

IBM System z Technology Summit



zEnterprise – The First System Of Systems

The Economics Of Workload Optimization

David Rhoderick


25th May 2011




Our Agenda Today

10 Minutes	<i>Welcome by Regional Sales Exec</i>
8:40 am	The Economics of Workload Optimization
9:50 am	Virtualization and Consolidation on zEnterprise
11:00 am	Reducing IT Labor Costs
12:00 pm	<i>Lunch</i>
1:00 pm	Deploying Web Applications
2:10 pm	Data Serving and Business Analytics on a Single Platform
3:20 pm	Unify Mainframe and Distributed Development


DEMO: Architecture



System x 3950
8 x 3.5GHz Xeon MP
64GB RAM



z10-BC
2098-E10
256GB RAM



System x VMware images running as desktop or server clients to System z

Smarter Computing With The IBM zEnterprise System

*World's first **multi-architecture platform***

*Workloads optimized on **best fit for purpose platforms at the lowest cost***

*Unified management of system-wide resources **drives down cost of acquisition and operation***



A System of Systems

Smarter Planet Solutions Usually Include Different Workloads

Transaction Processing and Database

- *Application Database*
- *Data Warehousing*
- *Online Transaction Processing*
- *Batch*

Analytics

- *Data Mining Applications*
- *Numerical*
- *Enterprise Search*


Business Applications

- *Enterprise Resource Planning*
- *Customer Relationship Management*
- *Application Development*


Web, Collaboration and Infrastructure

- *Systems Management*
- *Web Serving/Hosting*
- *Networking*
- *File and Print*


Different Workloads Have Different Characteristics



- Moderate transaction workload
- High I/O bandwidth
- High quality of service requirements

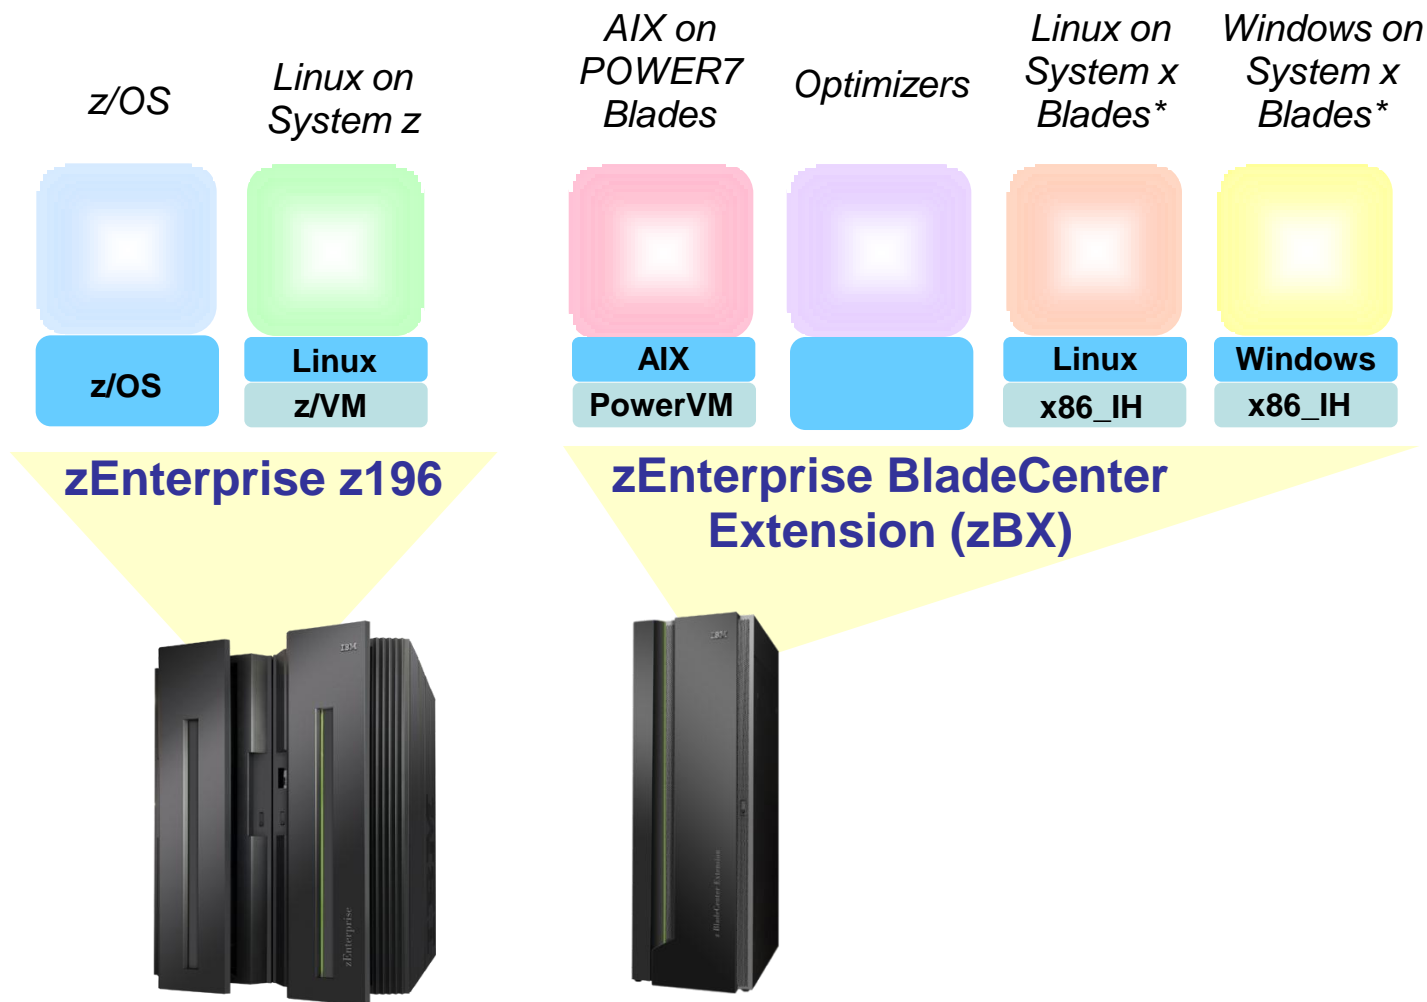


- High processing intensity
- Integer or floating point



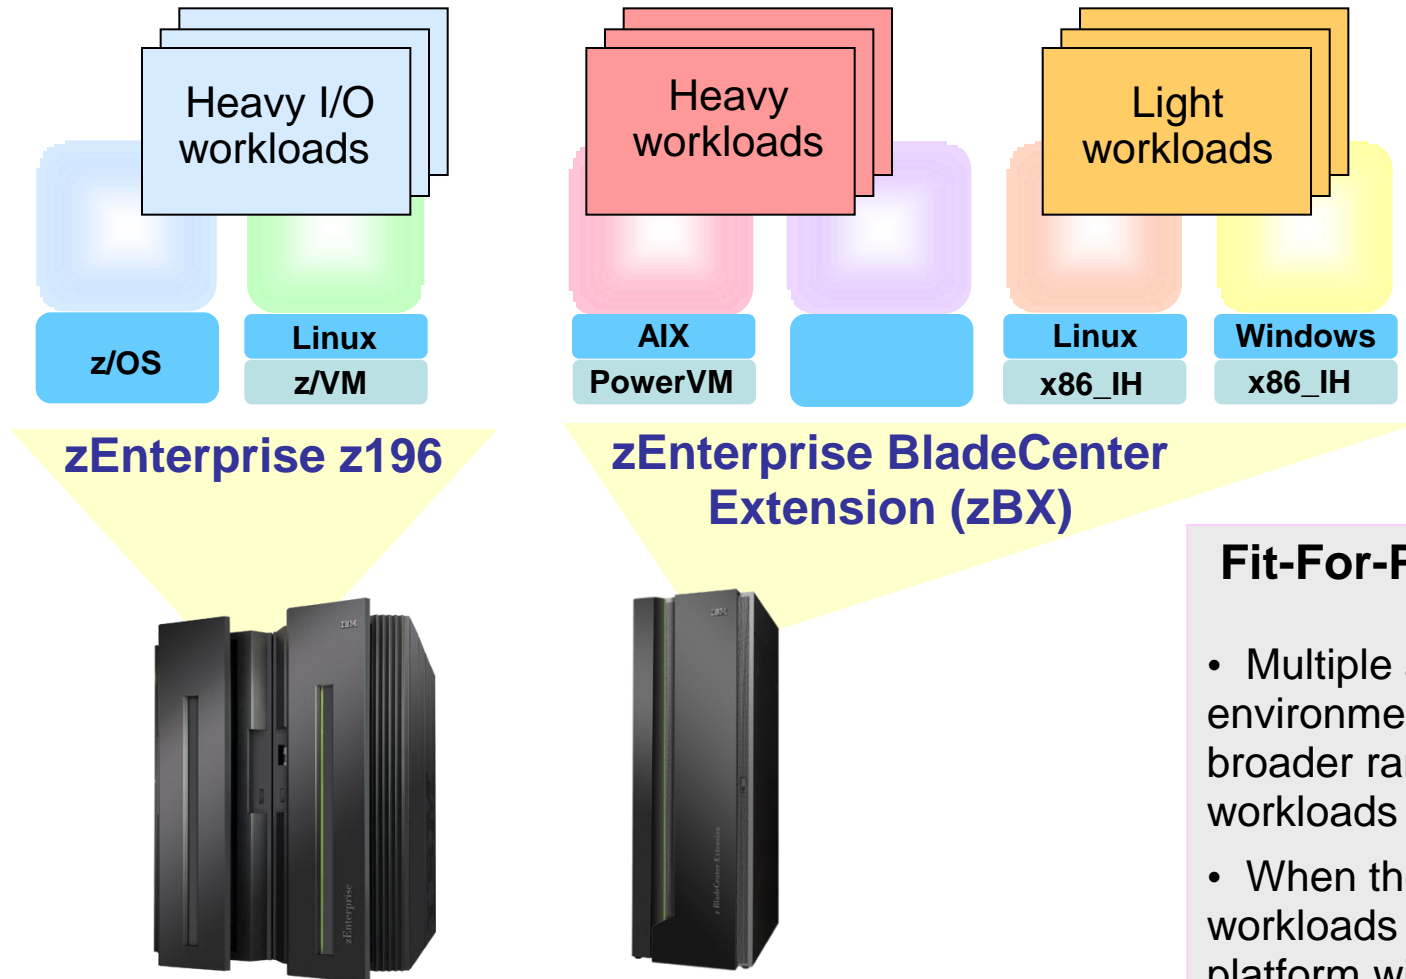
- Light to moderate processing
- Modest quality of service requirements

zEnterprise Has Different Environments For Different Workload Requirements



*All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represents goals and objectives only.

zEnterprise Provides Platforms For Workload Optimization




Fit-For-Purpose Strategy

- Multiple architecture environments to support a broader range of existing workloads
- When there is a choice, workloads can be assigned to platform with **lowest cost per workload**

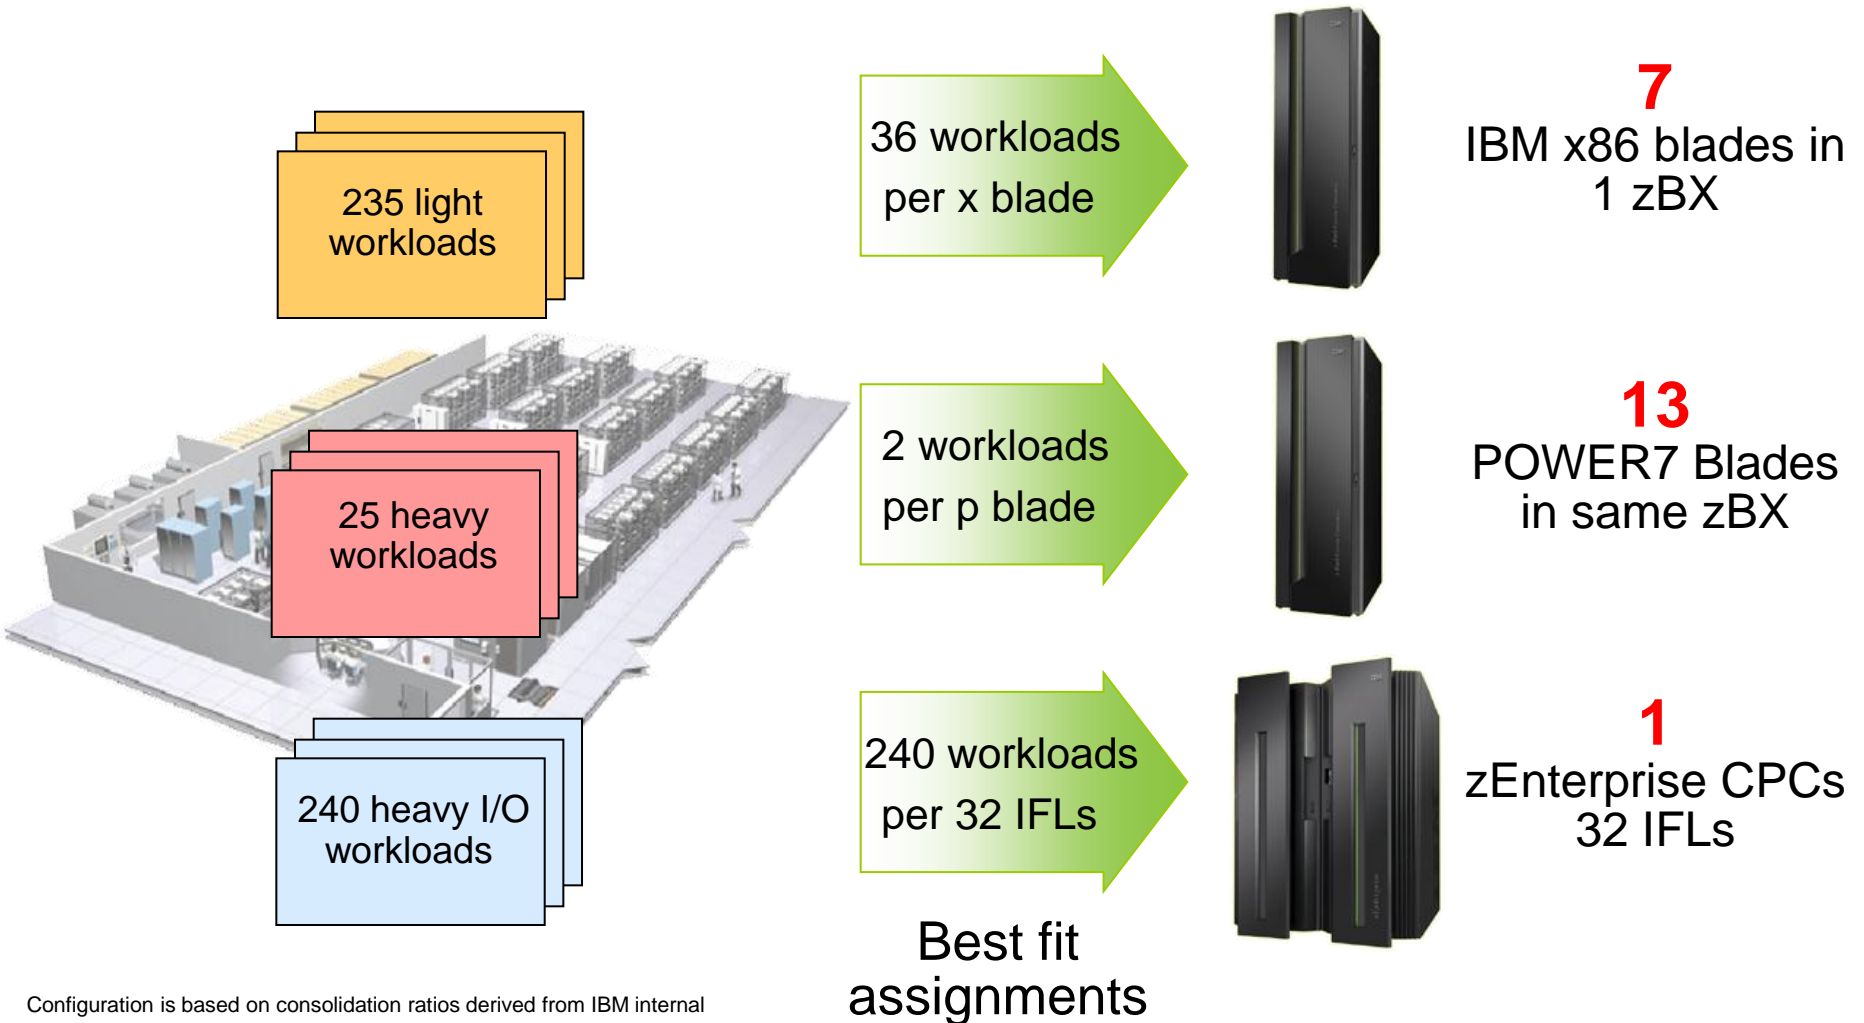
*All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represents goals and objectives only.

What Does It Cost To Deploy 500 Workloads On Virtualized Intel Servers?




IBM analysis of a customer scenario with 10,000 distributed workloads. Deployment configuration is based on consolidation ratios derived from IBM internal studies.

What Does It Cost To Deploy 500 Workloads On zEnterprise?



Configuration is based on consolidation ratios derived from IBM internal studies. z196 32-way performance projected from z196 8-way and z10 32-way measurements.
The zBX with x86 blades is a statement of direction only. Results may vary based on customer workload profiles/characteristics.


Compare Server Cost Of Acquisition



51%
less

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

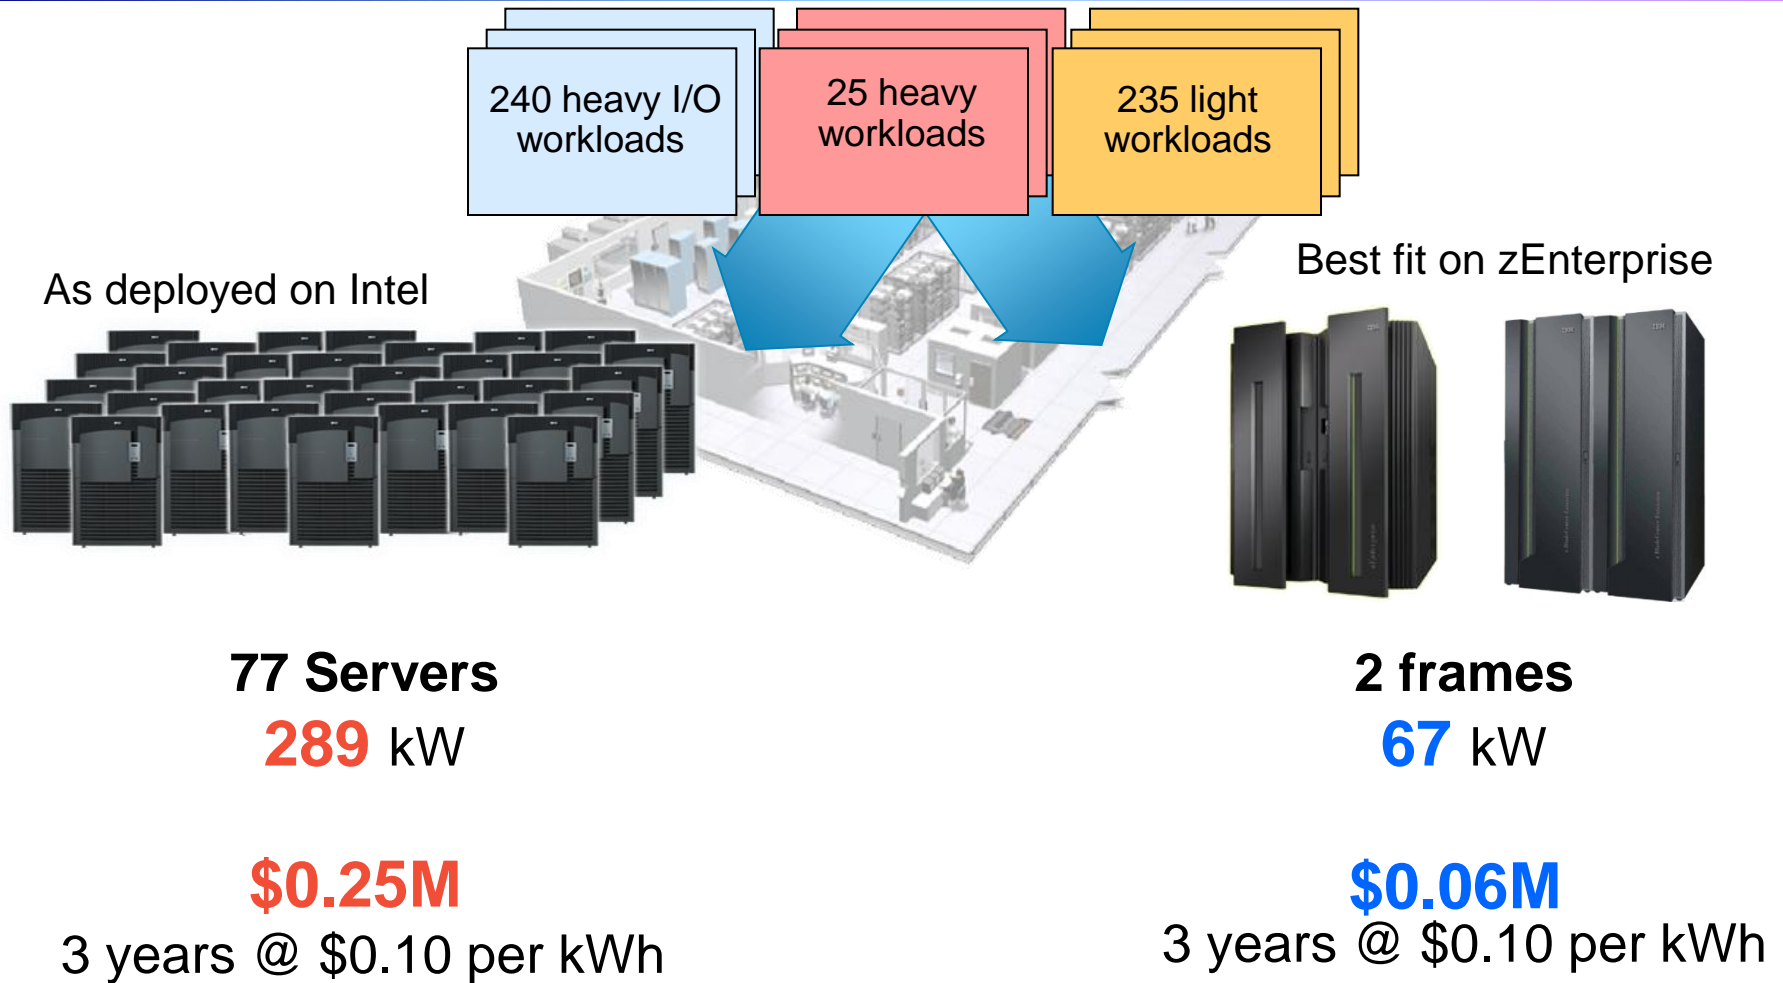
Compare Network Cost Of Acquisition



86%
less


Network configuration is based on IBM internal studies.
Prices are in US currency, prices will vary by country.h

Compare Power Consumption



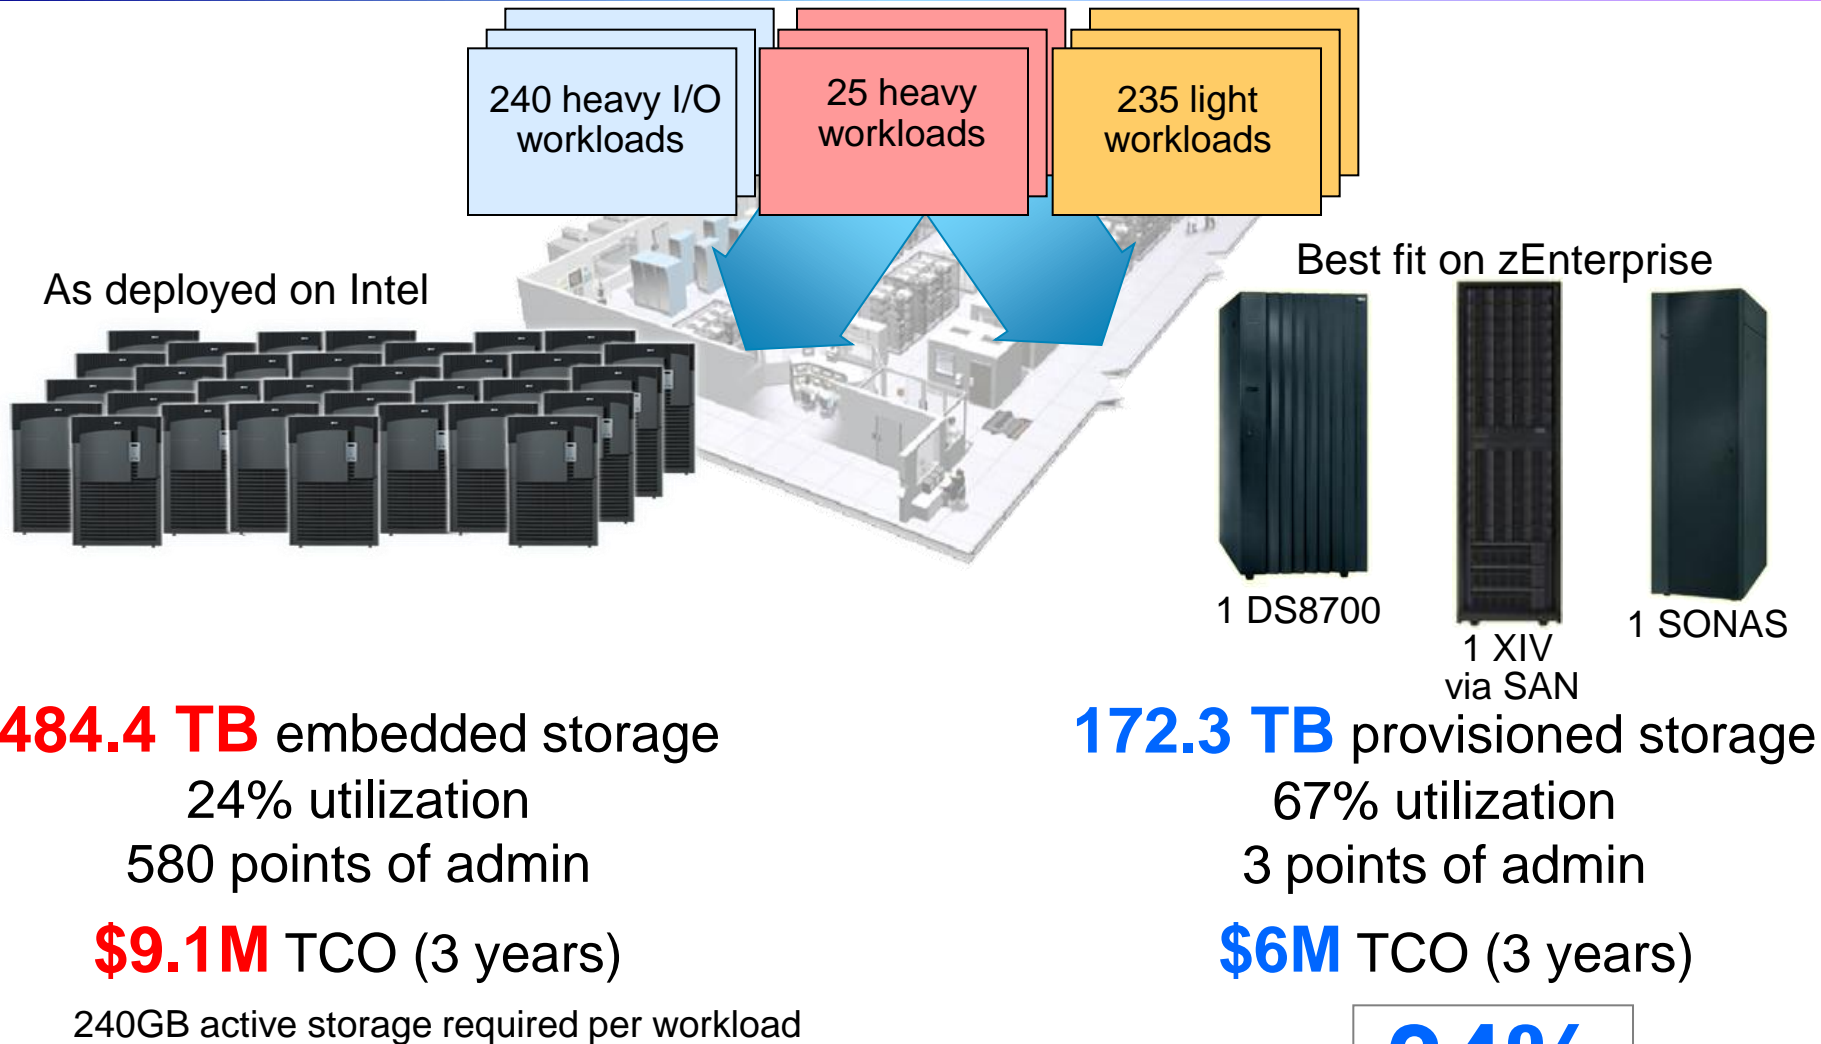
Server configuration based on IBM internal studies. Calculations for Intel servers based on published power ratings and industry standard rates. Prices are in US currency, prices will vary by country.

Compare Server Infrastructure Labor Costs




Configuration based on IBM internal studies. Labor model based on customer provided data from IBM studies. Labor rates will vary by country

Compare Storage Cost



Storage configuration is based on IBM internal studies.
Prices are in US currency, prices will vary by country.

Simplification – Fewer Parts To Assemble And Manage




As deployed on Intel

Best fit on zEnterprise

77	Servers	2 frames
664	Network (parts)	21
289	Power (kW)	67
7.2	Administrators	5.8
580	Storage admin points	3



The Savings Are Cumulative



<i>Three Year Cost of ...</i>	<i>Deploy on Intel</i>	<i>Best fit on zEnterprise</i>
Servers	\$ 15.2M	\$ 7.5M
Network	\$ 0.20M	\$ 0.03M
Power	\$ 0.25M	\$ 0.06M
Labor	\$ 3.45M	\$ 2.80M
Storage	\$ 9.1M	\$ 6.0M
Total	\$ 28.20M	\$ 16.39M
Total cost per workload	\$ 56K	\$ 33K

42%
less

Results may vary based on customer workload profiles/characteristics. Prices based on publicly available US list prices. Prices may vary by country

But what's really so special about zEnterprise?



Customer

Let's take a closer look at the value of zEnterprise...



IBM

zEnterprise Value

- zEnterprise z196 is STILL best for handling **core business workloads**

IBM zEnterprise System

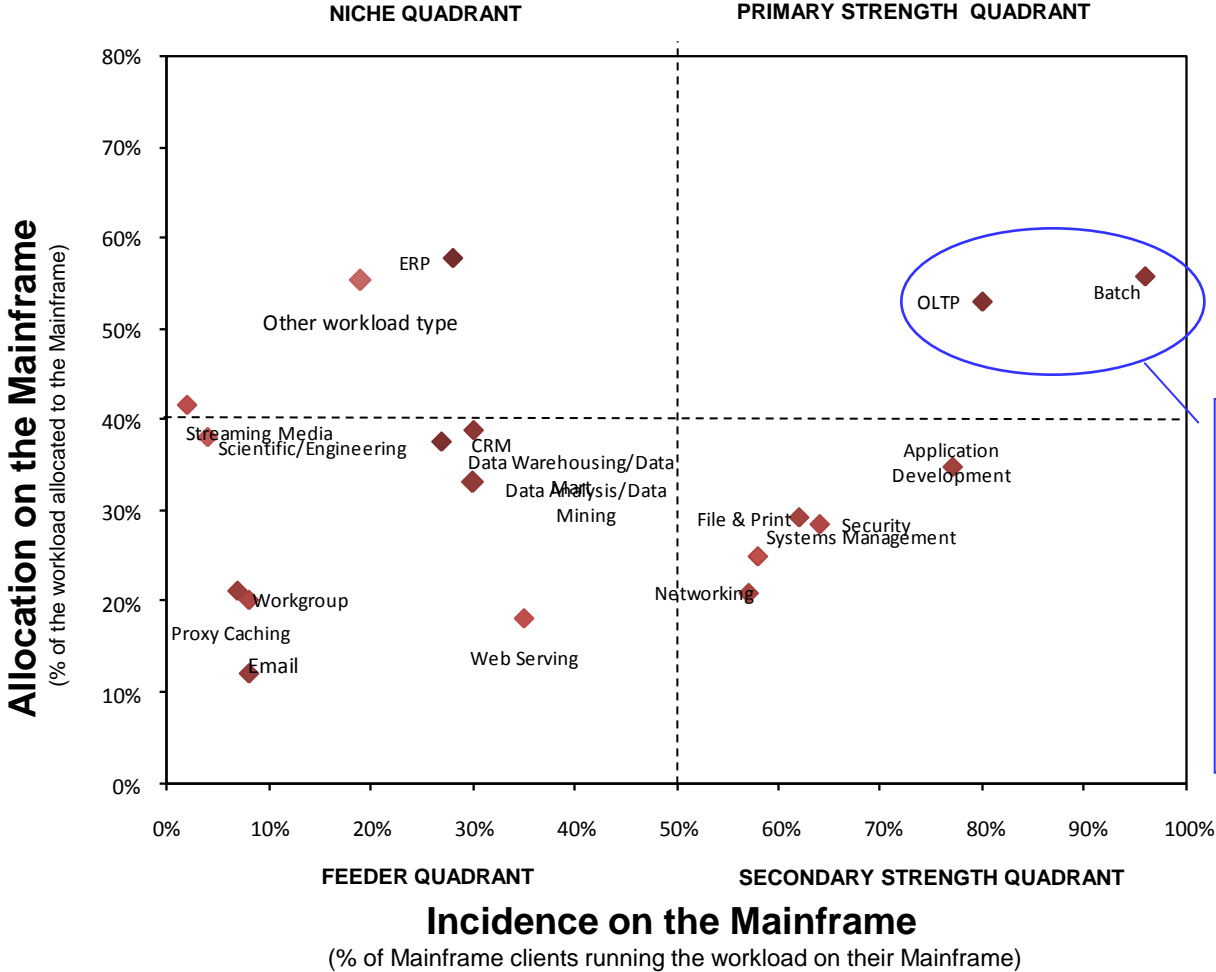


zEnterprise z196

*zEnterprise
BladeCenter
Extension (zBX)*

Batch And OLTP Are Prime Workloads For The Mainframe


Incidence of workload on the Mainframe vs. allocation on the Mainframe



High incidence plus high allocation means OLTP and Batch are core mainframe workloads

System z Platform Easily Handles Workload Peaks

Example: Core banking workloads running simultaneously with high utilization



Customer Data Shows Most Mainframe Workloads Are Already Best Fit

- IBM Eagle Team performs total cost of ownership (TCO) studies for customers
 - Includes detailed cost and risk analysis of mainframe vs. alternative platforms
- With over 200 customers evaluated, Eagle Team has shown **System z offers better TCO** than a distributed alternative... with very few exceptions
- Contact Craig Bender (csbender@us.ibm.com)



Now With zEnterprise z196, System z Is Better Than Ever

**zEnterprise z196
continues a tradition of
mainframe innovation**



System z z10



zEnterprise z196 – A new generation of mainframe

Faster clock speed!

4.4 GHz



5.2 GHz

More processors per MCM!

5



6

More total processors!

77 (64 configurable)



96 (80 configurable)

More memory!

Up to 1.5 TB



Up to 3TB

*More performance!***

920 MIPS



1202 MIPS

*More capacity!**

>30,000 MIPS



>50,000 MIPS

Same power!

1800 W per MCM



1800 W per MCM


• Based on LSPR ratings for fully configured system

** Single process performance

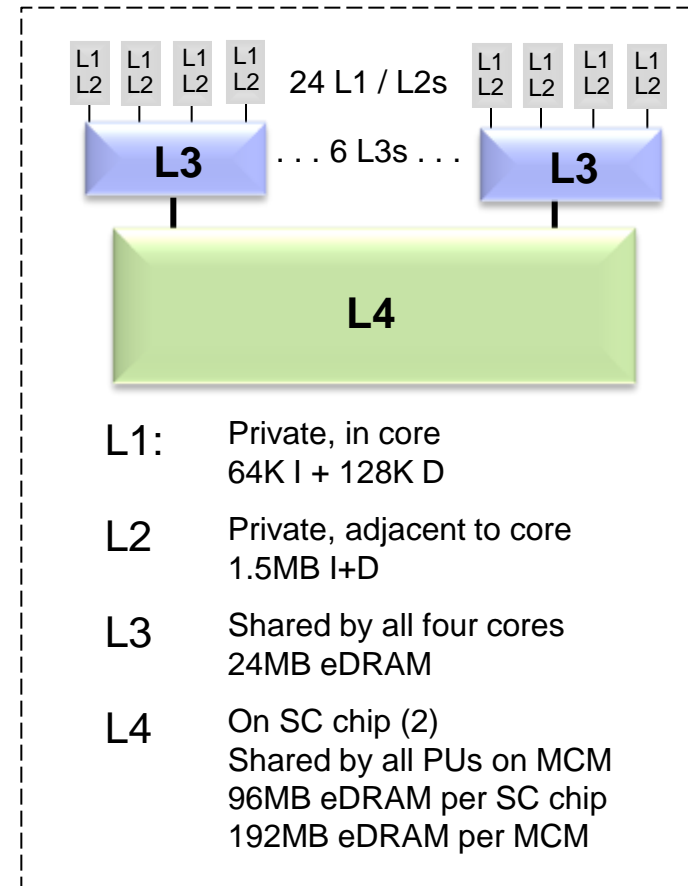
MCM = Multi-chip module

z196 Has Almost 8x More On-Chip Cache As z10 EC

One z10 EC MCM




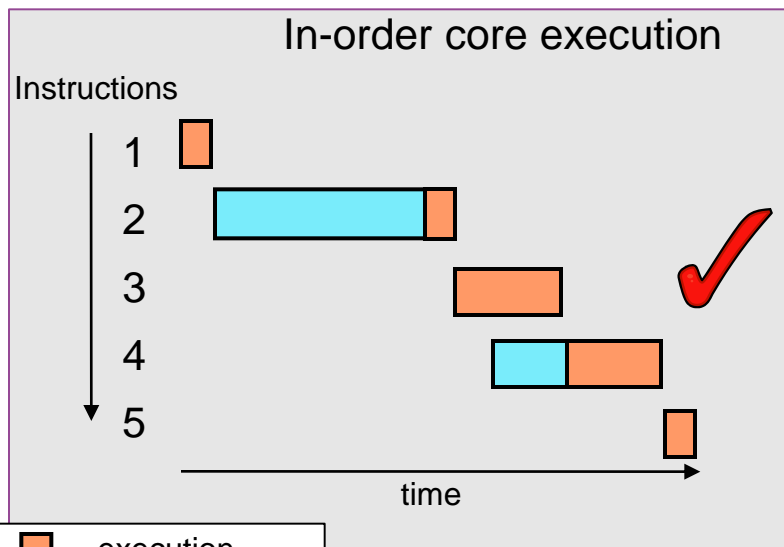
One z196 MCM



More cache means less latency

z196 Adds Out-Of-Order Processing To Its Superscalar Architecture

- Superscalar enhancements to z196:
 - ▶ Decodes up to 3 instructions per cycle (up from 2 on z10)
 - ▶ Executes up to 5 instructions per cycle (up from 2 on z10)
- >100 new instructions added
 - ▶ In particular, Instruction Cracking and Register Renaming which enable Out-of-Order (OOO) instruction execution
- Reduces instruction wait times, and benefits compute-intensive apps



How Does This Add Up?

Multiple Tests Show z196 Outperforms z10 EC

	Performance Ratio (z196 : z10 EC)
LSPR with z/OS V1R11	
z196 708 and z10 708*	1.37
z196 780 and z10 764**	1.64
CPO Banking Benchmark	
CICS – 3270 version	1.37
WAS on z/OS	1.32
WAS on Linux on System z	1.47
CPO COBOL Benchmark	
z/OS V1R11 Enterprise COBOL 4.1	1.41

* Customer average for z10 EC CEC is 9 GP processors

** Each as fully-configured systems

zEnterprise Value

- zEnterprise z196 is STILL best for handling **core business workloads**
- zEnterprise is more than a mainframe – it's a **complete multi-architecture platform**

IBM zEnterprise System




zEnterprise z196

*zEnterprise
BladeCenter
Extension (zBX)*

zEnterprise BladeCenter Extension (zBX) Adds New Platforms To System z

- zBX ordered as one fully built and tested System z “part”
 - ▶ Includes all necessary components – switches, chassis, power, and cabling
 - ▶ Blades and optimizers purchased separately
- Built from standard IBM Certified Components
- Full redundancy insures highest reliability
- System z product support for problem reporting, hardware and firmware updates



Each zBX can support*:

- 14 blades per chassis
- 2 chassis per rack

One fully loaded zBX is:

- 4 racks
- 112 blades

Selected IBM blades supported:

- IBM POWER7 blades
- IBM System x blades*
- Specialty Optimizers

* All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represents goals and objectives only.

Note that only 28 System x blades will be supported per zBX at initial offering.


zBX-supported POWER And System x Application Server Blades

*AIX on
POWER7
Blades*

Workload

AIX

PowerVM



*Linux on
System x
Blades**

Workload

Linux

x86_IH

*Windows on
System x
Blades**

Workload

Windows

x86_IH



■ POWER7 PS701 Express

- ▶ Single-width, 1ch/8co, 3.0 GHz
 - Up to 4 threads per core
- ▶ AIX OS 5.3 or greater
- ▶ PowerVM

■ System x HX5 (Westmere-EX)

- ▶ Single-width, 2ch/16 co, 2.13 GHz
 - Up to 2 threads per core
- ▶ Windows and Linux
- ▶ Integrated hypervisor

* All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represents goals and objectives only.

Blades Run Distributed Software Available Through Passport Advantage

IBM Information Management software

Cognos BI
Content Manager
DB2 UDB
Document Manager
Filenet
Informix
Information Integration
Information Server
InfoSphere
MDM
OmniFind
OpenPages
Optim
SPSS
...

Lotus software

Connections
Domino
Forms
ActiveInsight
Quickr
Web Content Manager
Workflow
Mashup Center
Sametime
...

- No MIPS or MSU rating for zBX software

WebSphere software

Application Workload Modeler
Communications Server
Decision Server
MQSeries
Process Integration Server
Application Server
BI Server
Business Integration
Commerce
ESB
Lombardi
Portal
Portlet Factory
Translation Server
Voice Server
...

Rational software

Team Concert
Requirements Composer
Asset Manager
BuildForge
ClearCase
AppScan
Quality Manager
Functional Test
Performance Test
...

Other

Unica
Systems Director
Sterling
...

Tivoli software

Directory Server
Maximo
Performance Analyzer
Composite Application Manager
Identity and Access Assurance
Access manager
Asset Manager
Change and Configuration Manager
Compliance Insight Manager
Directory Integrator
Federated Identity Manager
Identity and Access Manager
License Compliance Manager
Monitoring
Netcool
OMEGAMON
Provisioning
Security Compliance Manager
Service Automation Manager
Systems Automation
Workload Scheduler
...

zBX Optimizers Are Built-For-Purpose

- Delivered as Blades for use in zBX
- Fully-integrated, fully-contained – each targeted for specific workload functions
 - ▶ Pre-packaged, self-contained units including hardware, software, memory, etc.
- Designed for integration with and management by zEnterprise
- Two zBX optimizers available today:
 - ▶ **IBM Smart Analytics Optimizer**
 - ▶ **IBM WebSphere DataPower XI50 for zEnterprise**

Optimizers



But what is so unique about putting a BladeCenter next to a mainframe?



Customer

There's more to this than meets the eye!

The Unified Resource Manager – also called zManager – is the “secret sauce”.

It provides extensive management of resources and workloads across all zEnterprise platforms!



IBM

zManager Provides Platform And Resource Management Across zEnterprise Environments

Process	Typical Distributed Management Practices	zManager
Asset Management	<ul style="list-style-type: none"> Discover assets with ad hoc methods Manual entitlement management 	<ul style="list-style-type: none"> Automated discovery and management of entitlement assets
Deployment Management	<ul style="list-style-type: none"> Manually configure hypervisor and build networks 	<ul style="list-style-type: none"> Automated deployment of hypervisor and attachment to integrated networks
Security Management	<ul style="list-style-type: none"> Different ways to manage administrator access 	<ul style="list-style-type: none"> Centralized, fine-grained administrator access management
Change Management	<ul style="list-style-type: none"> No visibility into impact of changes 	<ul style="list-style-type: none"> Track dependencies for change impact
Capacity and Performance Management	<ul style="list-style-type: none"> No end-to-end transaction monitoring Manually adjust CPU resources to meet changing workload demands 	<ul style="list-style-type: none"> End-to-end transaction monitoring to isolate issues Automatic CPU resource adjustments to meet changing workload demands

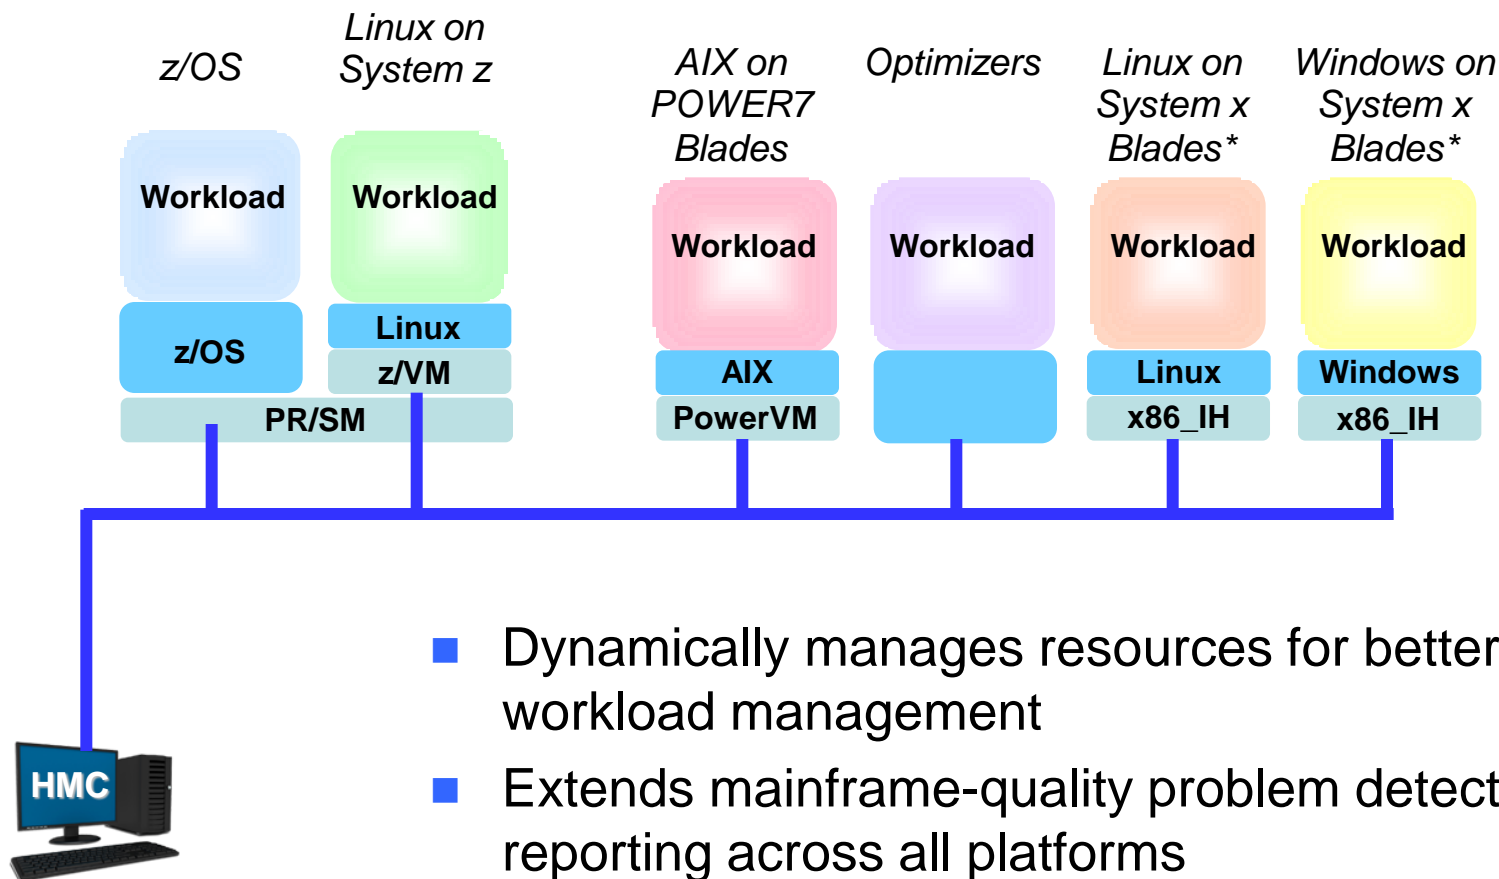
zEnterprise Network Simplification And Security

- “Network in a box” limits vulnerability to security breaches
 - ▶ Fully integrated concealed networks
 - No external switches or routers necessary – IBM-only equipment
 - Fully tested, pre-installed and pre-configured
 - ▶ Can reduce latency and the number of “hops”

- Security
 - ▶ Management Network:
 - Tightly restricted to zManager use only
 - ▶ Data Network:
 - Accessible only by authorized virtual machines
 - ▶ Logical security via virtualization
 - ▶ zManager includes strict “role-based” access control
 - ▶ No need for additional encryption or firewall

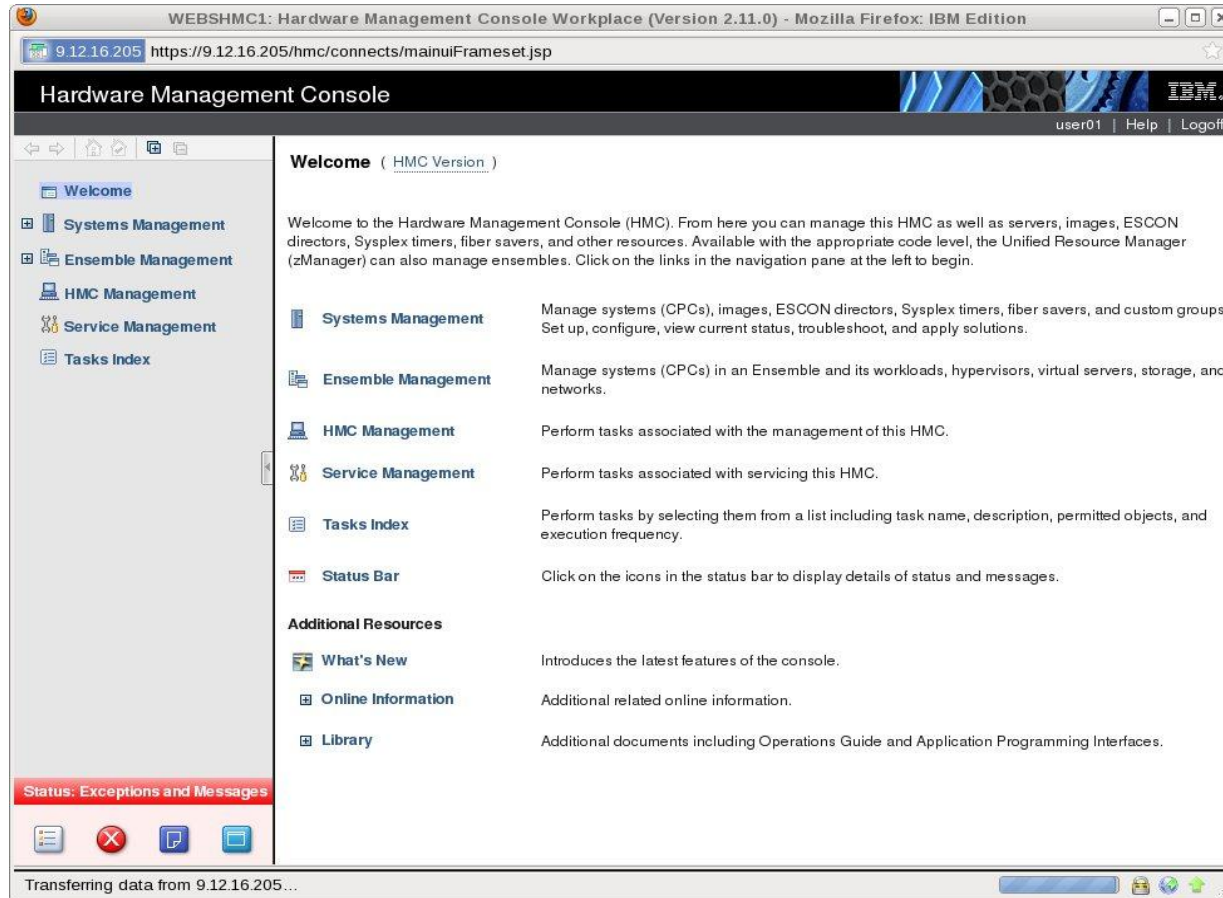


zManager Owns The Private Management Network For Hypervisor Communications



- Dynamically manages resources for better workload management
- Extends mainframe-quality problem detection and reporting across all platforms
- Monitors system-wide energy efficiency

DEMO: Manage Resources And Workloads Using zManager



- Hardware, workload and ensemble management

zEnterprise Value

- zEnterprise z196 is STILL best for handling **core business workloads**
- zEnterprise is more than a mainframe – it's a **complete multi-architecture platform**
- zEnterprise continues a tradition of **unmatched reliability** and **superior qualities of service**

IBM zEnterprise System



zEnterprise z196

*zEnterprise
BladeCenter
Extension (zBX)*

A Complex, Distributed-based Scale Out Strategy Has Its Risks

North America	Europe	Asia Pacific		Apr 26	Apr 25	Apr 24	Apr 23	Apr 22	Apr 21	Apr 20
Amazon CloudFront	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Amazon CloudWatch (N. California)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Amazon CloudWatch (N. Virginia)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Amazon EC2 (N. California)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Amazon EC2 (N. Virginia)	✓	✓	⚠	✖	✖	✖	✓	✓	✓	✓
Amazon EMR (N. California)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Amazon EMR (N. Virginia)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Amazon Flexible Payments Service	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Amazon Mechanical Turk (Requester)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Amazon Mechanical Turk (Worker)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Amazon RDS (N. California)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Amazon RDS (N. Virginia)	✓	✓	⚠	✖	✖	✖	✓	✓	✓	✓
Amazon Route 53	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Amazon Simple Email Service (N. Virginia)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Amazon SNS (N. California)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Amazon SNS (N. Virginia)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Amazon SQS (N. California)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Amazon SQS (N. Virginia)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Amazon S3 (N. California)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Amazon S3 (US Standard)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Amazon SimpleDB (N. California)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Amazon SimpleDB (N. Virginia)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Amazon VPC (N. Virginia)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Amazon public cloud platform suffered a 3+ day outage in April, 2011

- Distributed architecture designed “for durability and availability”
- Yet a complex *single point of failure* negated the advantage of rapid replacement of failed resources
- Numerous customers suffered significant and unrecoverable data loss

Caveat Emptor!

reddit is down.



System z Can Support Unprecedented Levels Of Availability

- Parallel Sysplex architecture designed for **99.999%** availability
 - ▶ Full redundancy yielding no single points of failure
 - All systems can have concurrent access to all critical applications and data
 - Automatic restart and recovery capabilities
 - ▶ Dynamic workload routing via z/OS Workload Manager and Sysplex Distributor
 - Work flow designed for best response times



zEnterprise Continues The Strategy Of Constant Improvements In Availability

- RAIM Memory
 - ▶ Provides more redundancy to protect against additional failure modes
 - Protects DIMM level components such as ASIC, power regulators, clock, and board
 - Protects memory channel failures such as signal lines, control lines, and drivers/receivers on the MCM
 - More robust than ECC, and more cost effective than 100% memory mirroring
 - No performance penalty
- Hot pluggable I/O drawer technology reduces planned down time
 - ▶ Perform maintenance while the system keeps running

The Data Center Of The Future Has Arrived

- World's first multi-architected platform
- A unique platform for workload optimization
- Lowest cost per workload
- The next generation of System z value



