



Pulse2011

Next Gen Virtualization Platform

John J Thomas


jothomas@us.ibm.com

Topics

- “Real world” workloads have variations in demand
 - ▶ Benefits of large scale virtualization
- Impact of IO and other considerations
- What if you have different workloads with different characteristics? (Next Gen Virtualization Platform)

Variability In Workloads


- Typically benchmarks are run at a constant demand for workload, driving the machines to 100% utilization
- Real world workloads have variability in workload demand
 - ▶ Peaks and lows over a period of time




Variability And CPU Utilization

- Variability in workload demand has an effect on what mean CPU utilization can be tolerated on the consolidation platform
 - ▶ We would like to run utilization high enough to achieve the highest consolidation, but low enough to allow for peaks caused by variance in the workload
- The headroom you need to provision depends on the variance and also on what SLA you are trying to achieve

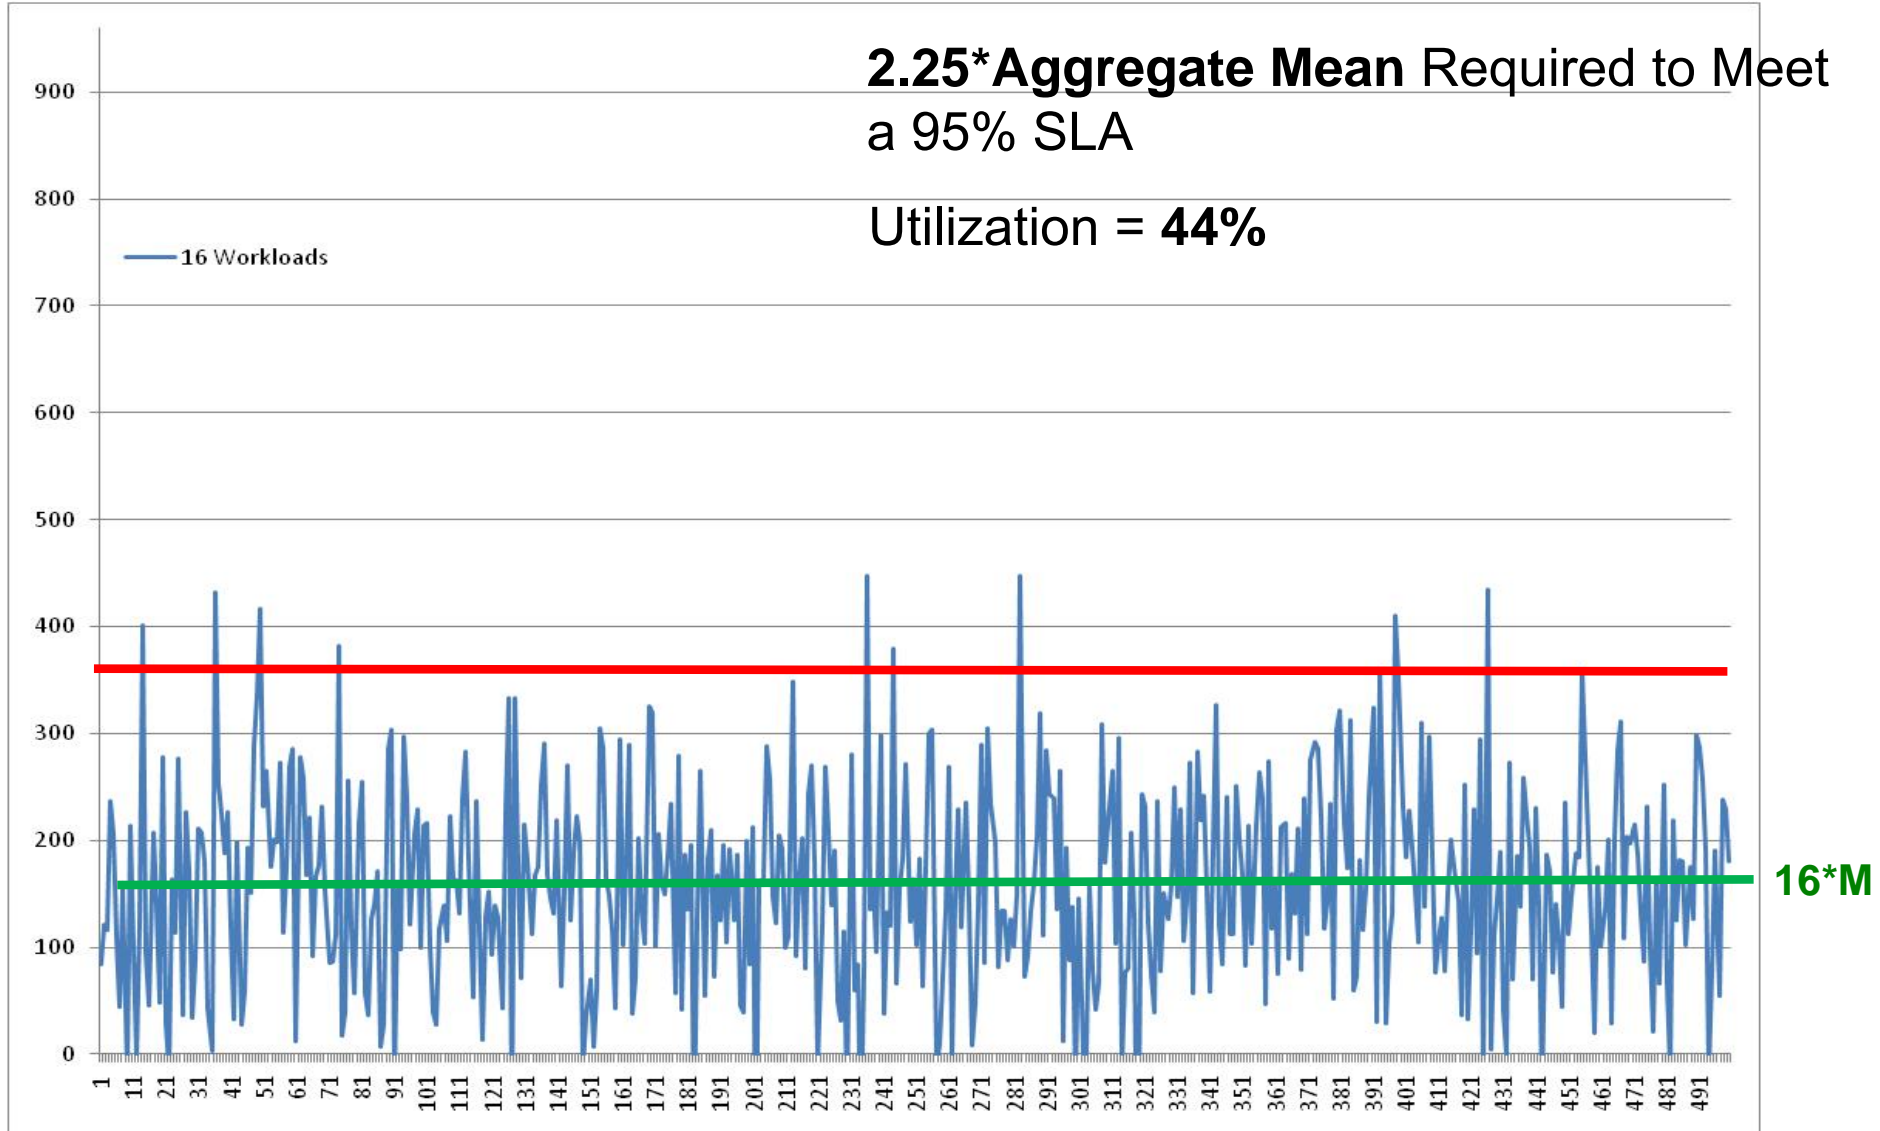
For A Single Workload, We Require A Machine Capacity 6.0x The Average Demand




When We Consolidate 4 Workloads We Only Require 3.5x Average Demand



When We Consolidate 16 Workloads We Only Require 2.25x Average Demand



When We Consolidate 144 Workloads We Only Require 1.42x Average Demand



Observations


- There is a benefit to large scale virtualization
 - ▶ The headroom required to accommodate variability goes up only by \sqrt{n} when n workloads are pooled
 - ▶ The larger the shared processor pool is, the more statistical benefit you get
 - Large scale virtualization platforms like 32-way z/VM LPARs, midrange and high end POWER systems etc. are able to consolidate large numbers of virtual machines because of this
- We need to refine “standard flat out” benchmarks to accommodate for real world variability in demand
 - ▶ Apply statistical methods to measured benchmark data

Example Of How This Methodology Is Applied - A 2009 Benchmark Study

We ran a benchmark to compare how many images can be consolidated

**Friendly Bank online banking benchmark
(WebSphere Application Server)**

**Intel servers x366
4 cores @ 3.66 GHz
12 GB memory**




**Workload for
each server:
5% utilization
22 tps**


**Existing non-virtualized
workload
on older servers**



**zLinux z10 EC
8 IFL cores @ 4.4 GHz**



**Consolidate VM
images on three
different platforms**




**Intel server x3950
8 cores @ 3.5 GHz**


**Each VM image
4 virtual cores
1 GB virtual memory**

Consolidation Ratios Derived From 2009 Study



Intel servers x366
4 cores @ 3.66 GHz
12 GB memory



Workload for each server:
5% utilization
22 tps



Linux on z10 EC
8 IFL cores @ 4.4 Ghz

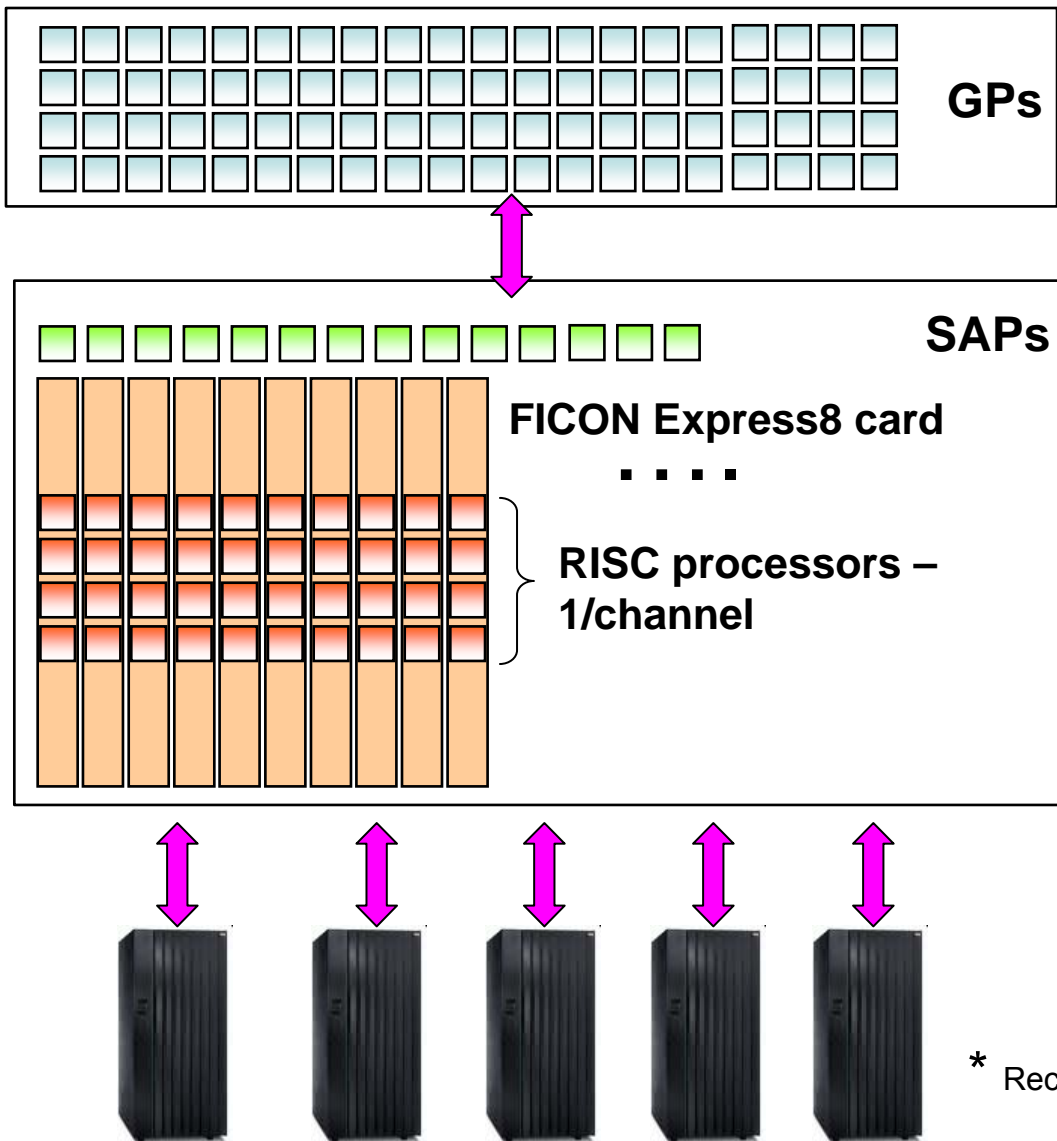


Intel server x3950
8 cores @ 3.5 Ghz

Topics

- “Real world” workloads have variations in demand
 - ▶ Benefits of large scale virtualization
- Impact of IO and other considerations
- What if you have different workloads with different characteristics?

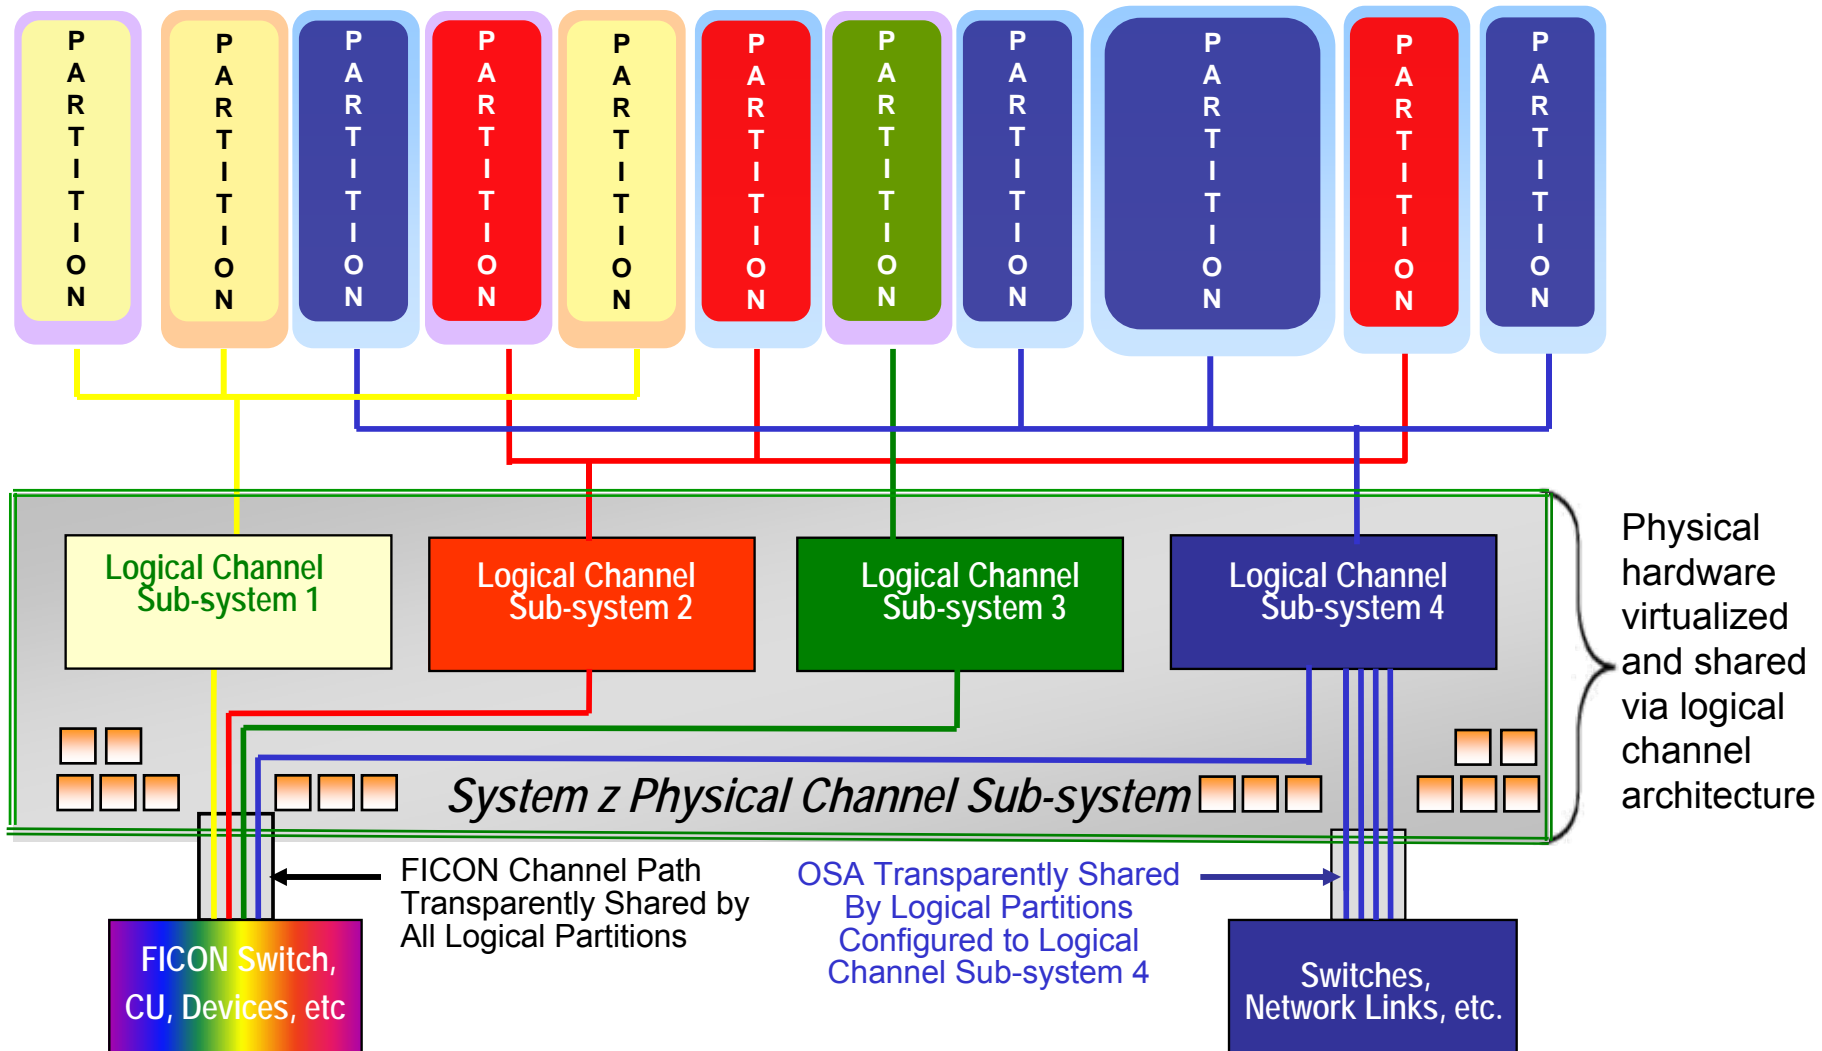
System z - Optimized For High I/O Bandwidth



- Up to 80 General Purpose (GP) or Specialty Engine processors
 - ▶ Execute business logic
- Up to 14 System Assist Processors (SAP) to manage I/O requests
 - ▶ Can sustain up to **2.2M IOPS*** operations per second
- Up to 84 physical FICON cards for I/O transfers
 - ▶ Up to **336 RISC channel I/O processors**
 - ▶ Up to 1024 logical channels
- IBM DS8800 Storage System
 - ▶ Up to **440K IOPS capability**


* Recommend 70% max SAP Utilization – 1.5M IOPS

Physical I/O Adapters And Channels Are Virtualized And Shared By The Consolidated Workloads





CPU Behavior In Systems Not Optimized For IO

Intel CPU As IO Load Increases



Case Study: Consolidate Heavy IO Workloads

Benchmark to determine how many workloads can be consolidated



Deploy as VMs on 8-core virtualized Intel Nehalem blade



Deploy as LPARs On PS701 8-core blade



How many VMs can be fit on to each container?

Online banking workloads, each driving **22** transactions per second, with **1 MB** asynchronous I/O per transaction


Deploy as Linux VMs on 32-way z/VM



Benchmark To Determine Relative Capacity

CPU Utilization – Friendly Bank with Heavy I/O Workload

Intel Blade Linux 8c POWER7 Blade AIX 8c z196 z/VM 32 CPU



Consolidation ratios derived from IBM internal studies. Results may vary based on customer workload profiles/characteristics.


Derive Consolidation Ratios

CPU Utilization – Friendly Bank with Heavy I/O Workload

Intel Blade Linux 8c

POWER7 Blade AIX 8c

z196 z/VM 32 CPU



Consolidation ratios derived from IBM internal studies. Results may vary based on customer workload profiles/characteristics.


Virtualization With Linux and z/VM On z196

- Benefits

- Larger scale of shared processor pools (32 cores)
 - ▶ Statistical benefit of sharing a larger pool of processors
- Software priced per core
- Enterprise Linux Server / Solution Edition pricing
- Dedicated I/O Sub-system offloads I/O processing
- Greater I/O bandwidth
- Virtualization of I/O processing resources
- Rock solid security
- Dynamically add processing to meet increase in demand

z/VM Security For Virtualization

- Operates without interference/harm from guest virtual machines
- Virtual machines cannot circumvent system security features
- Protects virtual machines from each other
- Ensures that a user only has access to resources specifically permitted
- Tracks who is accessing all system resources
- LPAR certified Common Criteria EAL5
- z/VM certified at Common Criteria EAL4+
- HiperSockets for highly secure internal networking
- Access to System z Crypto features
 - ▶ CPACF, CryptoExpress3

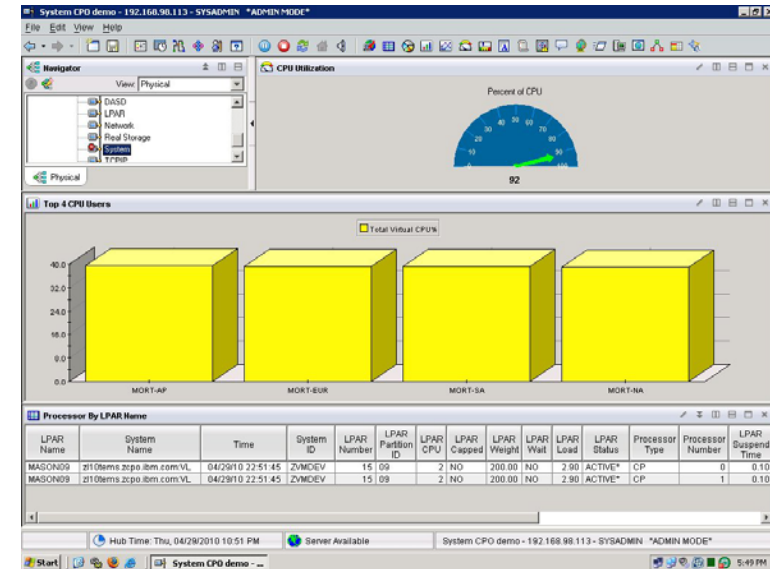


Linux On System z Workloads Inherit System z Qualities Of Service

- Reliability, availability, serviceability characteristics of System z
- Site failover for disaster recovery
- Capacity on demand upgrades
- Add physical processors to Linux environment without disruption

DEMO: Dynamically Add New Processor To z/VM LPAR To Handle Increased Workload

1. A customer has in-house Risk Analysis program running on Linux on System z
2. Increased workload to all 4 Linux guests is causing z/VM LPAR utilization of 90%+
3. Customer determines this is a long term trend - additional physical capacity needed
4. New capacity made available to LPAR as new Logical CPU, available for work
 - ▶ Without disruption in service




VMware can't recognize and take advantage of additional physical processors without bringing down and rebooting the system

Note: Assumes available processors on installed books

Topics

- “Real world” workloads have variations in demand
 - ▶ Benefits of large scale virtualization
- Impact of IO
- What if you have different workloads with different characteristics? (Next Gen Virtualization Platform)

zEnterprise - Environments Optimized For Different Workloads



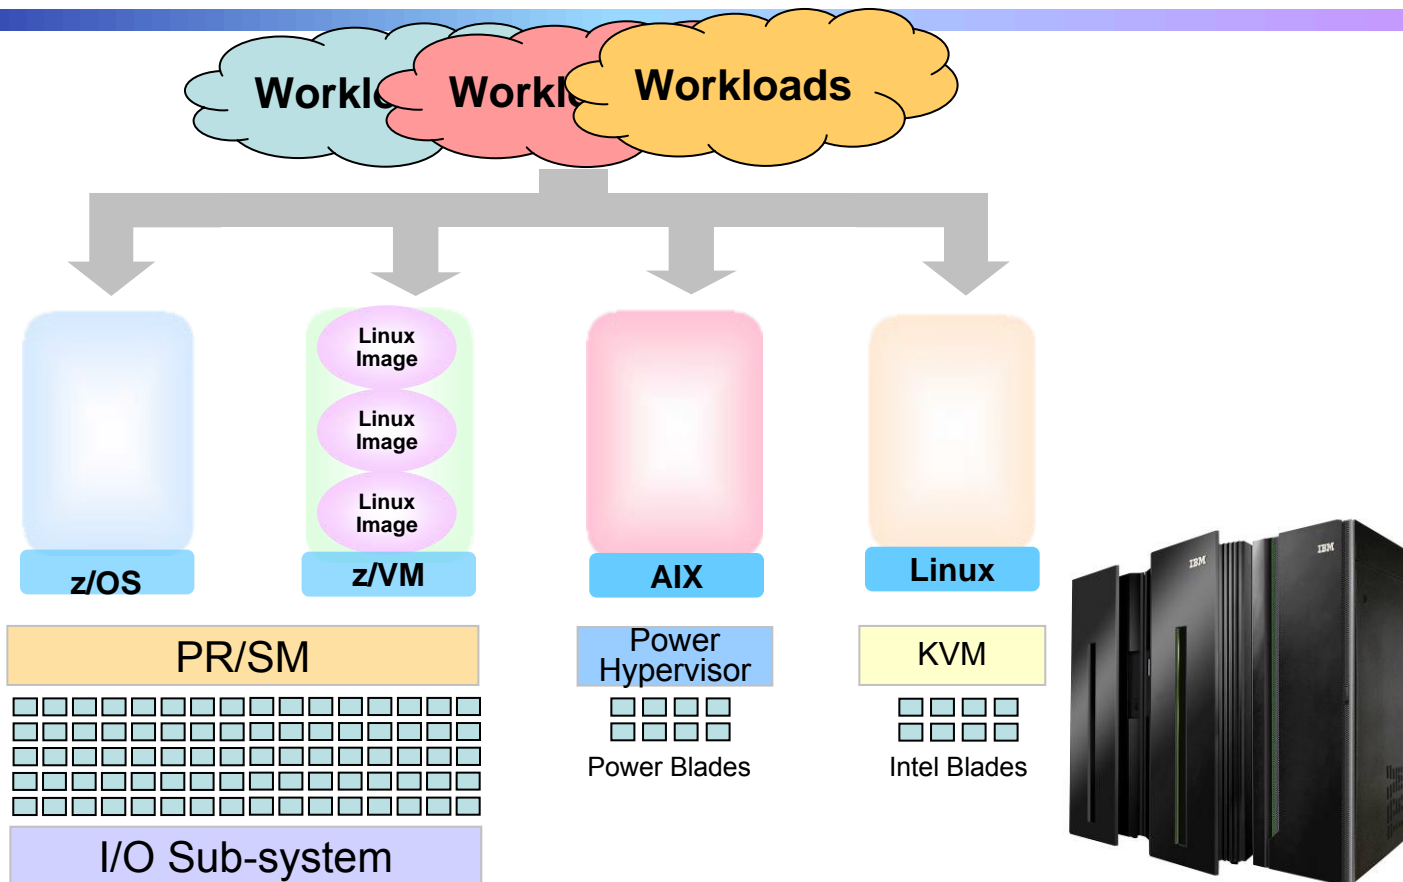
- Scale up to 80 cores in a frame
- Parallel Sysplex
- Dedicated I/O Sub System
- Superior qualities of service

- Scales to 8 cores per blade
- Fast processing threads
- Floating point accelerators

- Scales to 16 cores per blade
- Commodity I/O
- Lower qualities of service

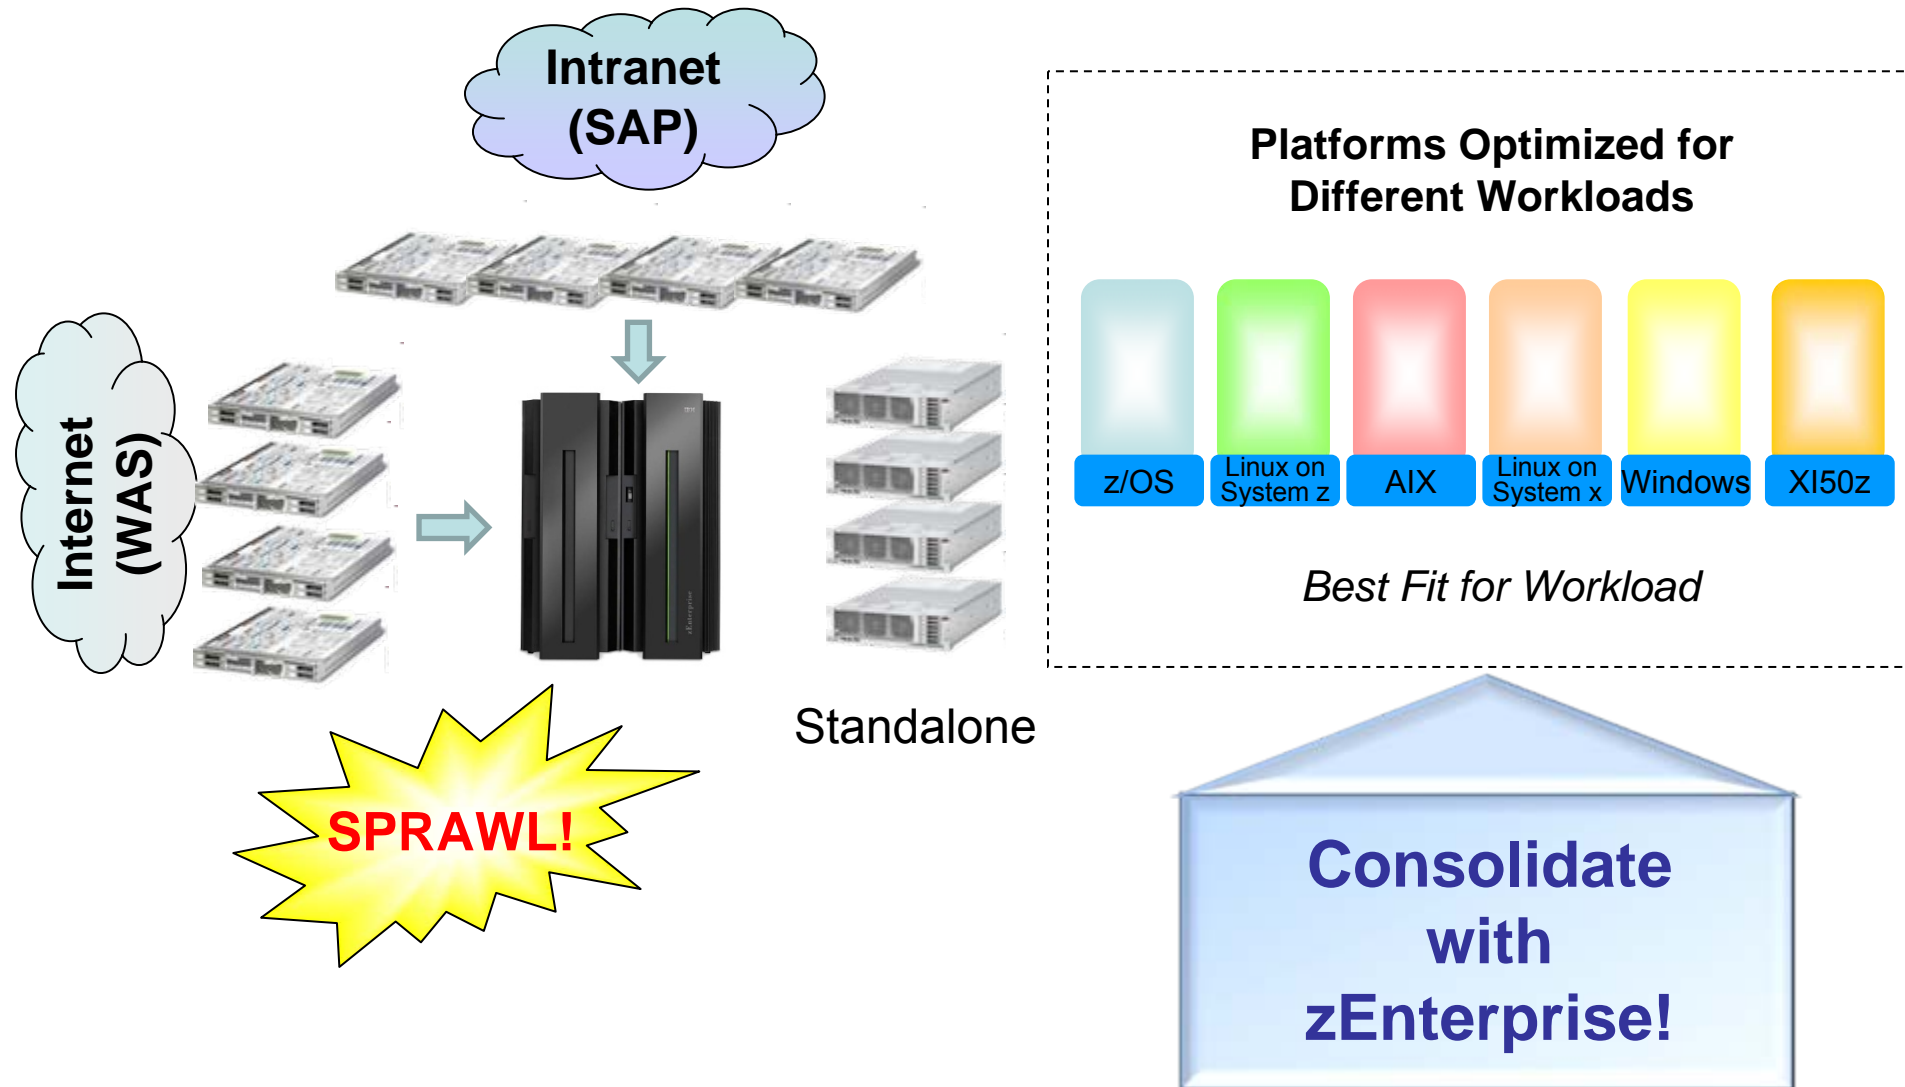
← Structured Management →

Workload Characteristics Influence The Optimal Deployment Decision




Deploy or consolidate workloads on the environment best suited for each workload, manage everything centrally

Address Sprawl With zEnterprise Multi-Architecture Environment



Simplifying Hardware Infrastructure

Our front end infrastructure is too complex!



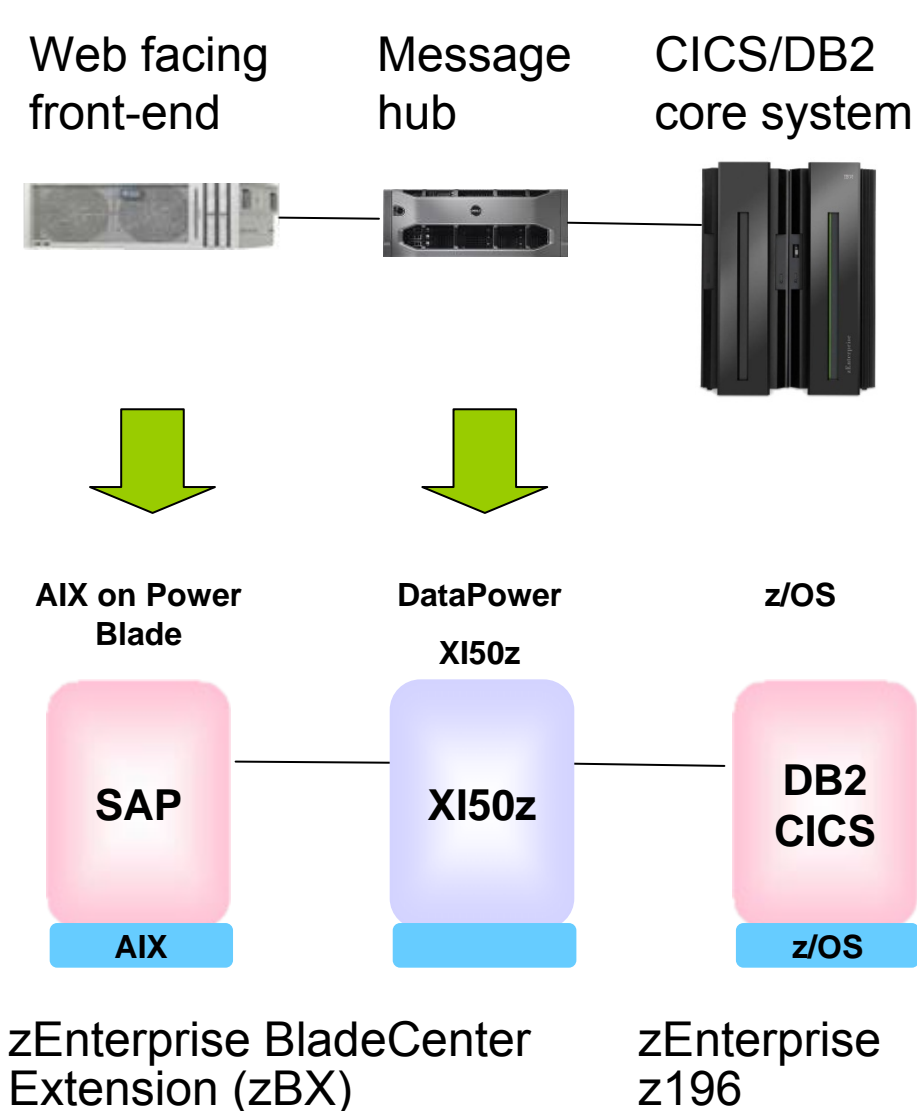
Customer

You can simplify by consolidating everything onto a single platform!



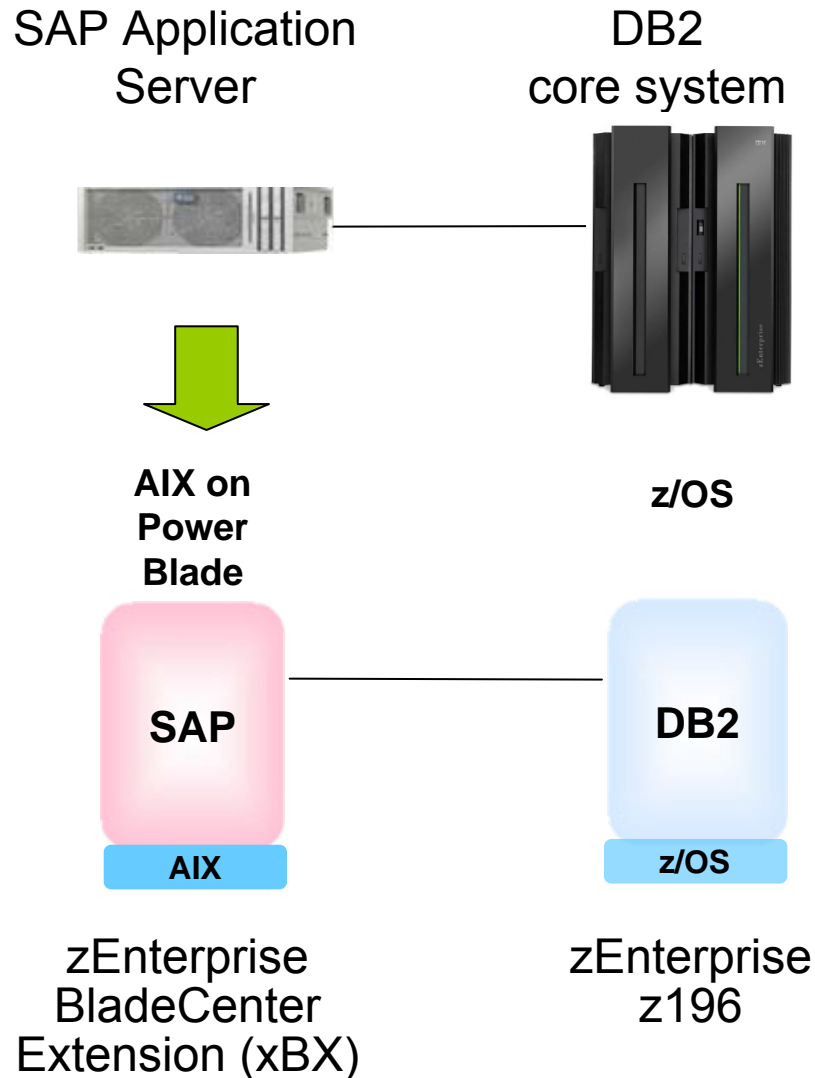
IBM

Collapse Web Front End Workloads Into zEnterprise Platform



- Run as ensemble of virtual servers
- Unified management of virtual machines
- Manage ensemble as a single workload with service goals
- Dynamic adjustment of CPU resources drives 10% higher utilization
- Assign best fit to Power blade and XI50z for lowest cost per workload
- Embedded pre-configured data network

Collapse SAP Front End Applications Into zEnterprise Platform



- Run as ensemble of virtual servers
- Unified management of virtual machines
- Manage ensemble as a single workload with service goals
- Dynamic adjustment of CPU resources drives 10% higher utilization
- Assign best fit to Power blade and XI50z for lowest cost per workload
- Embedded pre-configured data network

Case Study – Consolidate 880 Standalone Workloads Onto zEnterprise

- Distributed workload profile is a mix of
 - 784 light weight
 - 56 heavy weight (cpu intensive)
 - 40 heavy I/O
- What is the most cost effective way to consolidate/deploy all these workloads?

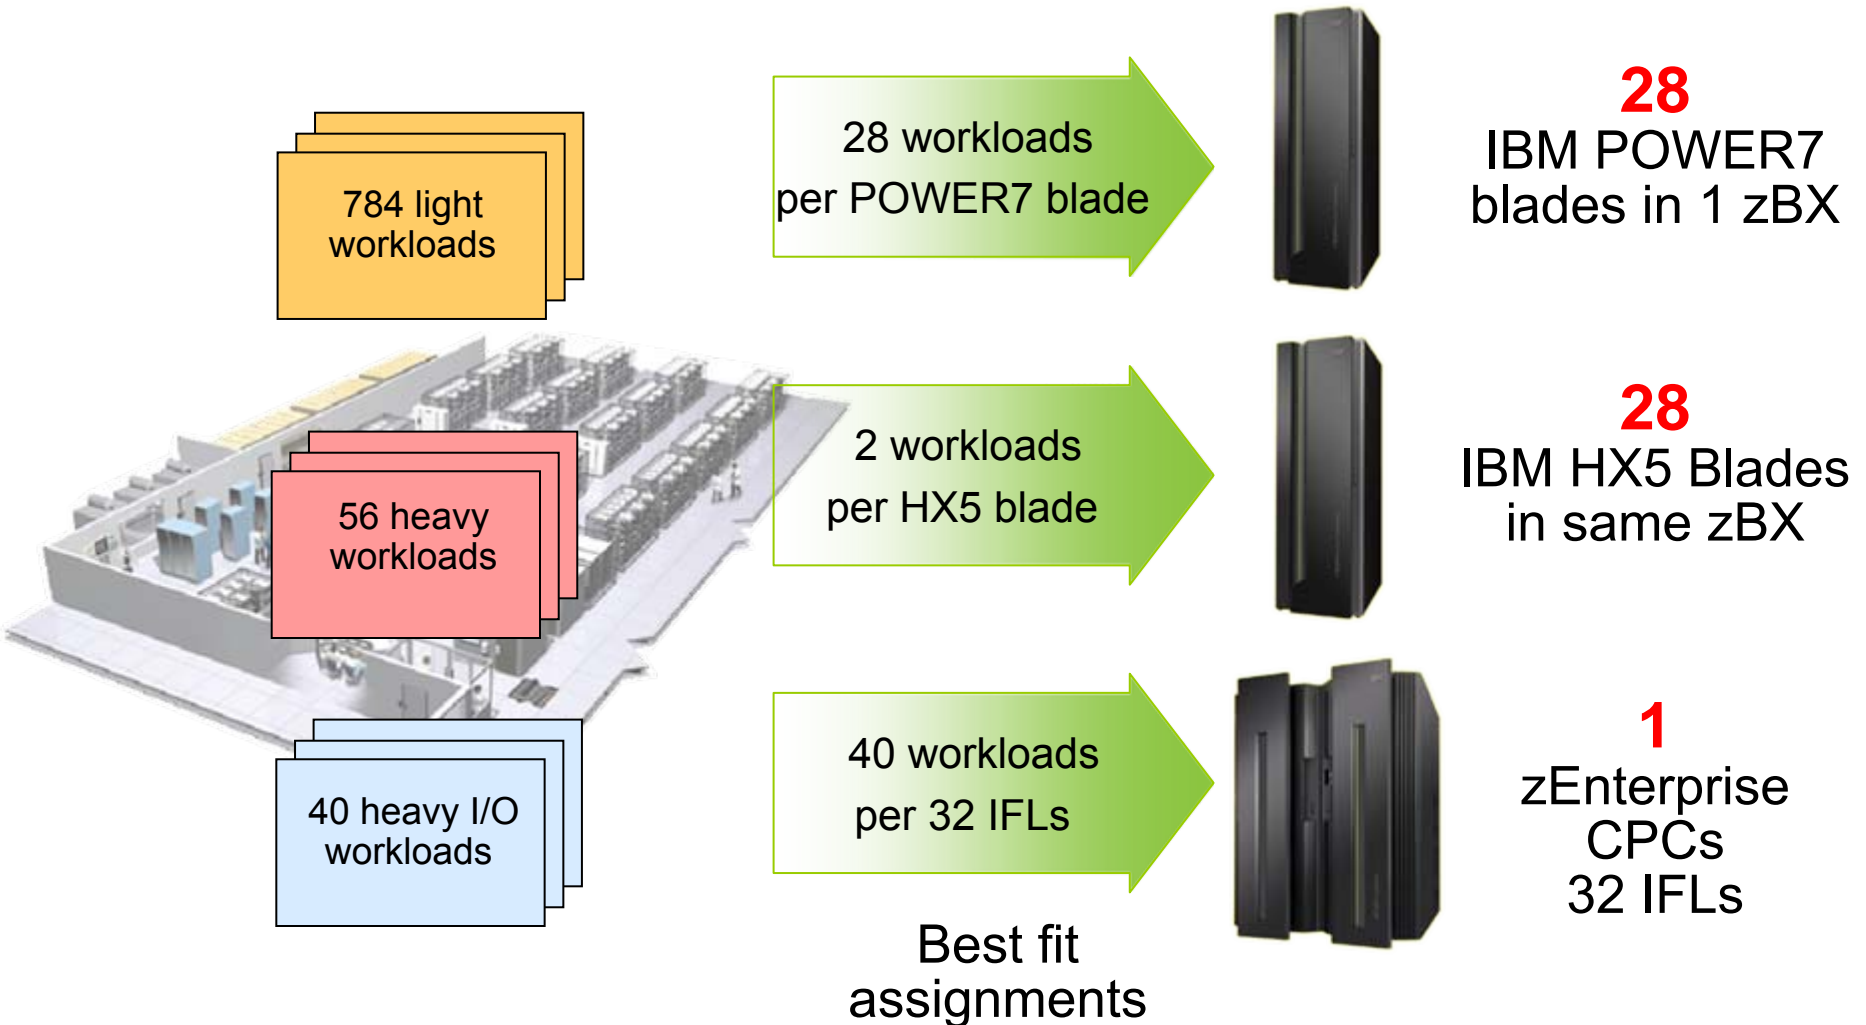
Competitor



zEnterprise




What Does It Cost To Deploy 880 Workloads On zEnterprise?



Server configurations are based on consolidation ratios derived from IBM internal studies. Prices are in US currency, prices will vary by country

Options For Deploying Workloads With Different Characteristics – Best Fit Strategy On zEnterprise



\$37.8M TCA (3 years)

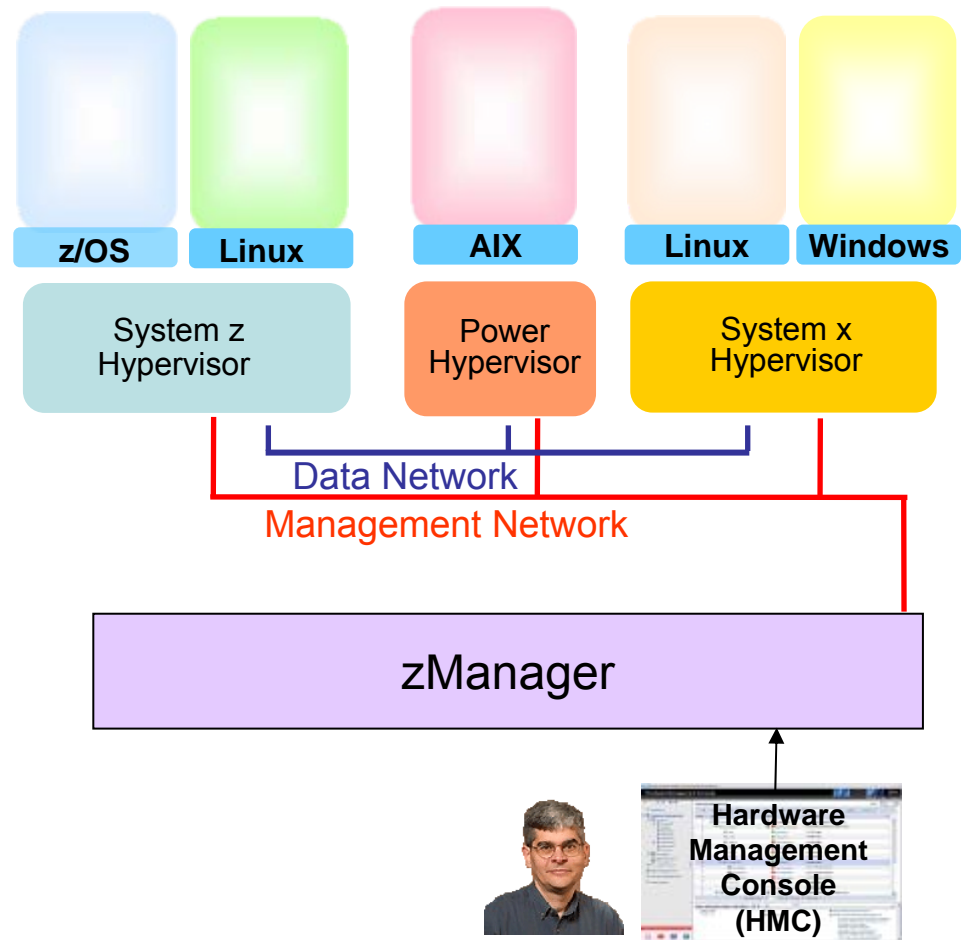
\$19.8M TCA (3 years)

Server configurations are based on consolidation ratios derived from IBM internal studies. Prices are in US currency, prices will vary by country

48% less


Manage All Virtual Servers With zManager

- From one console, create virtual machines in z/VM and in zBX hypervisors
- Start / stop / delete virtual machines under zManager control
- Create virtual networks
- Monitor resource usage
 - ▶ CPU, Memory, Power consumption




Leverage Multi-Architecture Virtualization Capabilities Of zEnterprise To Cut IT Costs

\$\$\$\$\$\$\$\$



\$\$



**Mainframe workloads
+
Distributed workloads
deployed on
zEnterprise with
Best Fit for Cost**