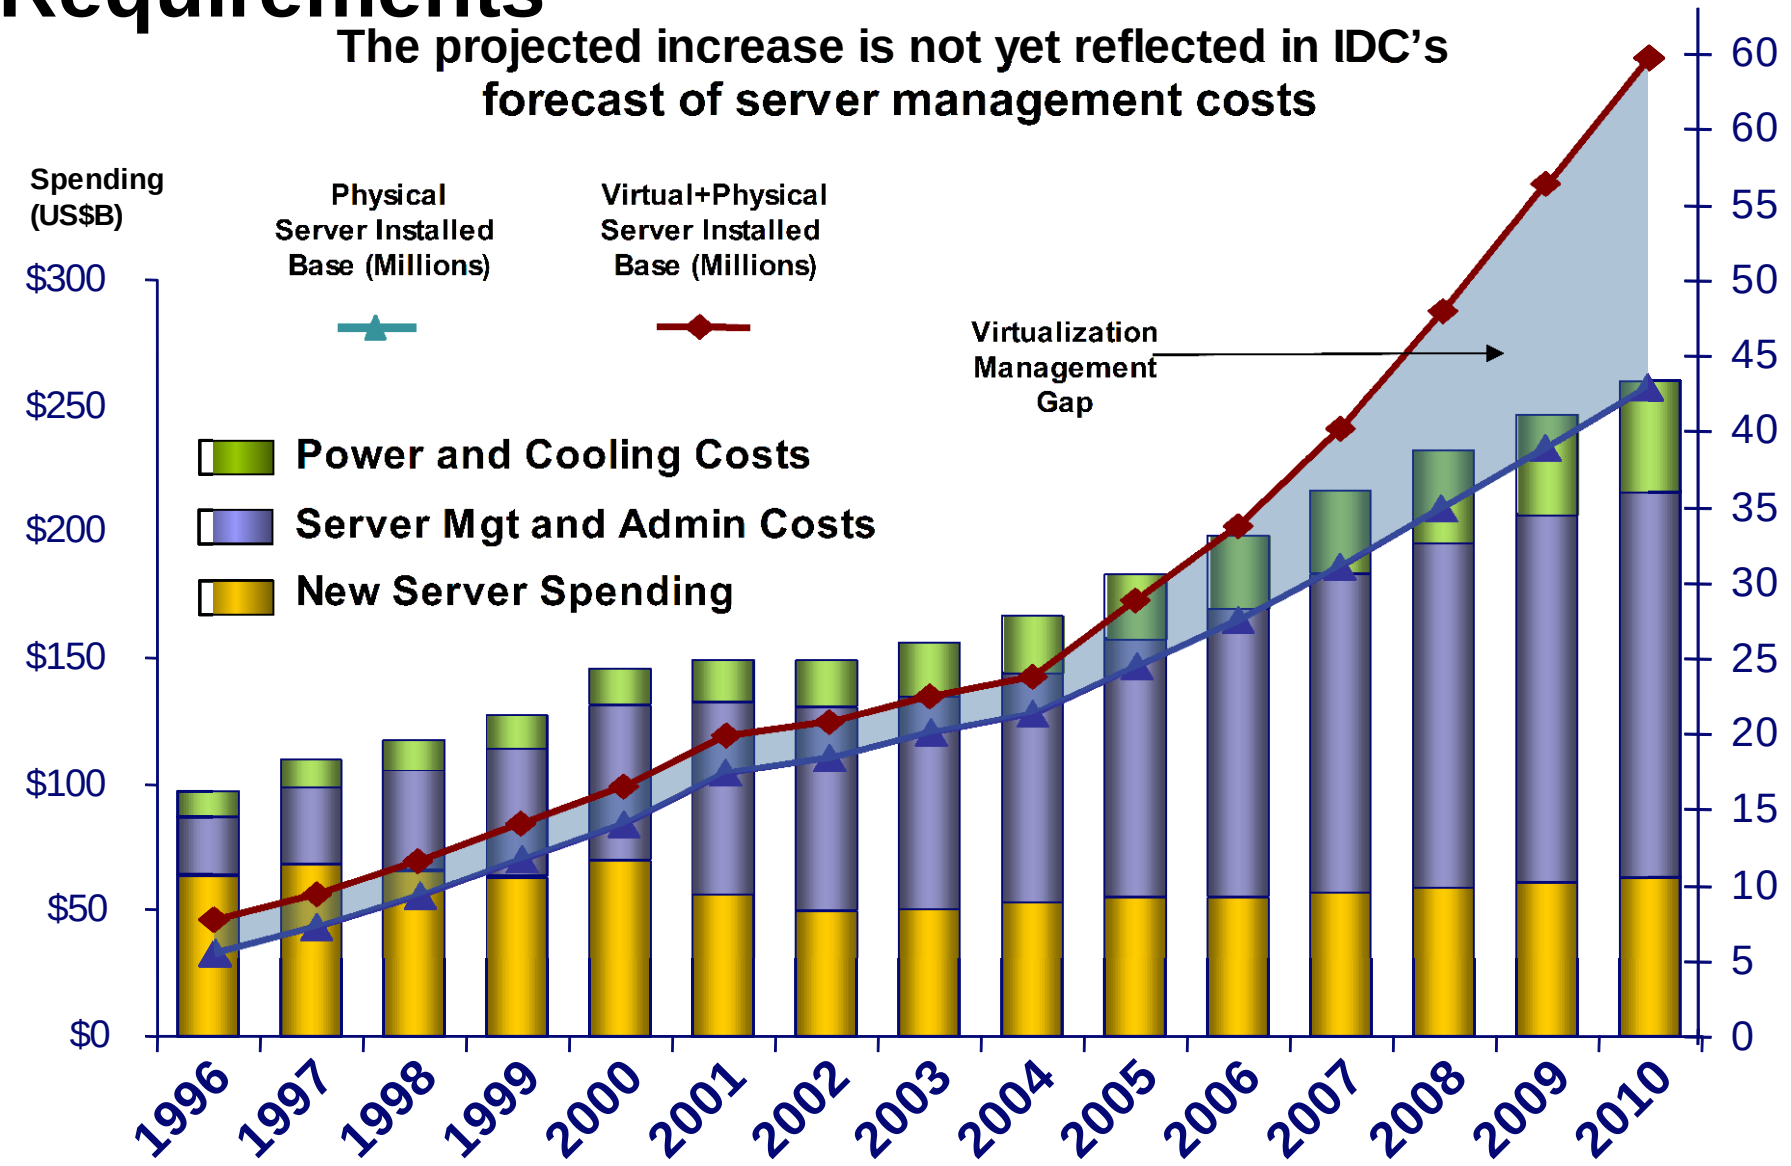

Source: IDC, May 2006

IDC Workload Segments

Transaction Processing and Database

Analytics and High Performance

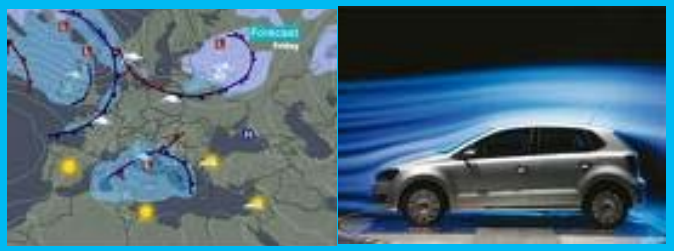

Characteristics vary: ***Application Usage Patterns SLA Data Structure Integration***

Business Applications

Web, Collaboration and Infrastructure

One size DOES NOT fit all

Next: Coping with physical limits

- **The industry is hitting fundamental physical limits:**
  - Size
  - Speed of electromagnetic propagation
  - Heat transfer rates
- **Large CPU speed increases are a thing of the past, across the industry**
- **Capacity increases will increasingly come from higher n-way, more multithreading, and NUMA optimization**
- **Demand for lower latency will drive co-location of hybrid transaction processing elements**

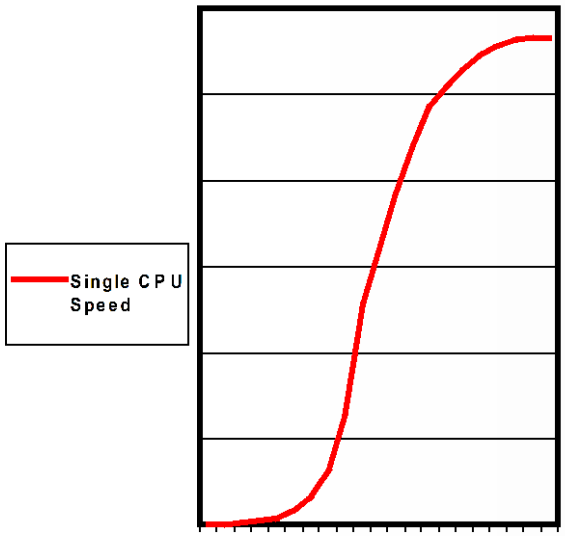

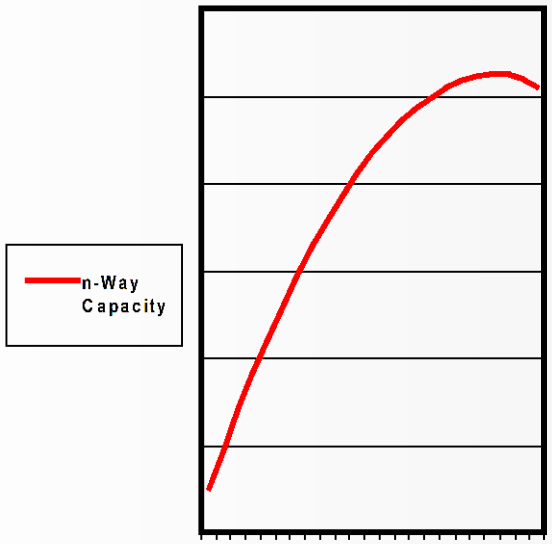

"In terms of size [of transistor] you can see that we're approaching the size of atoms which is a fundamental barrier, …."

Gordon Moore, April 2005*

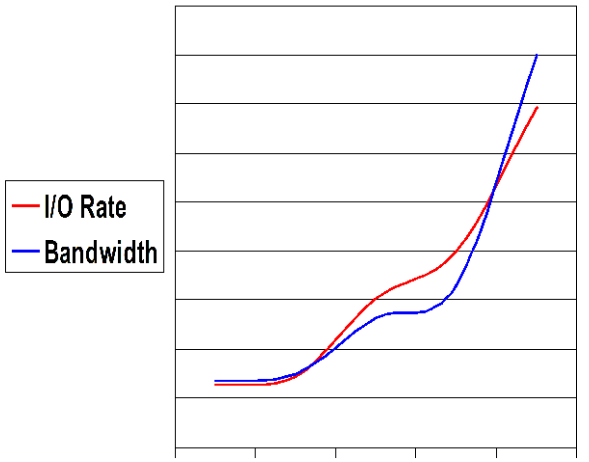

* Techworld, Operating Systems and Servers News, 13 April 2005

Dynamic Infrastructure, Fit for Purpose, Workload Optimization

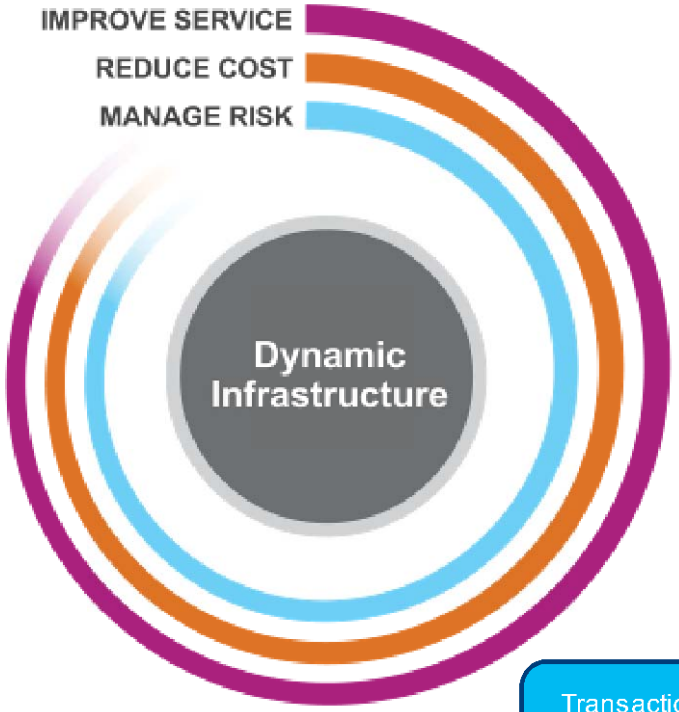

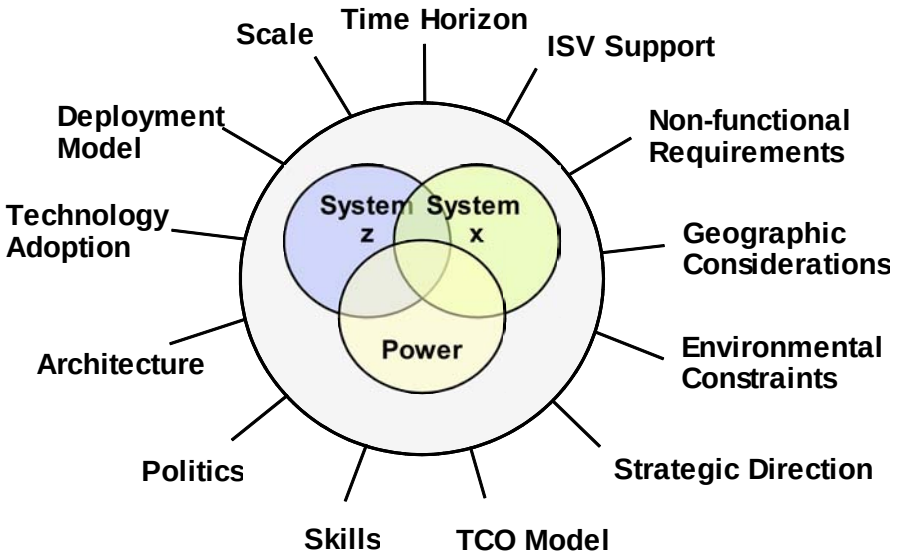

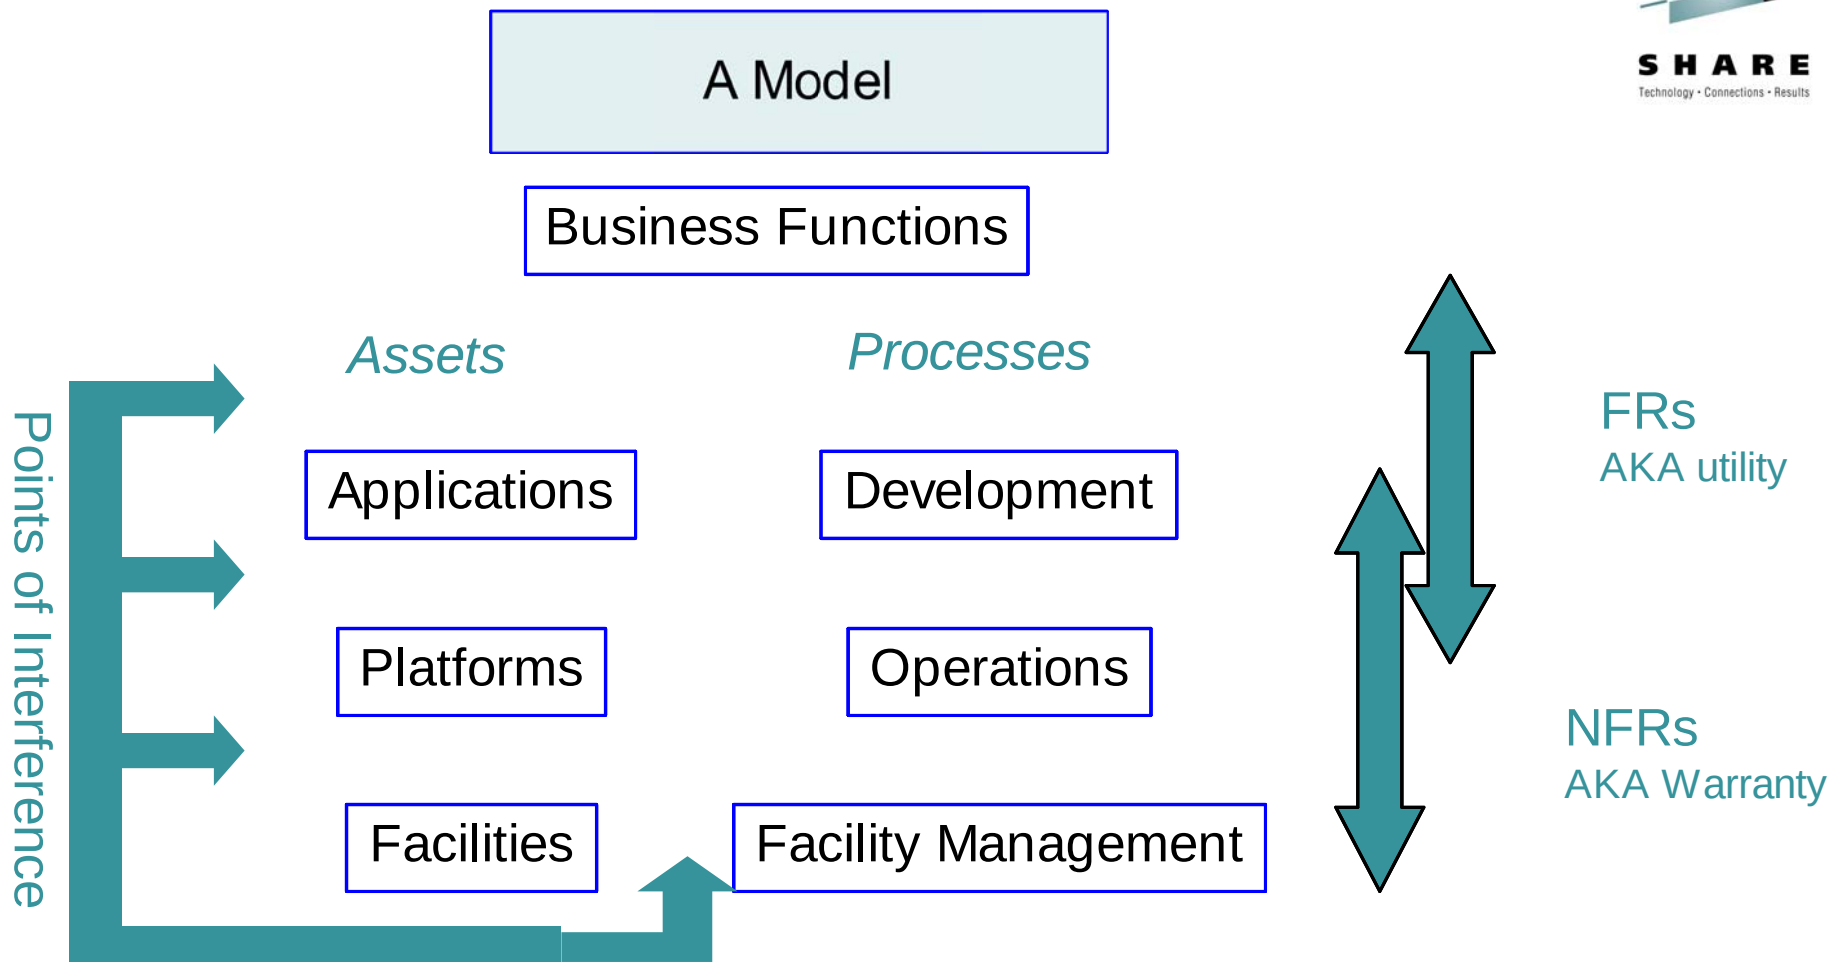

Architecture creates abstractions at the layer boundaries.

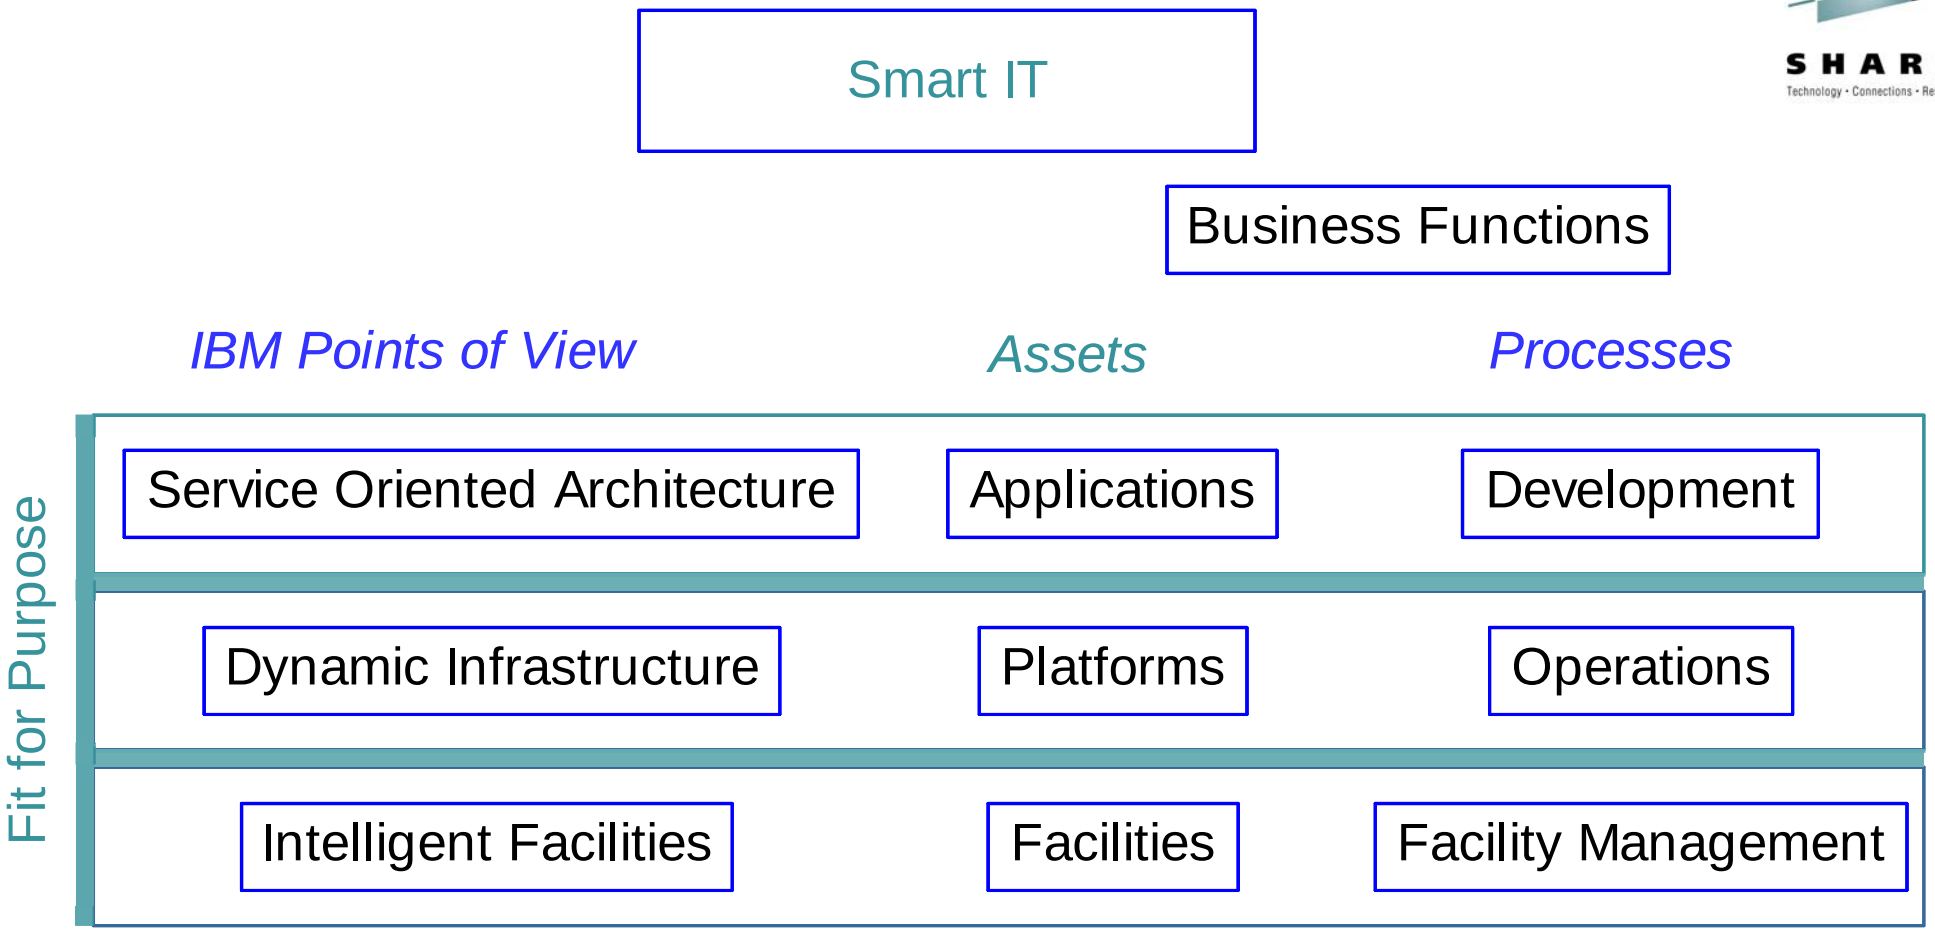

Fit for purpose is the glue that deals with interference at the boundaries
Applies to all IT whether you accept IBM's points of view or not

Fit for Purpose

Server Selection and Positioning

SHARE in Boston

Many Factors Affect Choice

Would you purchase a family car solely on one factor?

Car	Server
Purchase price	Purchase price
Gas mileage, cost of repairs, insurance cost	Cost of operation, power consumption, floor space
Reliability	Reliability
Safety, maneuverability, visibility, vendor service	Availability, disaster recovery, vendor service
Storage capacity, number of seats, towing capacity	Scalability, throughput
Horsepower	Chip performance
Dash board layout Steering wheel location	Skills
Handling, comfort, features	Manageability
Looks, styling, size	Peer and industry recognition

Fit for purpose Highlights:

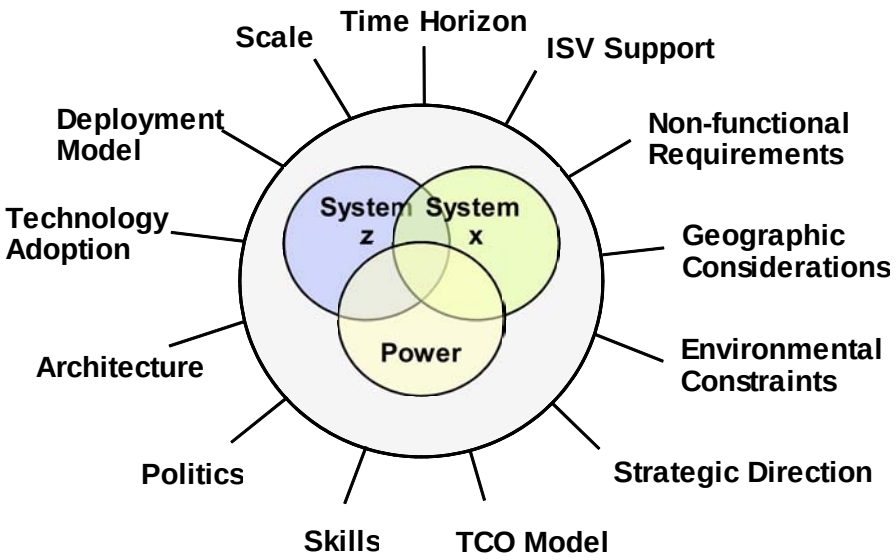

Local Factors Matter

- Skills
- Technology adoption
- Management
- Volume of servers
- Organizational

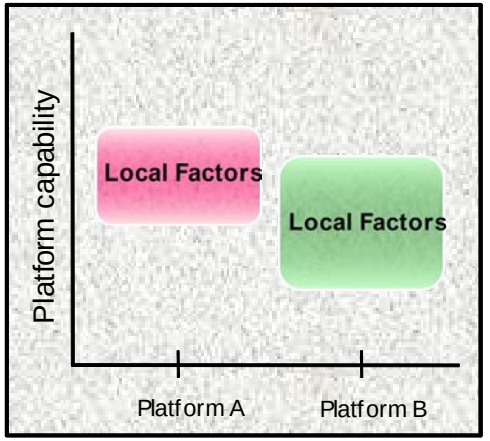

Infrastructure Size Matters

- Changes people dynamics
- Increases handoffs
- Affects testing, patching, etc

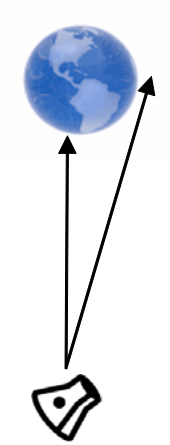

Apollo 13

Workload Matters

Fit for purpose Highlights:

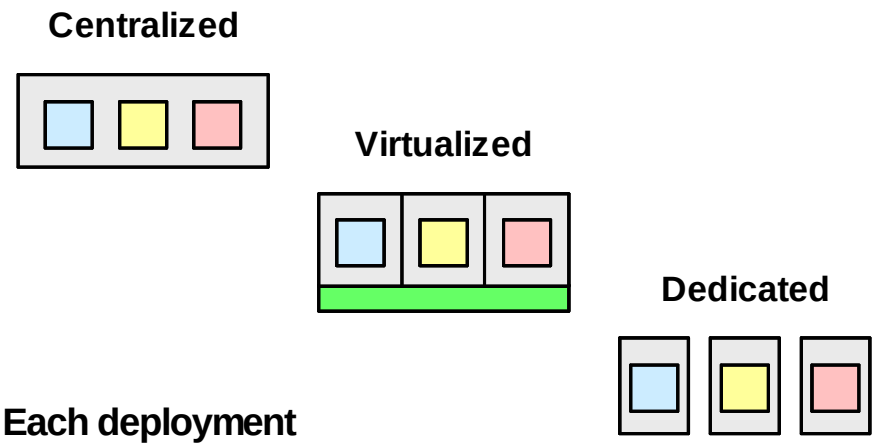

Each deployment model has its place

Non-Functional
"How well it does it"

- Availability requirements
- Transactions per minute
- Security requirements
- Ease of provisioning and support
- Disaster recovery requirements
- Future growth

Non-functional requirements are the significant element of platform selection

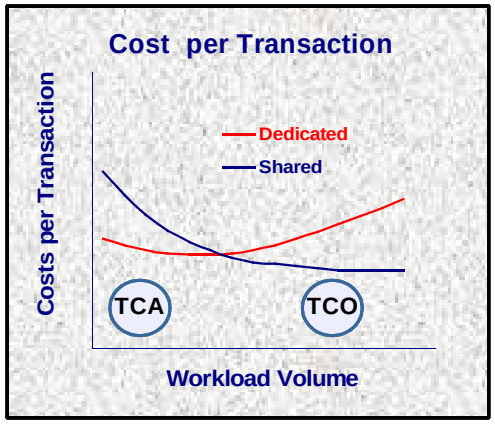

Beware of the hidden cost of sub optimization

Chargeback models often distort the selection process

Key ideas:

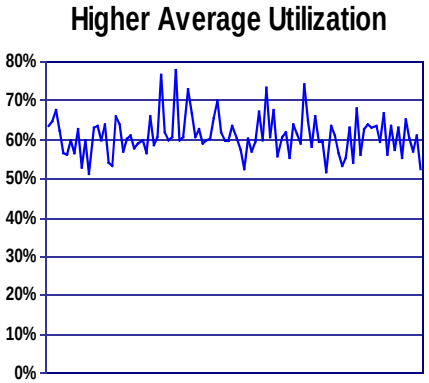

Increased Productivity

More Efficient Data Center

Better RAS

Large servers offer advantages

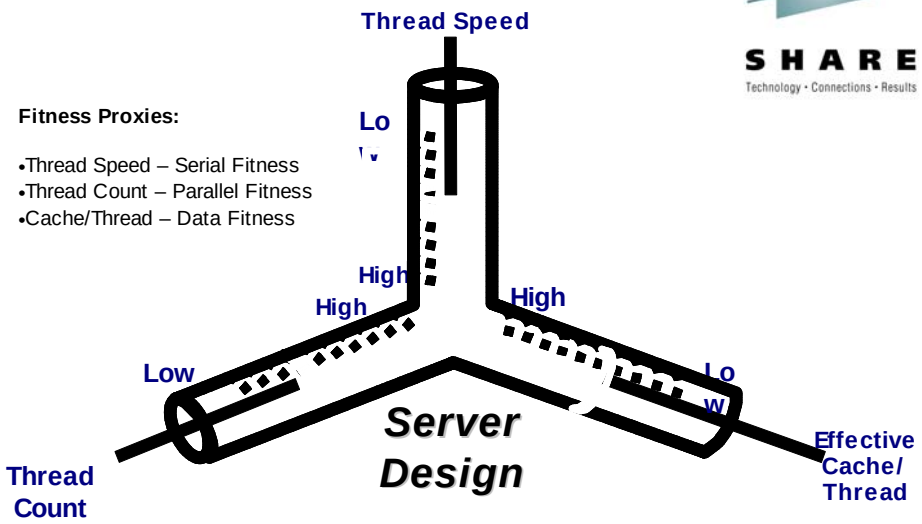

Tradeoffs mean that one size does not fit all

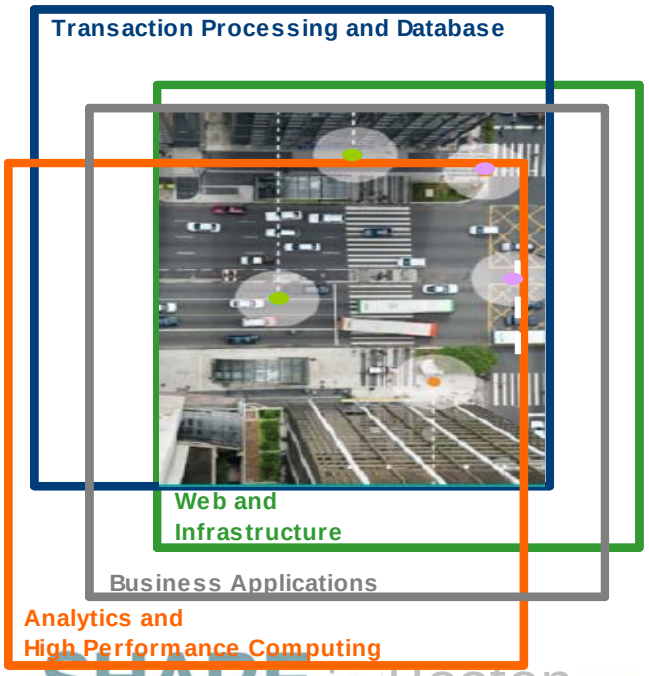

Smarter Planet solutions are hybrids

Classic BI

- Analyze and report on historical data
- Move data into a data warehouse
- Strategic, long-term analysis
- Processing scheduled, typically timed to meet hreporting deadlines

Emerging requirements

- Predict the future
- BI data co-located with transactional data
- Results drive immediate, sometimes automated action
- Analyze and act in real time

Workload requirements shift with time

Workload Optimization

Joe Temple
jliitemp@us.ibm.com

SHARE in Boston

Workload Requirements Shift over Time

Classic BI

- Analyze and report on historical data
- Move data into a data warehouse
- Strategic, long-term analysis
- Processing scheduled, typically timed to meet reporting deadlines

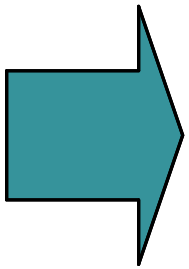

Emerging requirements

- Predict the future
- *BI data co-located with transactional data*
- *Results drive immediate, sometimes automated action*
- Analyze and act in real time

Data analytics require more parallel fitness from System z, or faster Extract, Transform and Load (ETL) from classic BI solutions

In some cases analytics get embedded in transactions.
This increases path length and cache footprint and/or spawns competing threads, creating a “mixed workload”.

Pfister's Paradigm and "Temple's Assertion"

From: In Search of Clusters, The ongoing battle in lowly parallel computing, Greg Pfister, p461

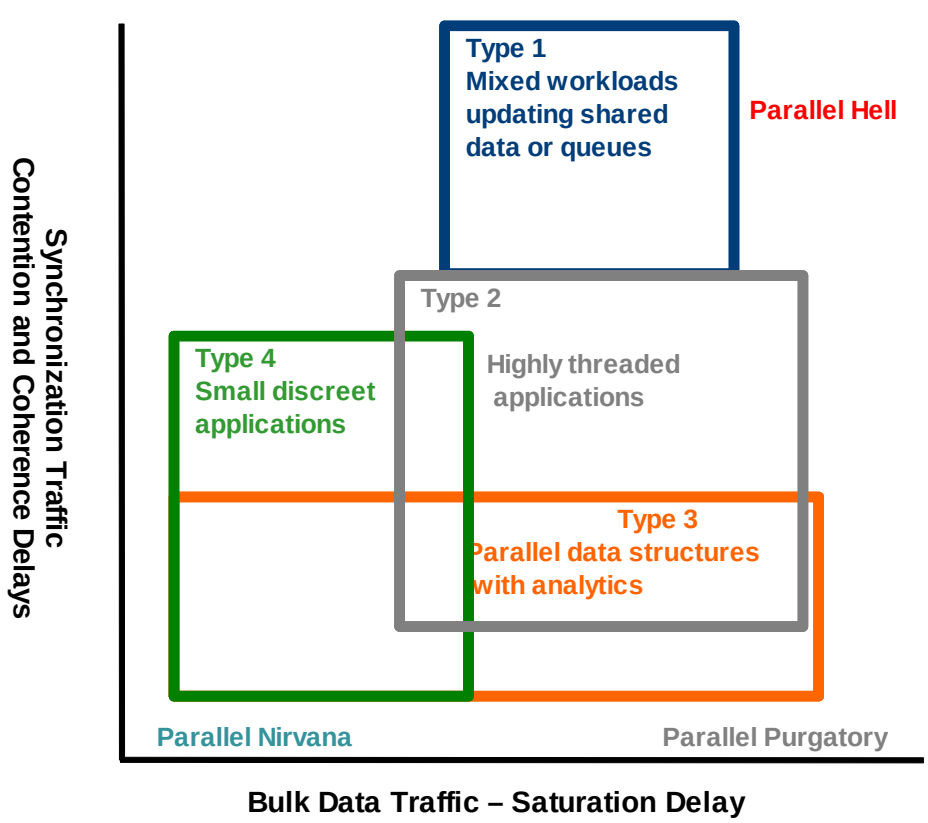

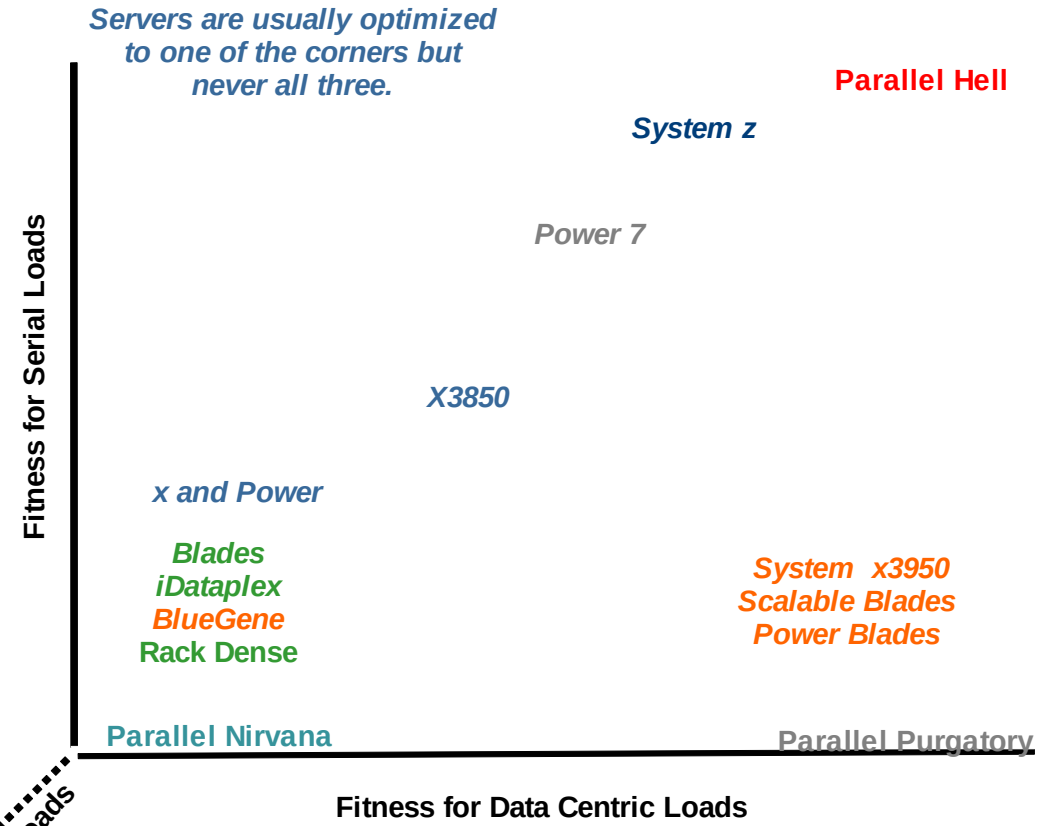

Matching workload types to machines

Trade offs in the fitness space

Fitness Proxies:

- Thread Speed – Serial Fitness
- Thread Count – Parallel Fitness
- Cache/Thread – Data Fitness

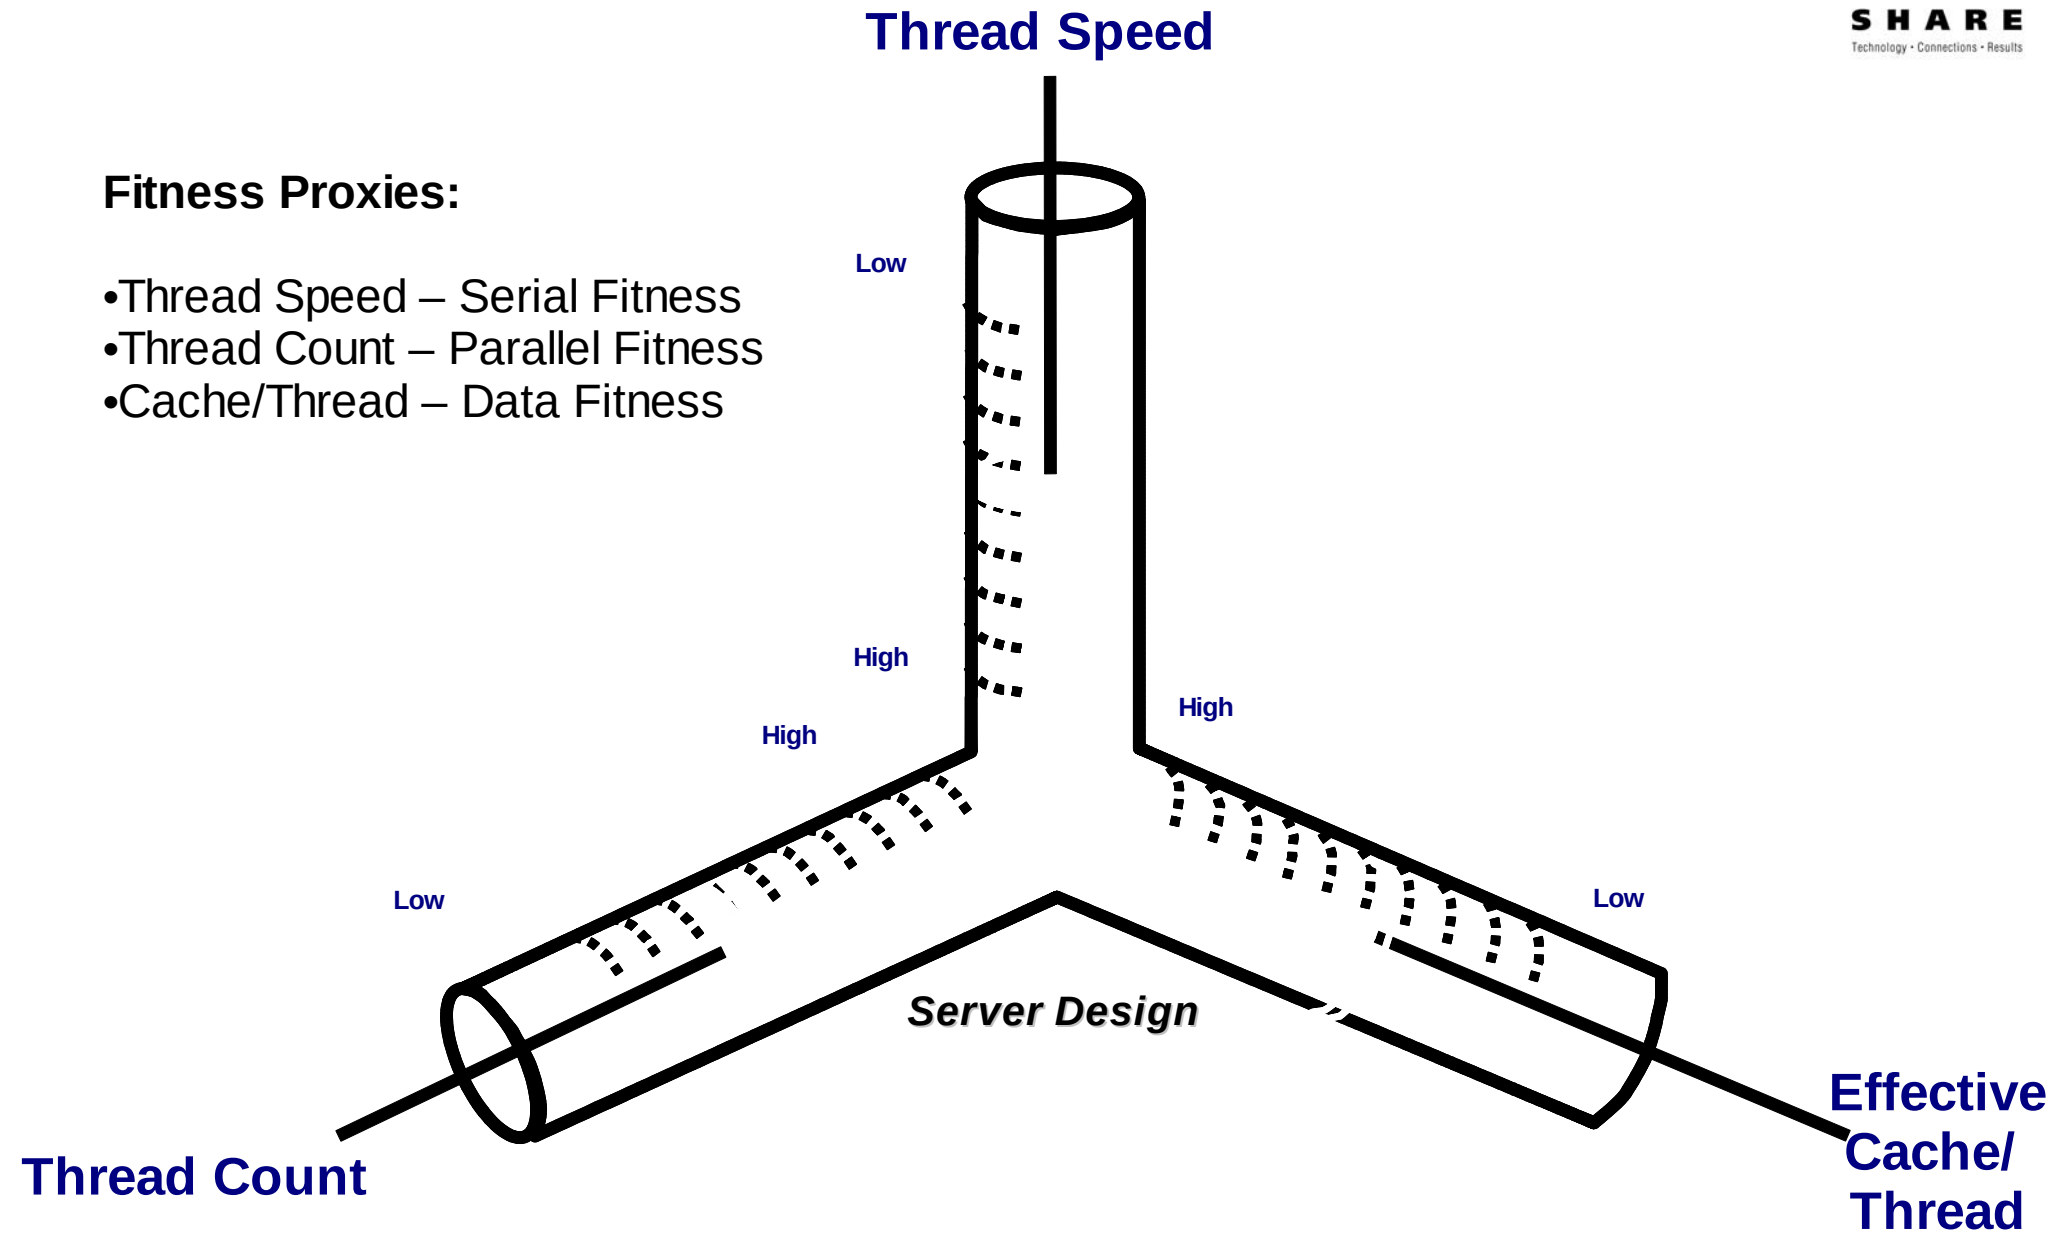

The fitness of 2010 server designs

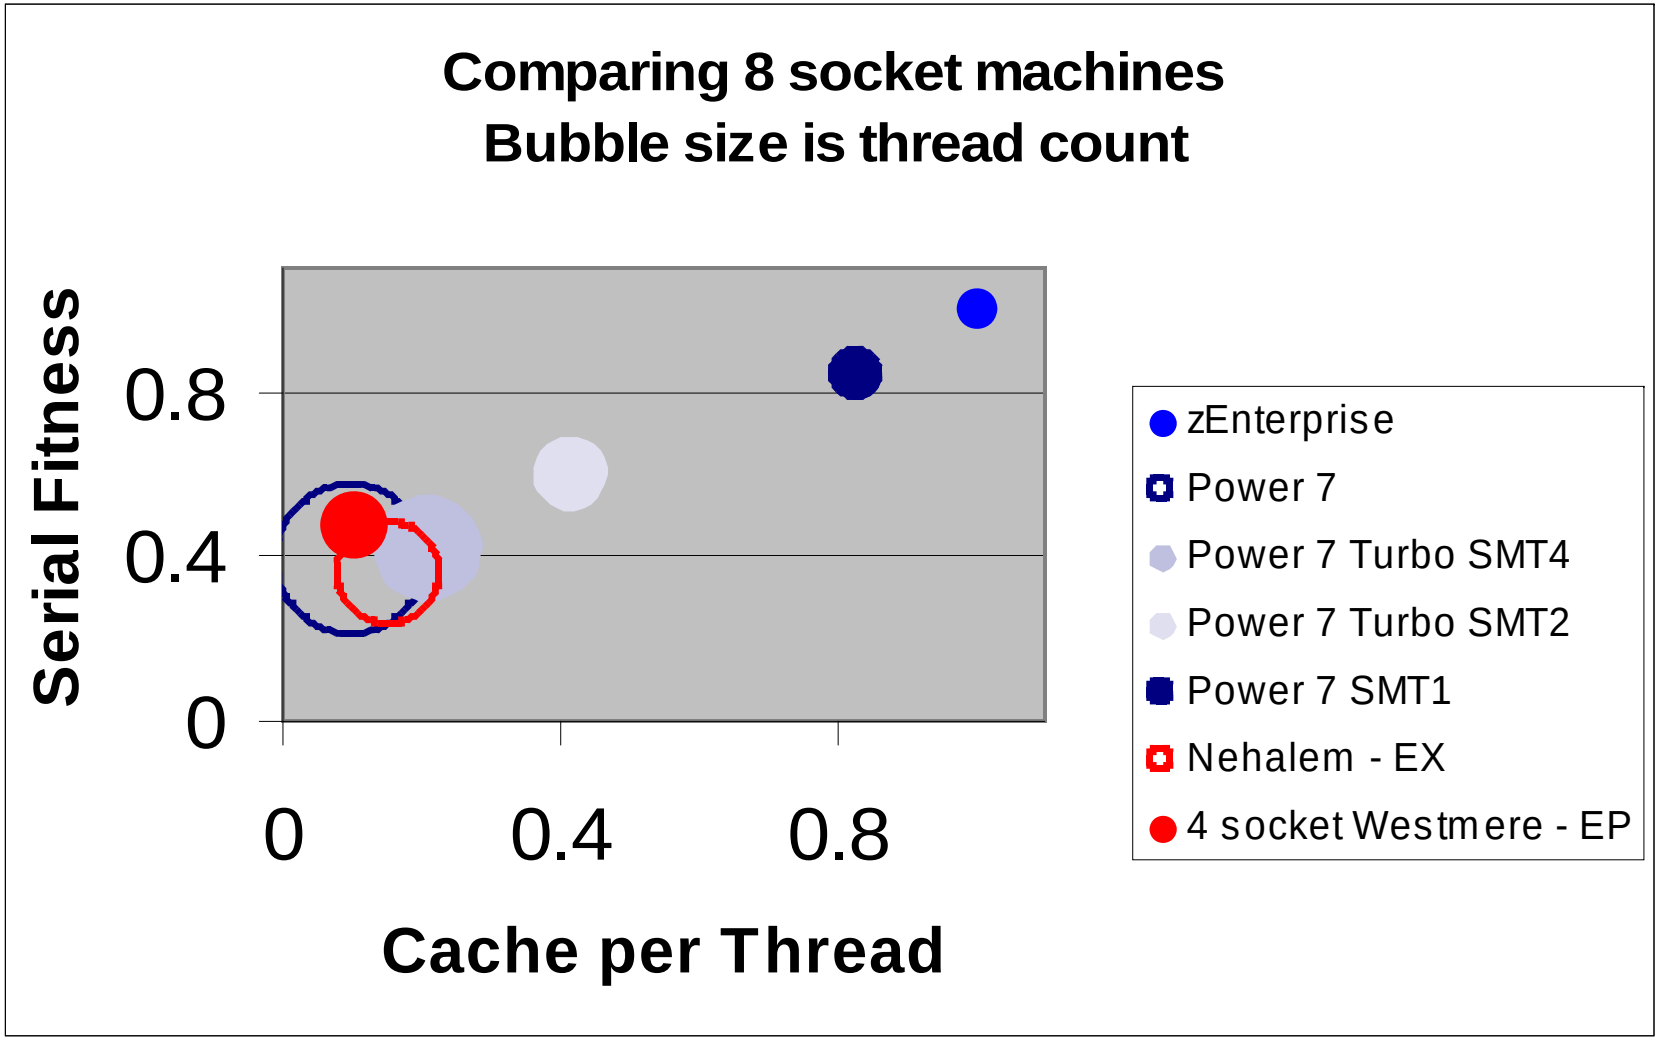

System z's bubble is so small and there is only 1.

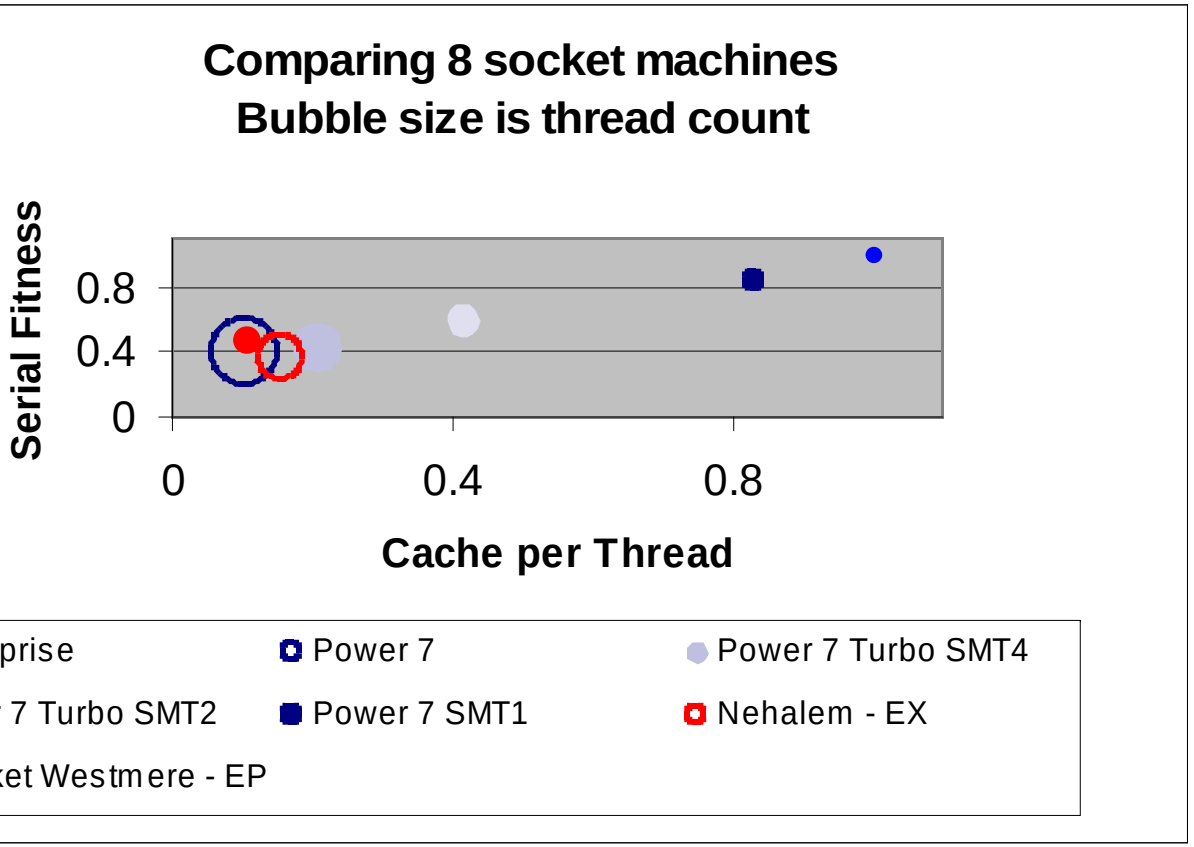

How come Power isn't shown as faster than z?

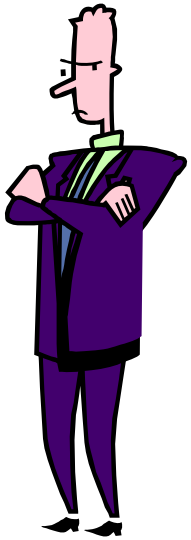

What is this? This is too much like math.
Can't you just tell me what to run where?

Power Products and Fit for purpose

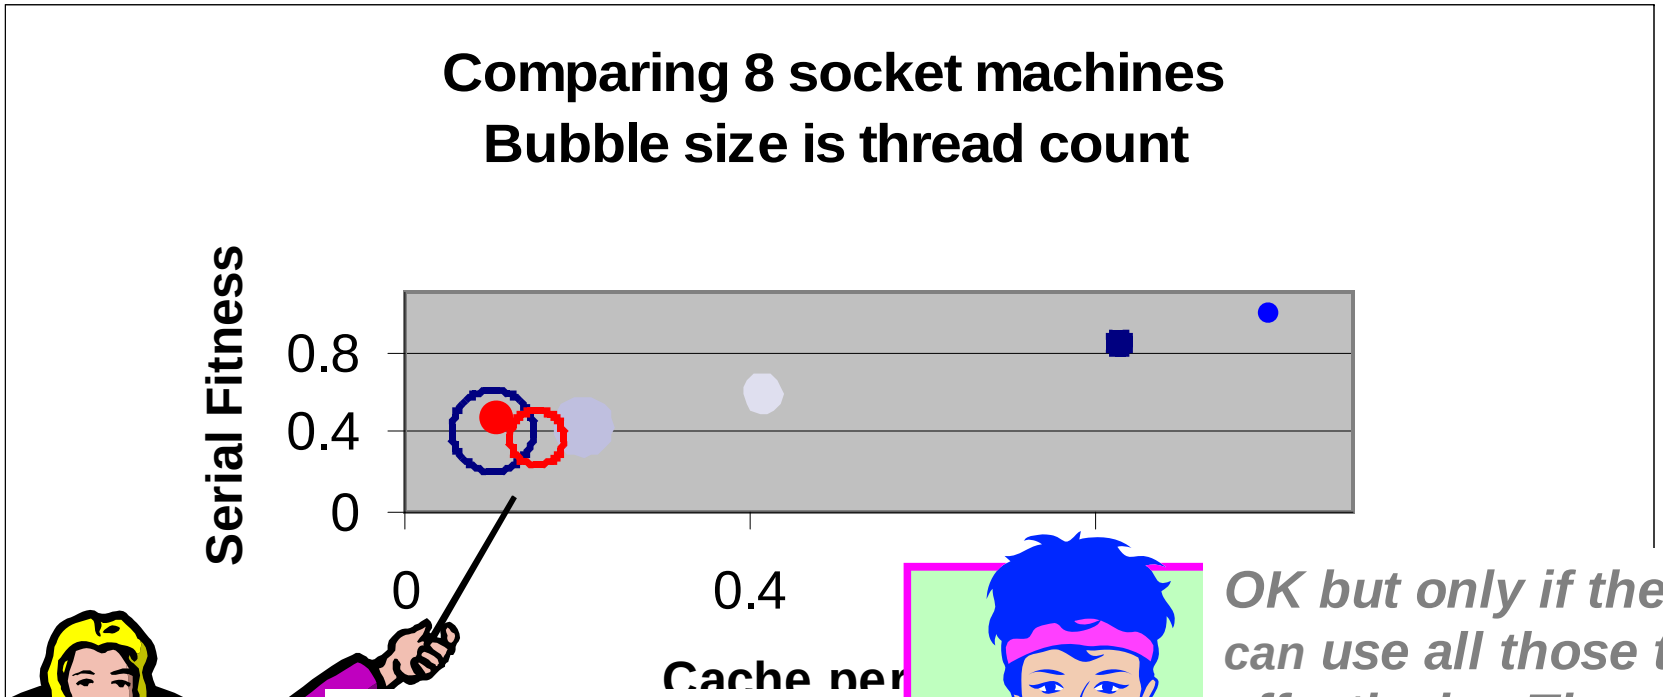

Power has the largest bubble, it must be better. Also benchmarks indicate better thread speed than this shows.
Note that there is a mode that fits any workload.

OK but only if the load can use all those threads effectively. The serial fitness axis includes locking and context switching that is usually avoided in standard benchmarks. Local factors can mitigate the raw fitness of any platform

System X products and Fit for Purpose

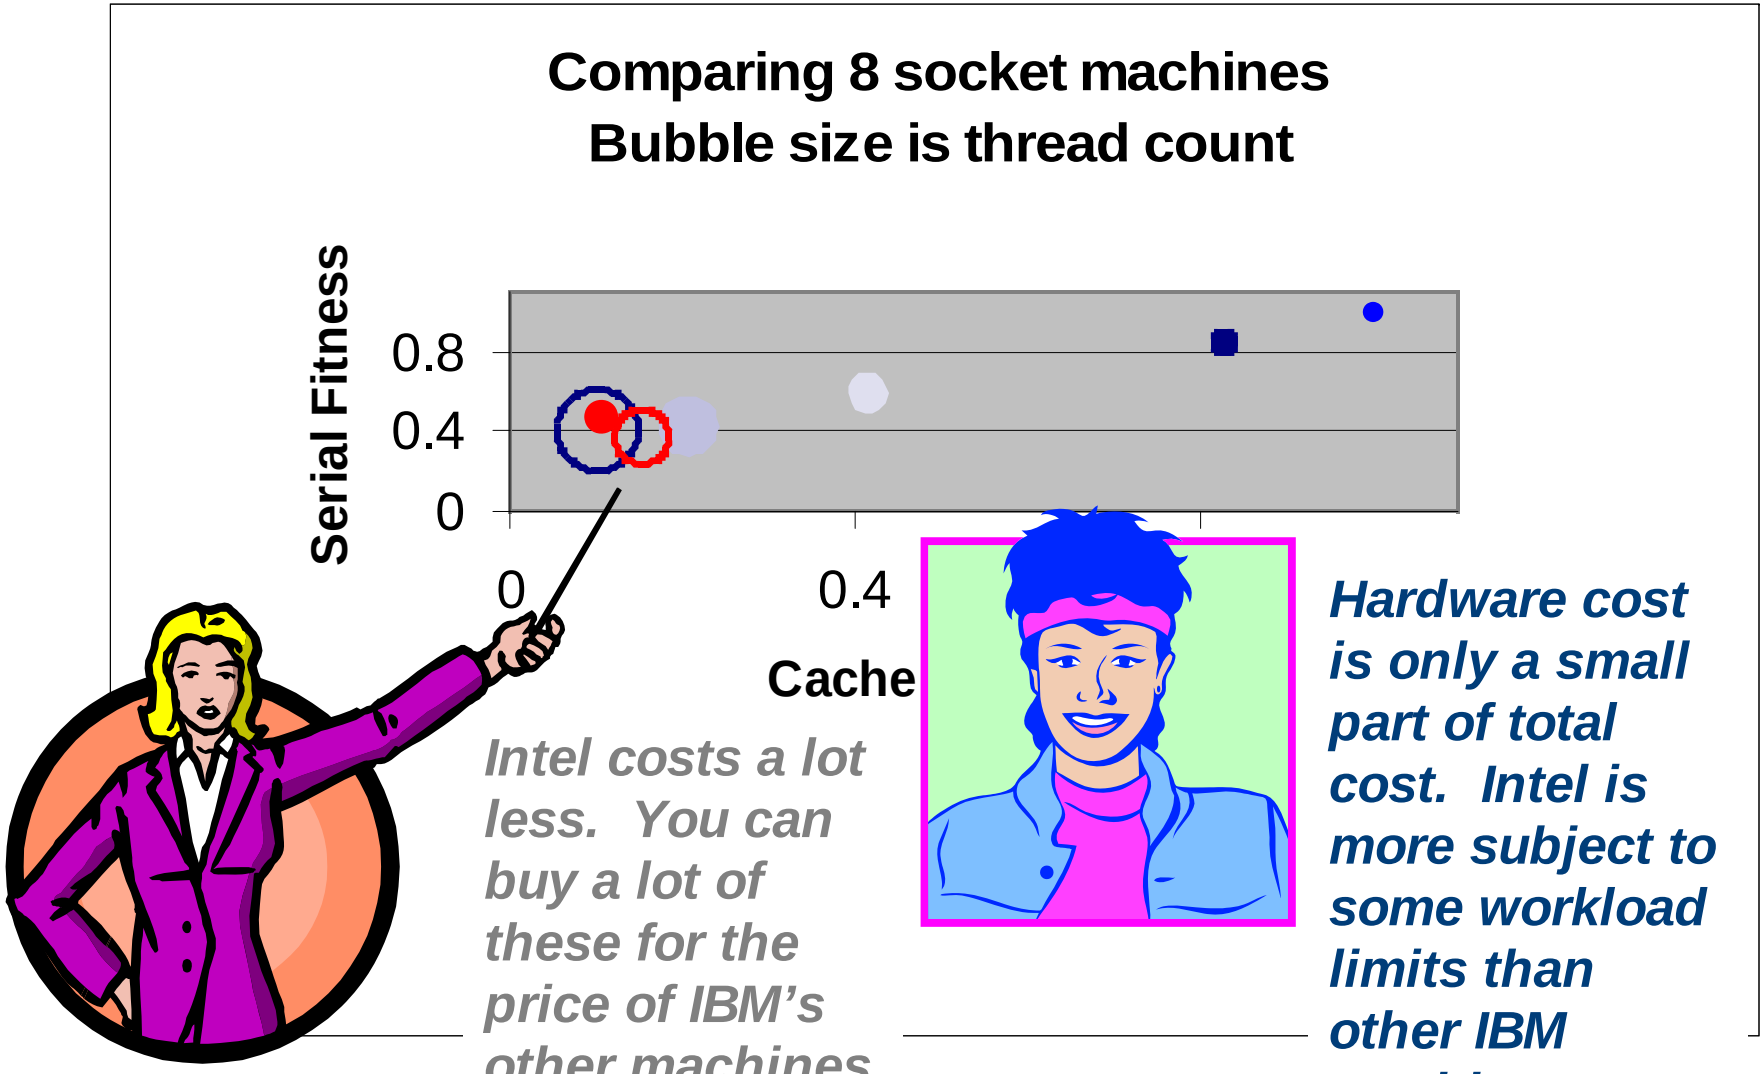

zEnterprise Product and Fit for Purpose

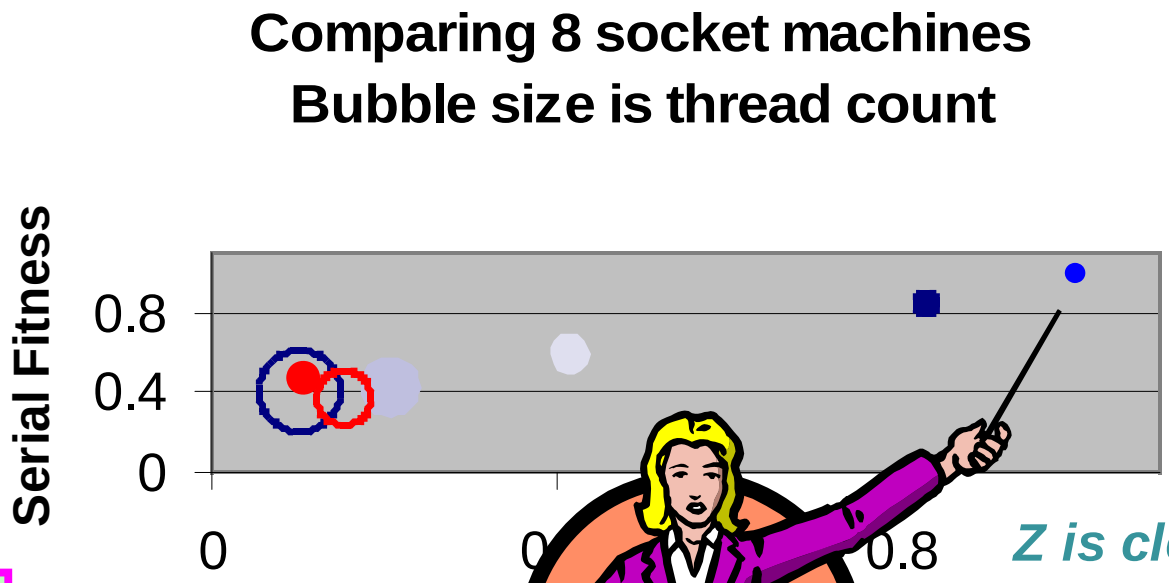

Ok but it is better only when the load does not have strong parallel leverage. In some cases Power blades in an zEnterprise Ensemble can help.

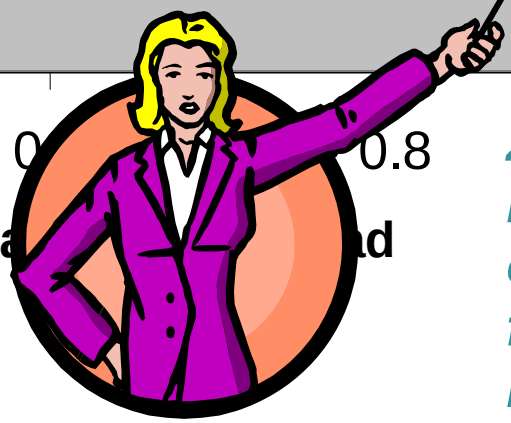

Z is clearly best of breed on 2 out of 3 fitness axes. It must be better!

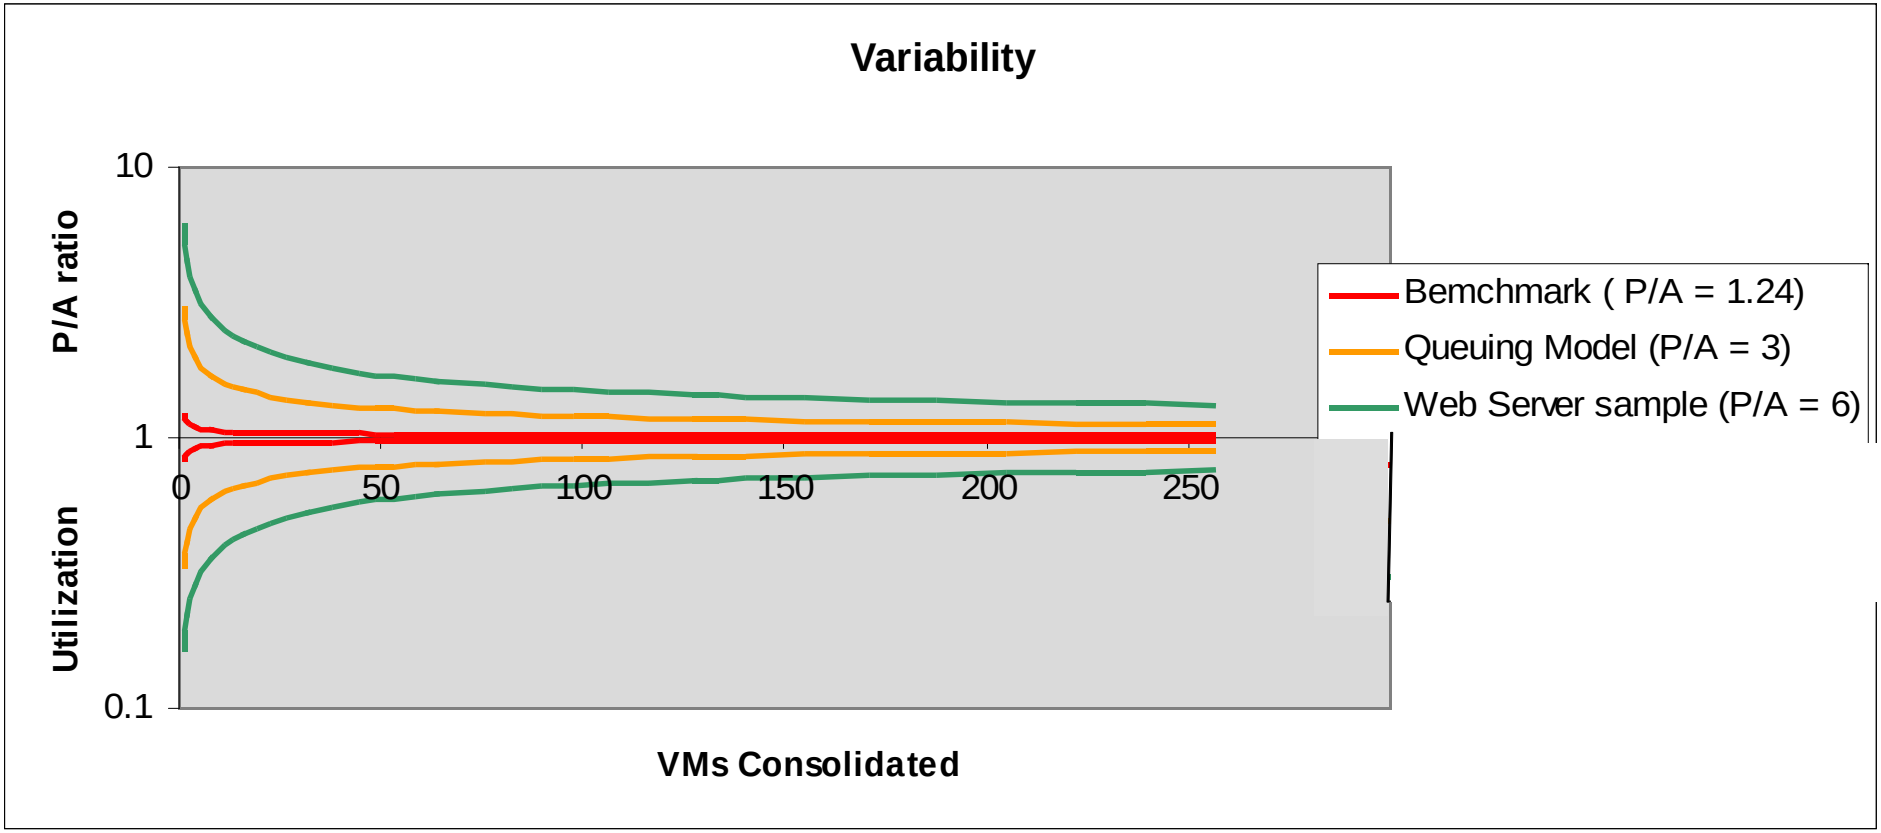

Variability causes the average utilization of individual loads to drop..
Consolidation raises average utilization by sharing head room.

Parallelization and virtualization are at odds

Parallelization trades off utilization efficiency for response time
Virtualization trades off response time for utilization efficiency

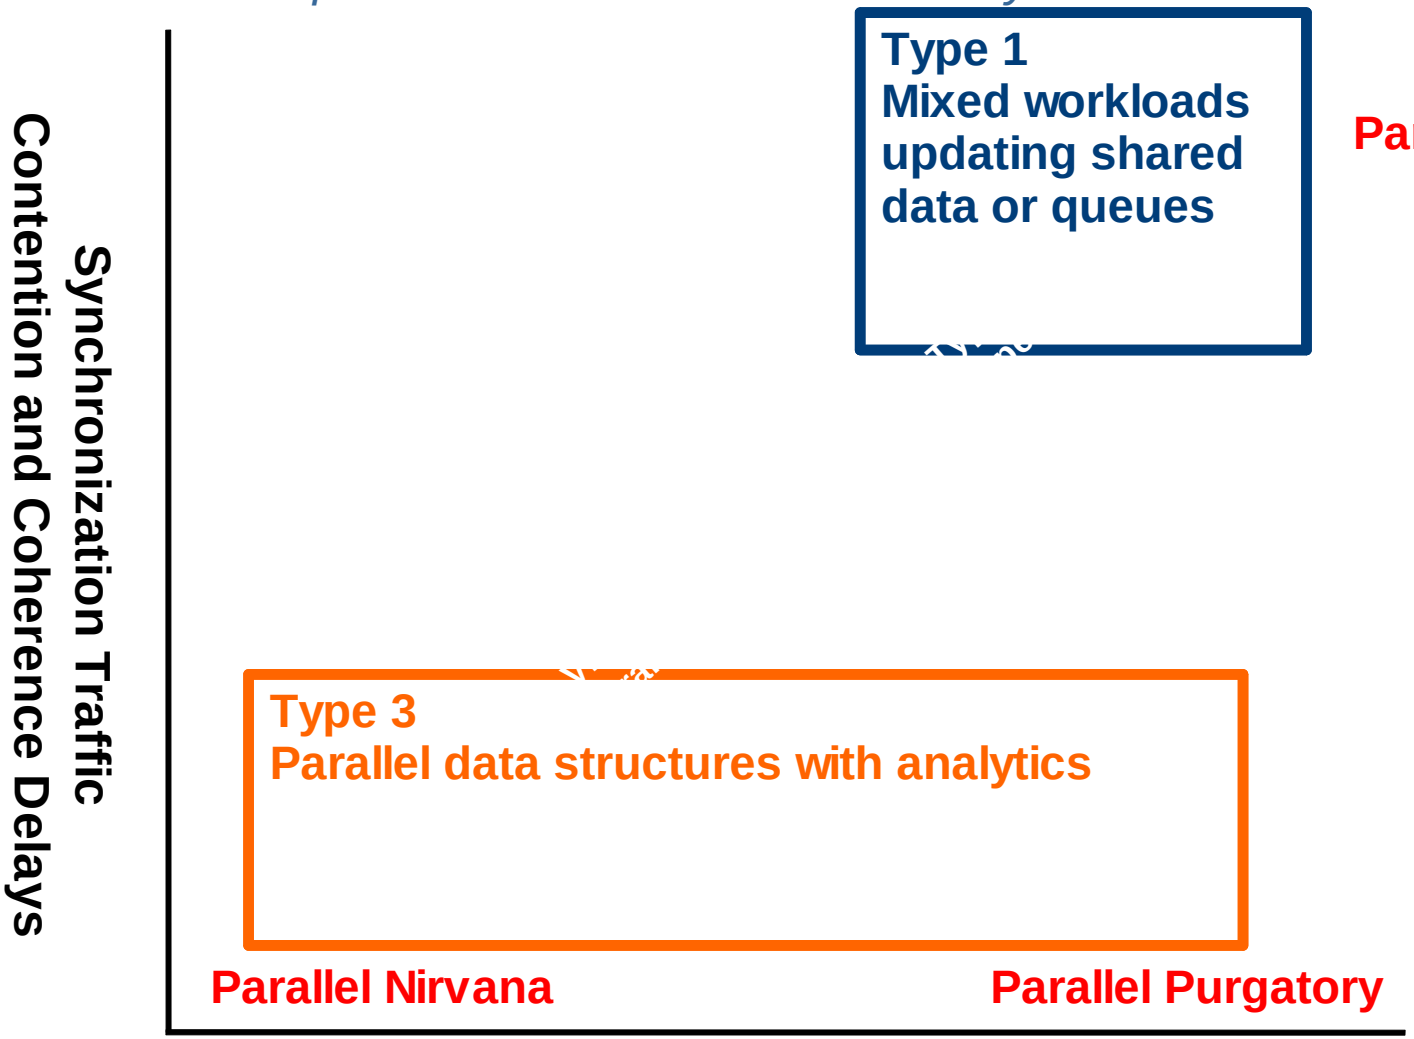

Virtualization is most effective with variable load, parallelization with steady state load.

Summary of workload optimization

- Workloads consist of:
  - application function (application design)
  - data structures (application design)
  - usage patterns (a local factor)
  - service levels (a local factor)
  - level of integration with other work (a local factor
- There are strong tradeoffs that result from variability in usage patterns, parallelism in the application/data and the service level.
  - You can't fully exploit parallelism and virtualization at the same time
- Machines have fitness for parallelism, fitness for data handling and fitness for serialization
  - You cannot maximize all three in the same design

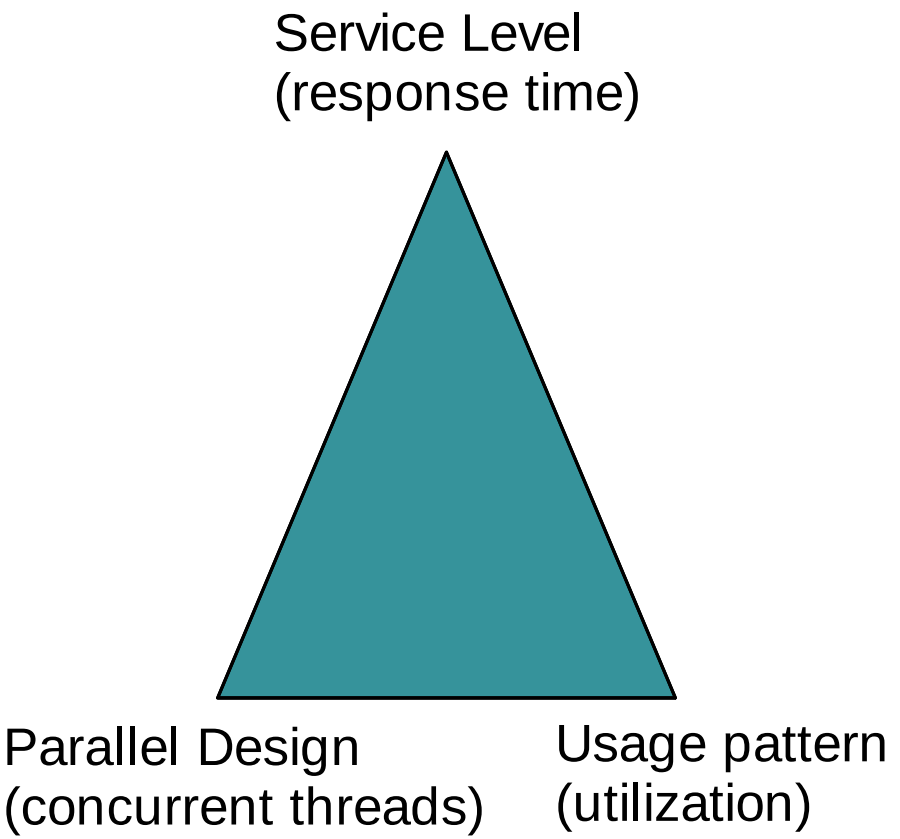

If you drive up resource sharing to gain utilization you reduce parallelism and increase response time. If the load is highly variable, adding parallelism will reduce the utilization dramatically. Notice that the design is only one corner of the triangle. Also note that integration increases sharing creating serializations and impacting service levels.