

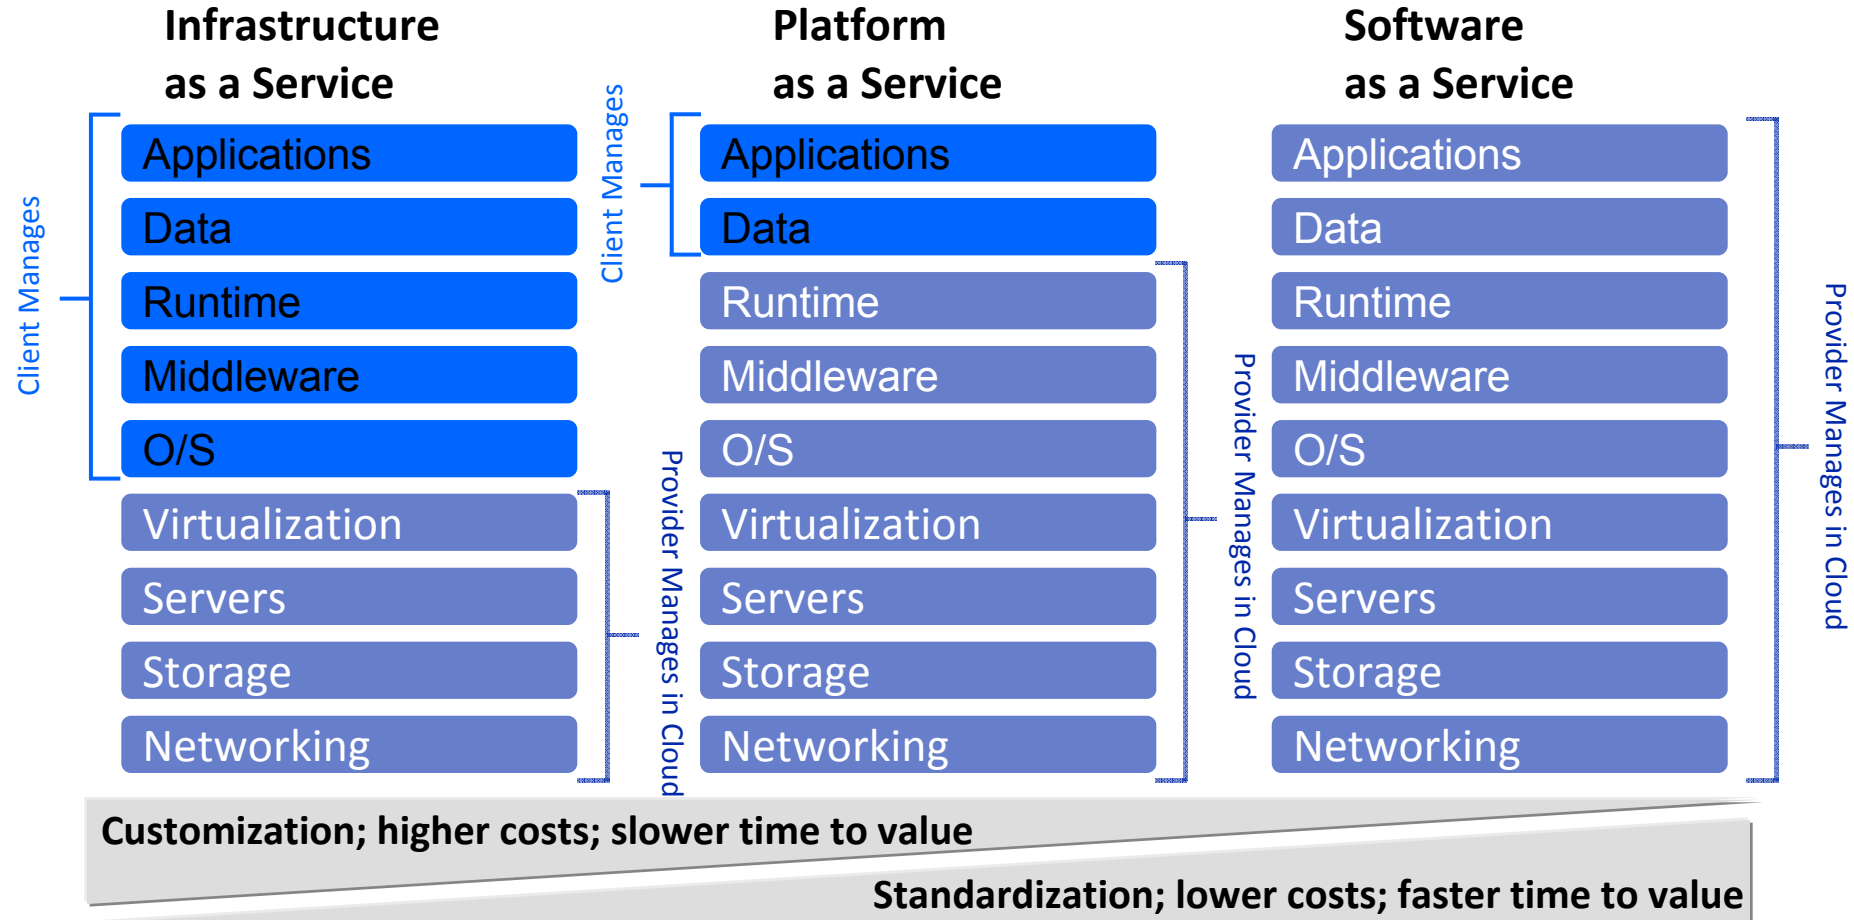
Smarter Computing:
What's Next. Ready Now.

A Step by Step Approach to Cloud

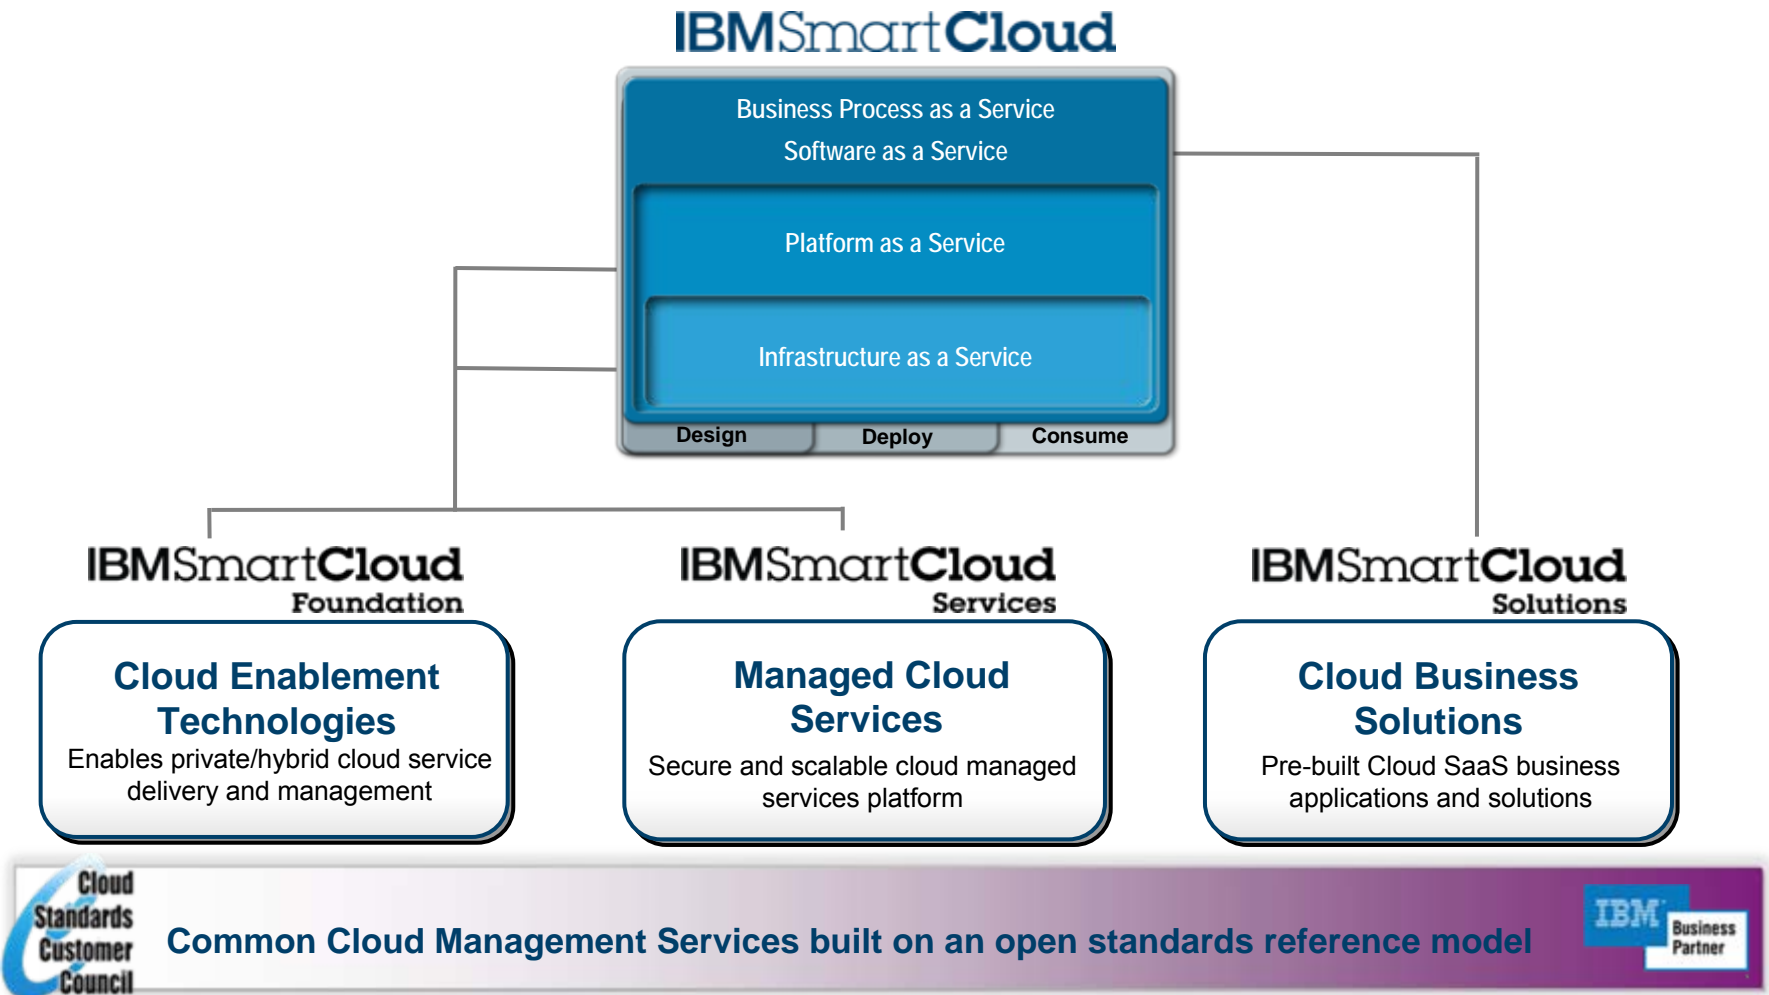
Smarter Computing Briefing 2013



Some Basics - Cloud Service Models



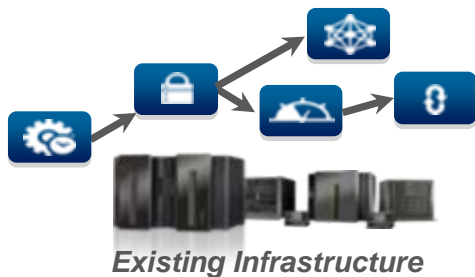
IBMSmartCloud – Comprehensive Cloud Capabilities From IBM



A Practical Cloud Strategy

- Customers already have existing investments in images
 - ▶ Images typically on multiple platforms (x, p, z)
- A practical cloud strategy should try to preserve and re-use these assets
 - ▶ Migrating architectures is often costly and risky
- Instead of artificial and forced migrations, look at how to re-use existing image investments in a cloud model
- Focus on improving price/performance and improving labor

A Practical Strategy With IBM SmartCloud



Hybrid Cloud

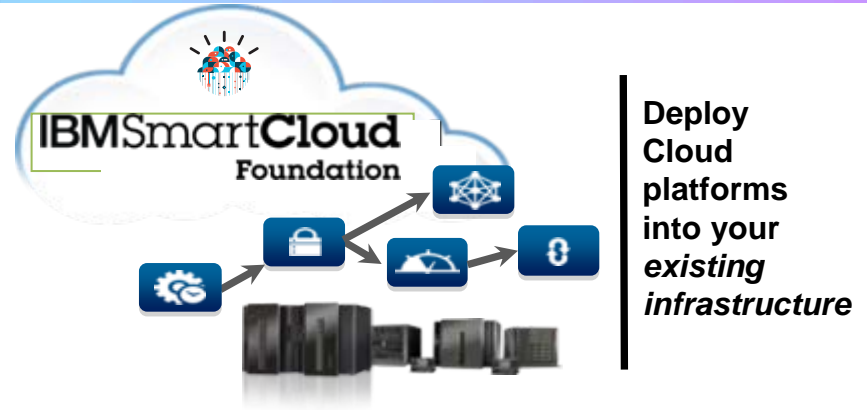
1. Adapt existing infrastructure to a private cloud model

2. Access managed cloud services as needed

3. Leverage consistent technologies in a hybrid model



SmartCloud Foundation



- Adapt existing infrastructures into a Private Cloud environment
- Supports heterogeneous server, networking, storage & middleware
- Capabilities designed to interoperate based on a common open standards Cloud reference model

Platform as a Service Technologies



Lifecycle



Resources



Environments



Management



Integration

Infrastructure as a Service Technologies



Infrastructure



Management



Performance



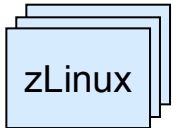
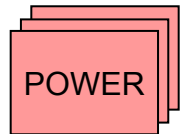
Security



Usage

A Step By Step Approach With Defined Benefits

Existing workloads



Simplification
Standardization
Optimized
Price/Performance

**Integrated Systems
for
IaaS**

Labor efficiency
Agility

**Integrated Systems
for
PaaS**

Unified management
across hosting platforms

**Unified Private
And Public Cloud**

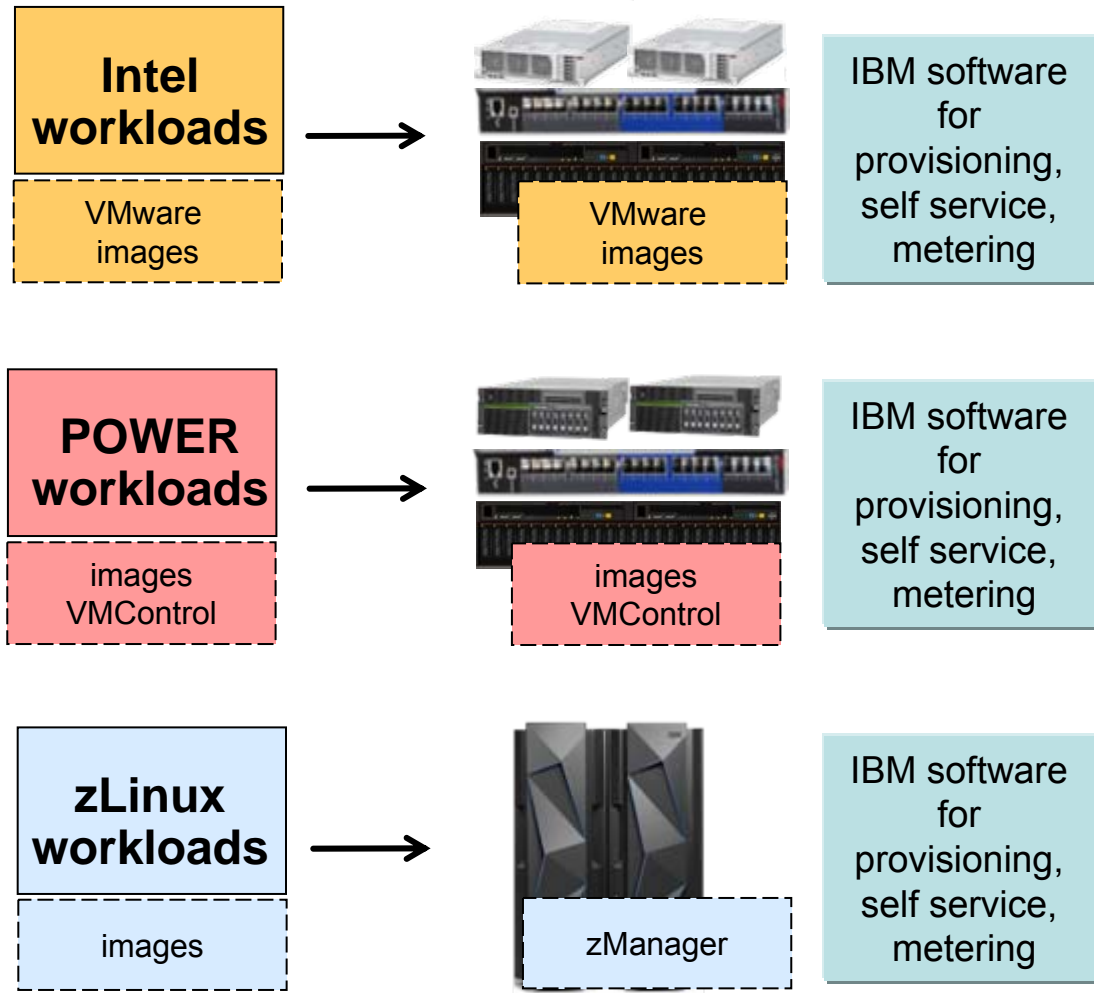


Existing Infrastructure

IBM Provides A Continuum Of Options To Build A Private Cloud

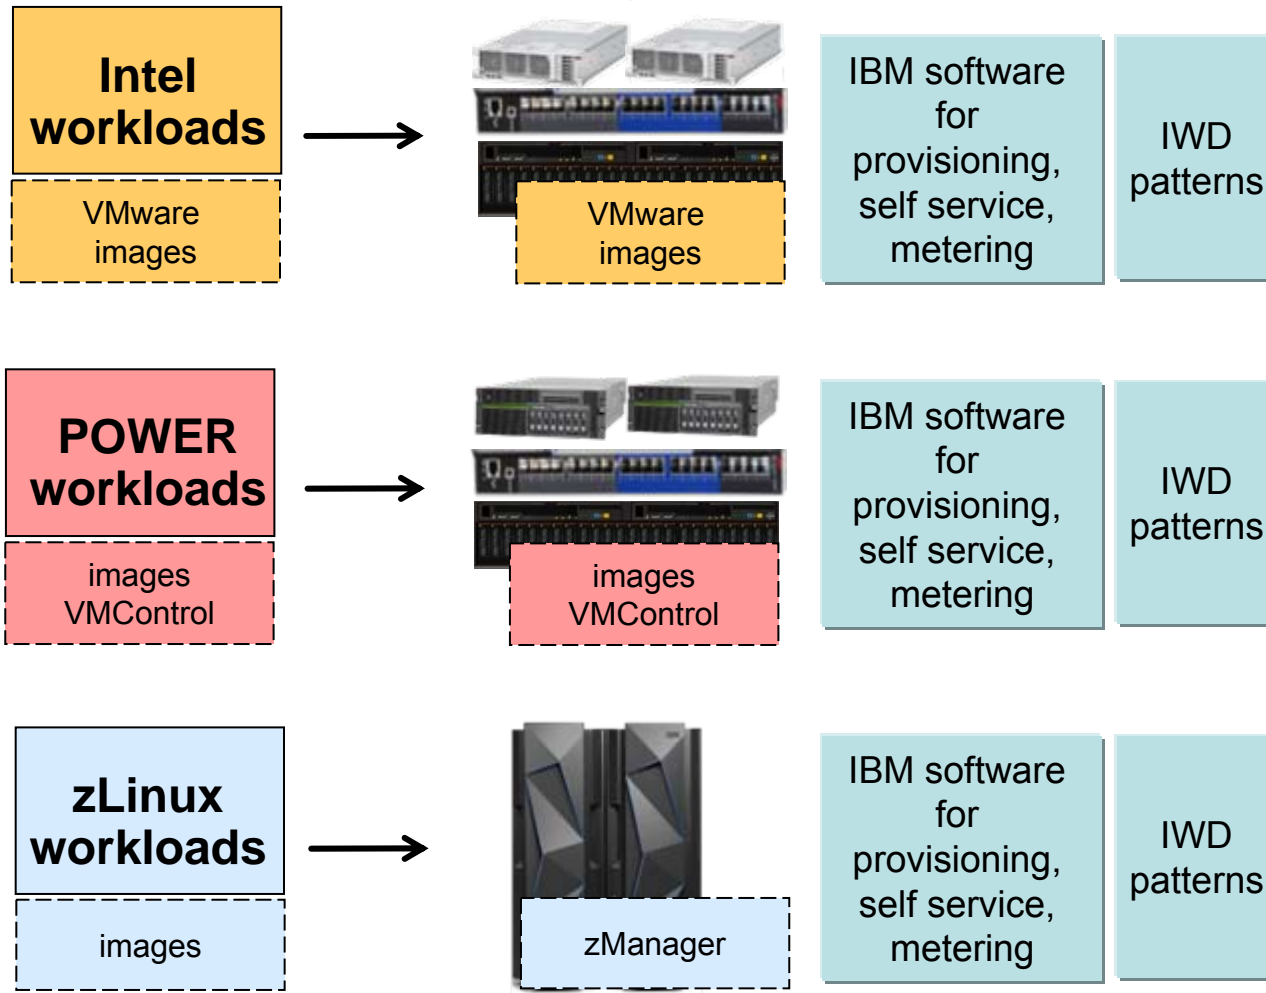
- Move up and down this continuum with a consistent set of technologies
 - ▶ From do-it-yourself to pre-integrated built-in expertise
 - ▶ Preserve existing image investments
- Multi-platform architecture capabilities along the continuum
 - ▶ Avoid costly cross-platform migrations
 - ▶ Maintain “Best Fit” for
 - Architecture
 - Workload Optimization
 - Qualities of Service
- IBM is working to deliver common management capabilities across all hosting platforms

Do-It-Yourself IaaS



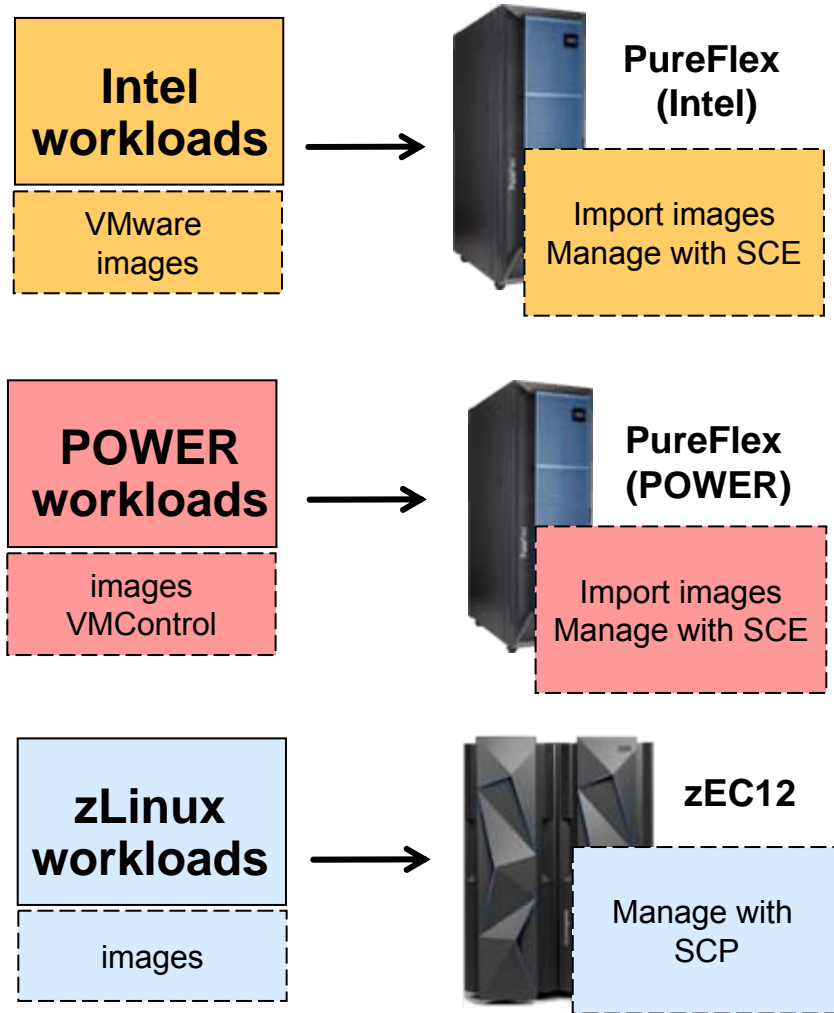
**Existing workloads
and platforms**

Do-It-Yourself PaaS



Existing workloads
and platforms

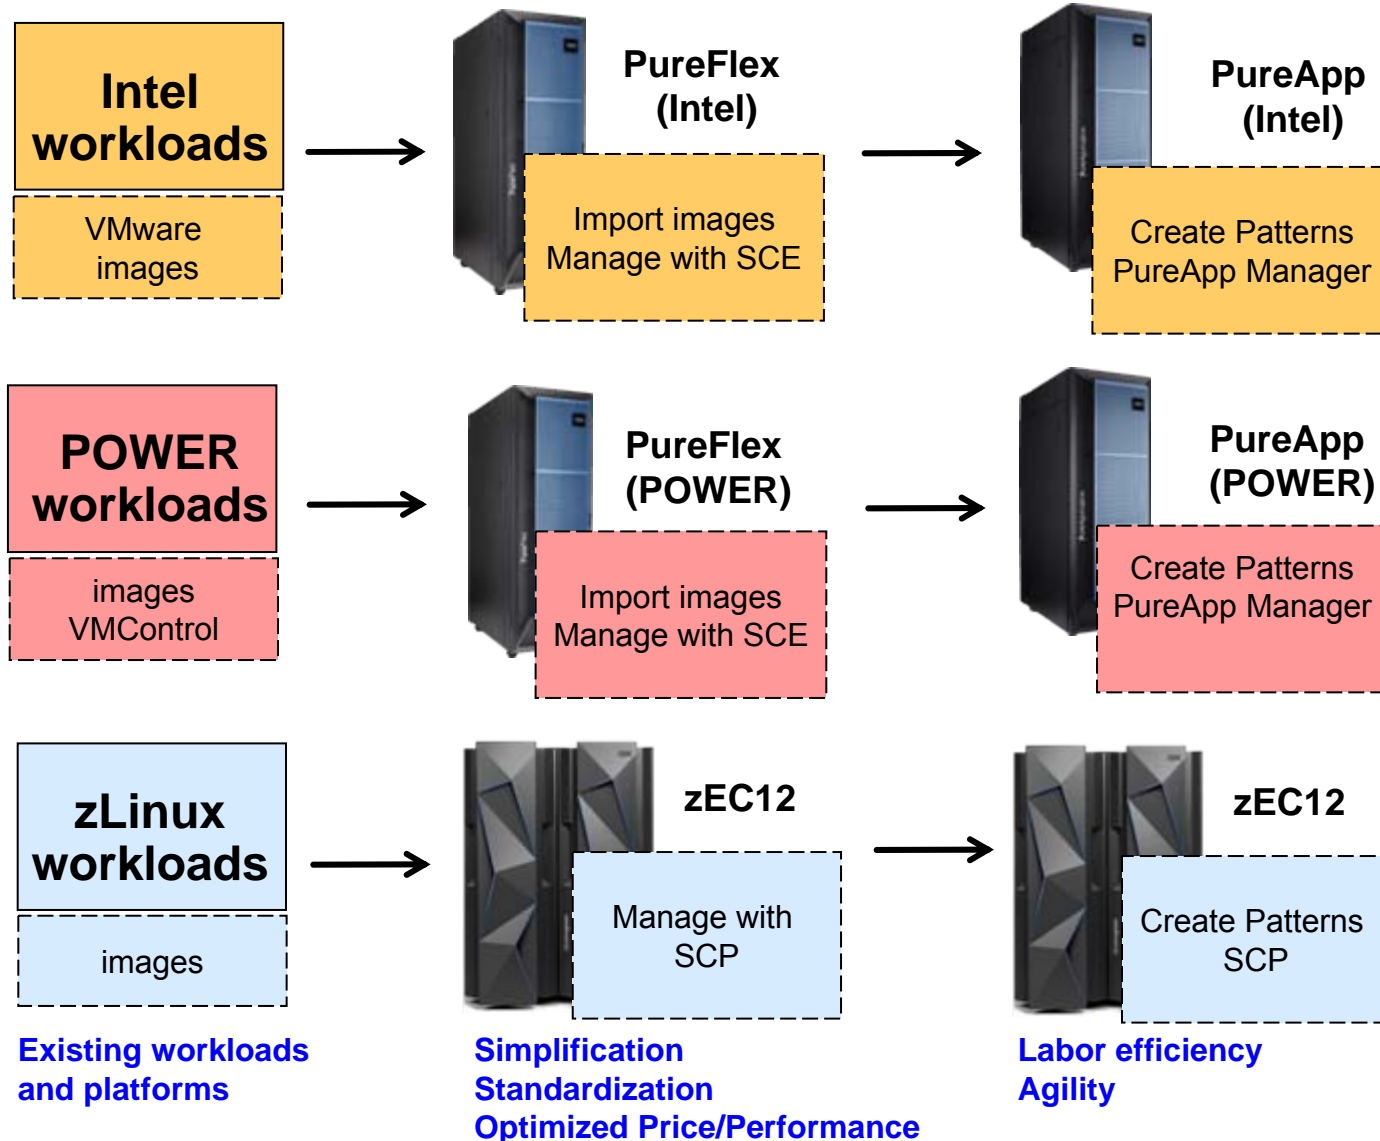
Move To Integrated Systems For IaaS



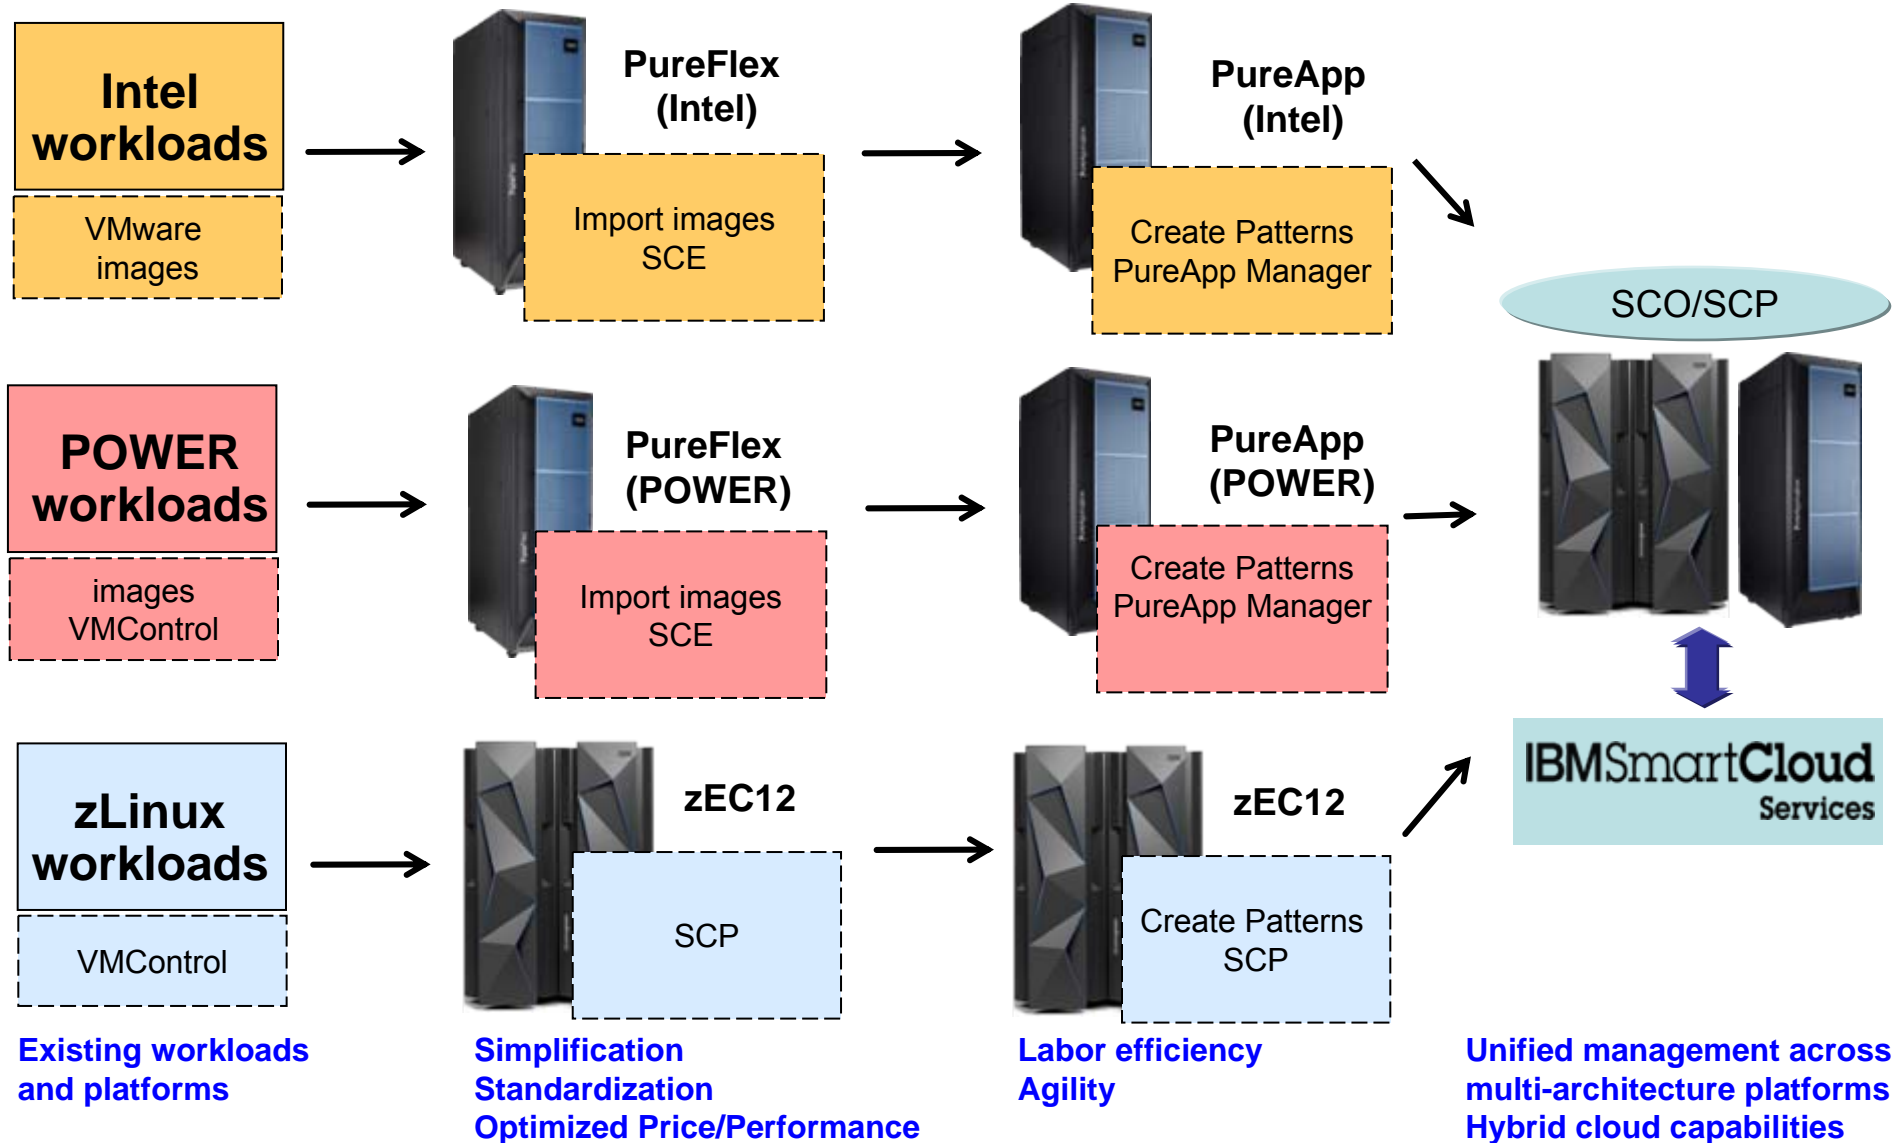
Existing workloads
and platforms

Simplification
Standardization
Optimized Price/Performance

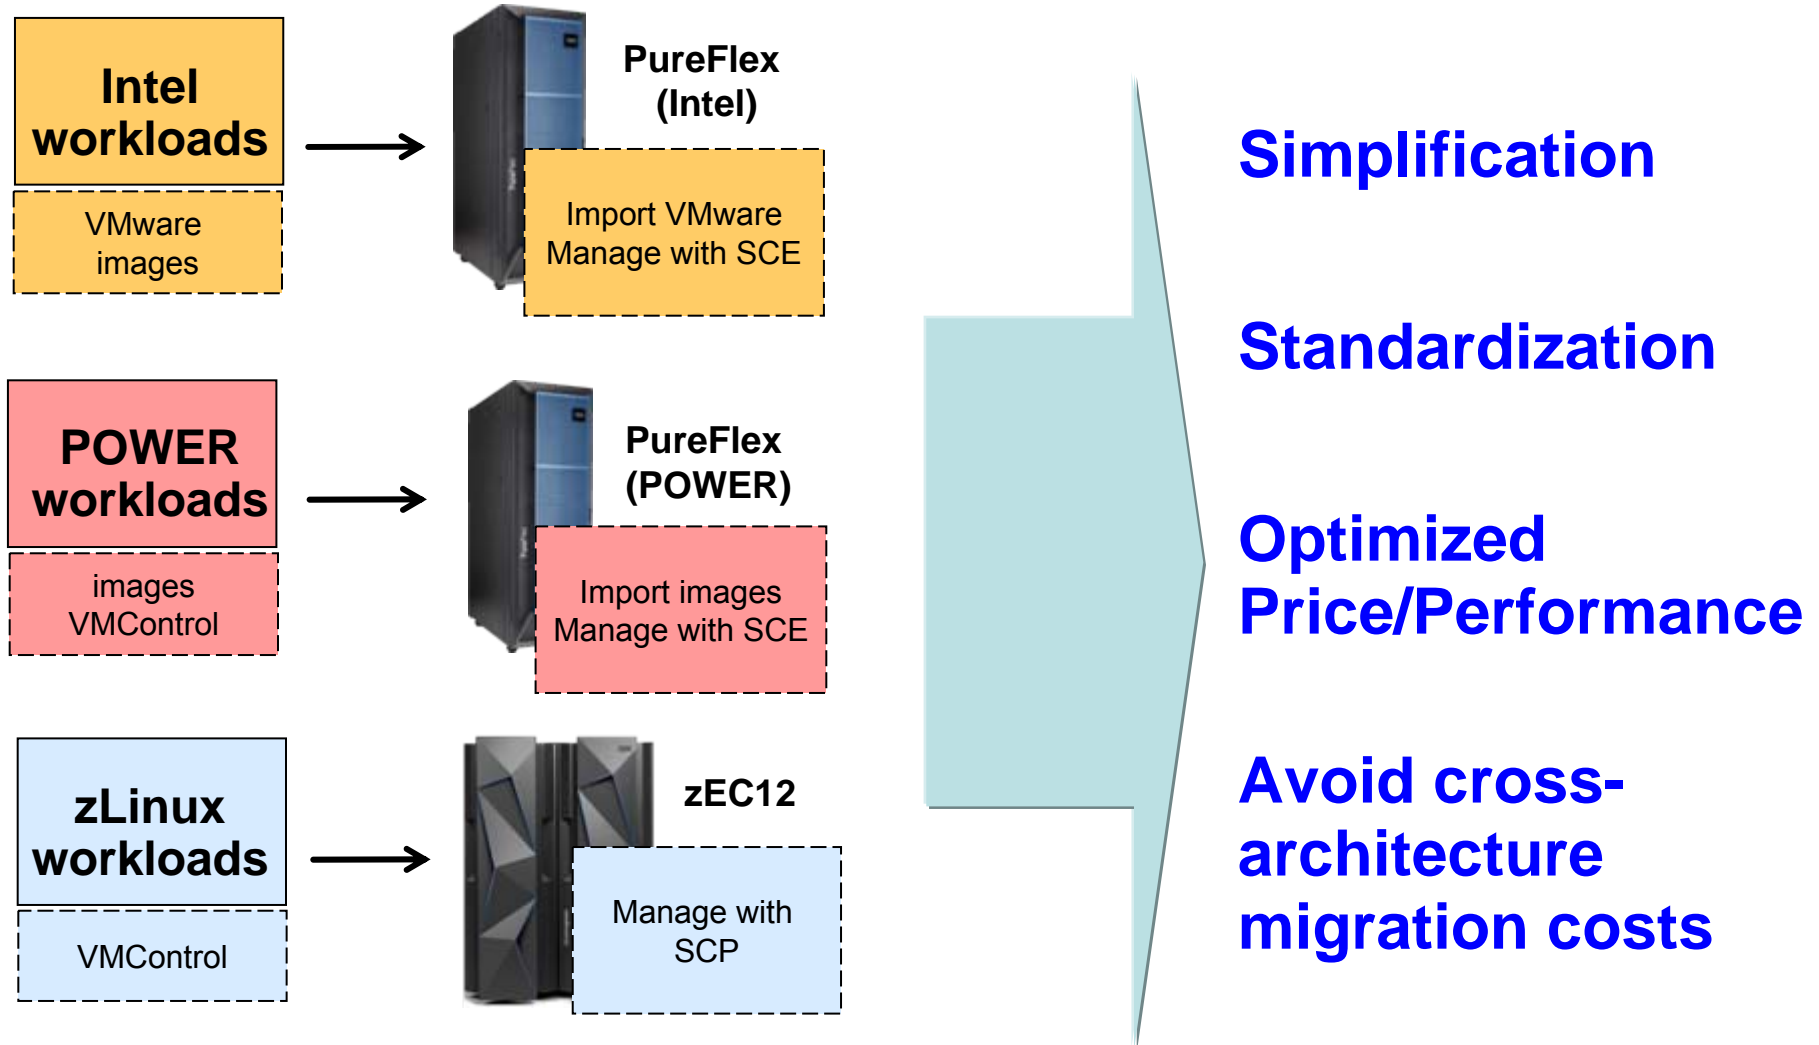
Move To Integrated Systems For PaaS



Move To A Unified Private And Public Cloud



Integrated Systems For IaaS – Benefits



Existing workloads
and platforms

PureFlex Delivers A **Simplified Experience** By Integrating Various Components

Building Blocks: IBM Flex System™ components

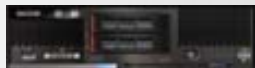
Chassis

14 half-wide
bays for nodes



Compute Nodes

Power 2S/4S
x86 2S/4S



Storage Node

V7000



Management Appliance



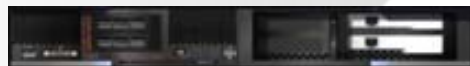
Networking

10/40GbE,
FCoE, IB
8/16Gb FC



Expansion

PCIe
Storage



Build to Order

PureFlex System

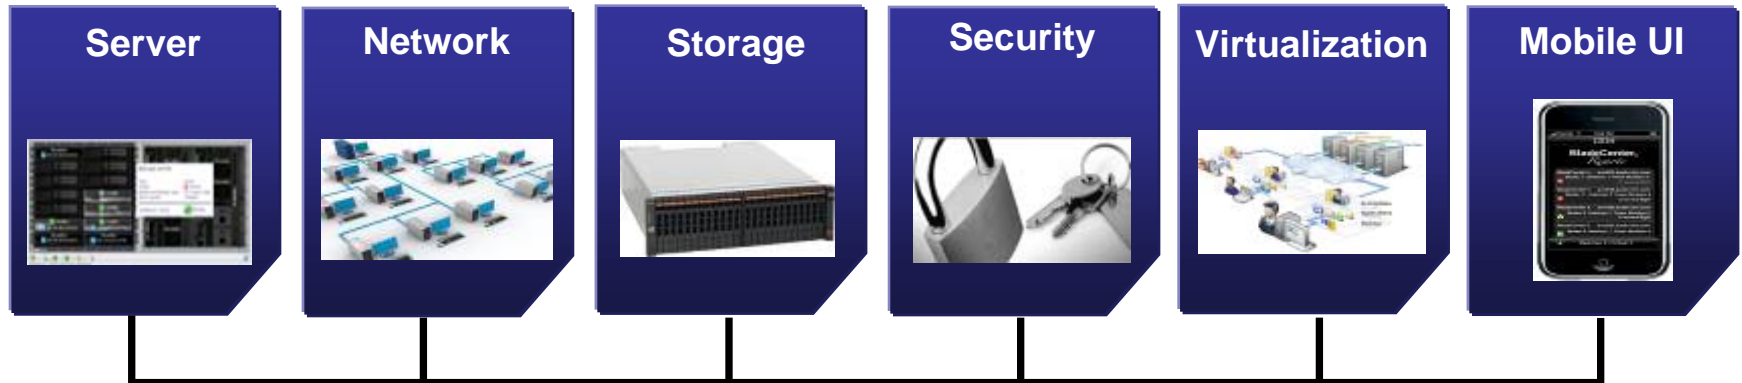


- Flexible choice
- Integrated design
- Pre-assembled hardware
- On-site set up services
- Faster time to value

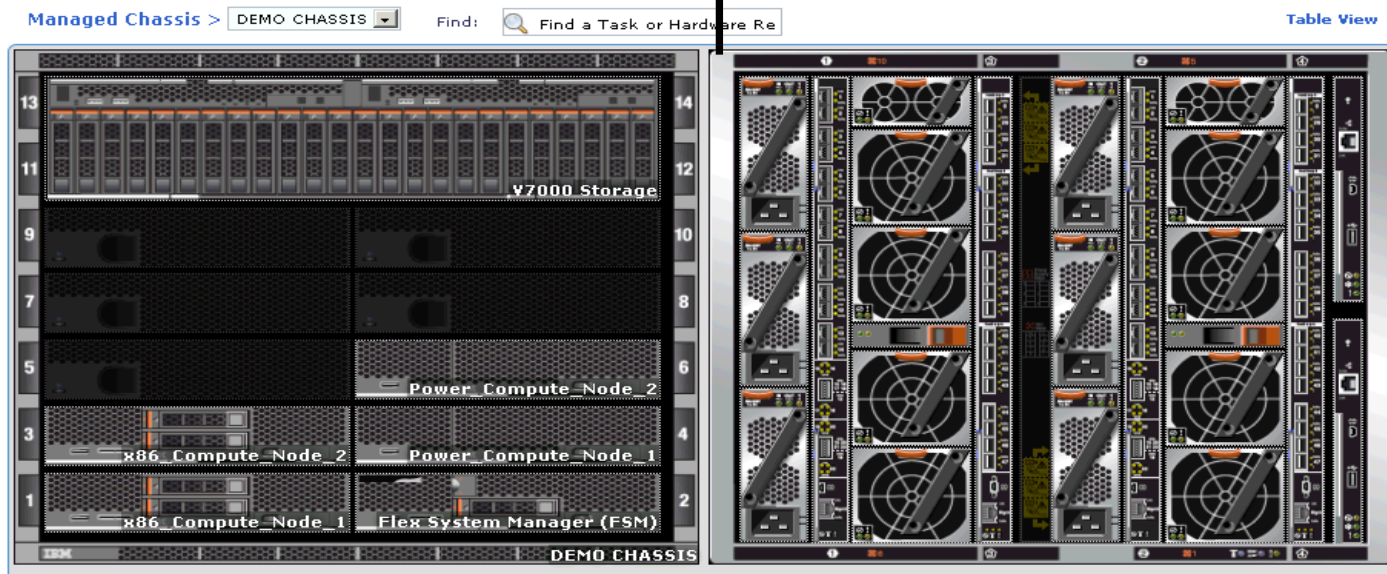
**Express, Standard and
Enterprise Configurations**

DEMO: Management Is Now A Simplified Experience

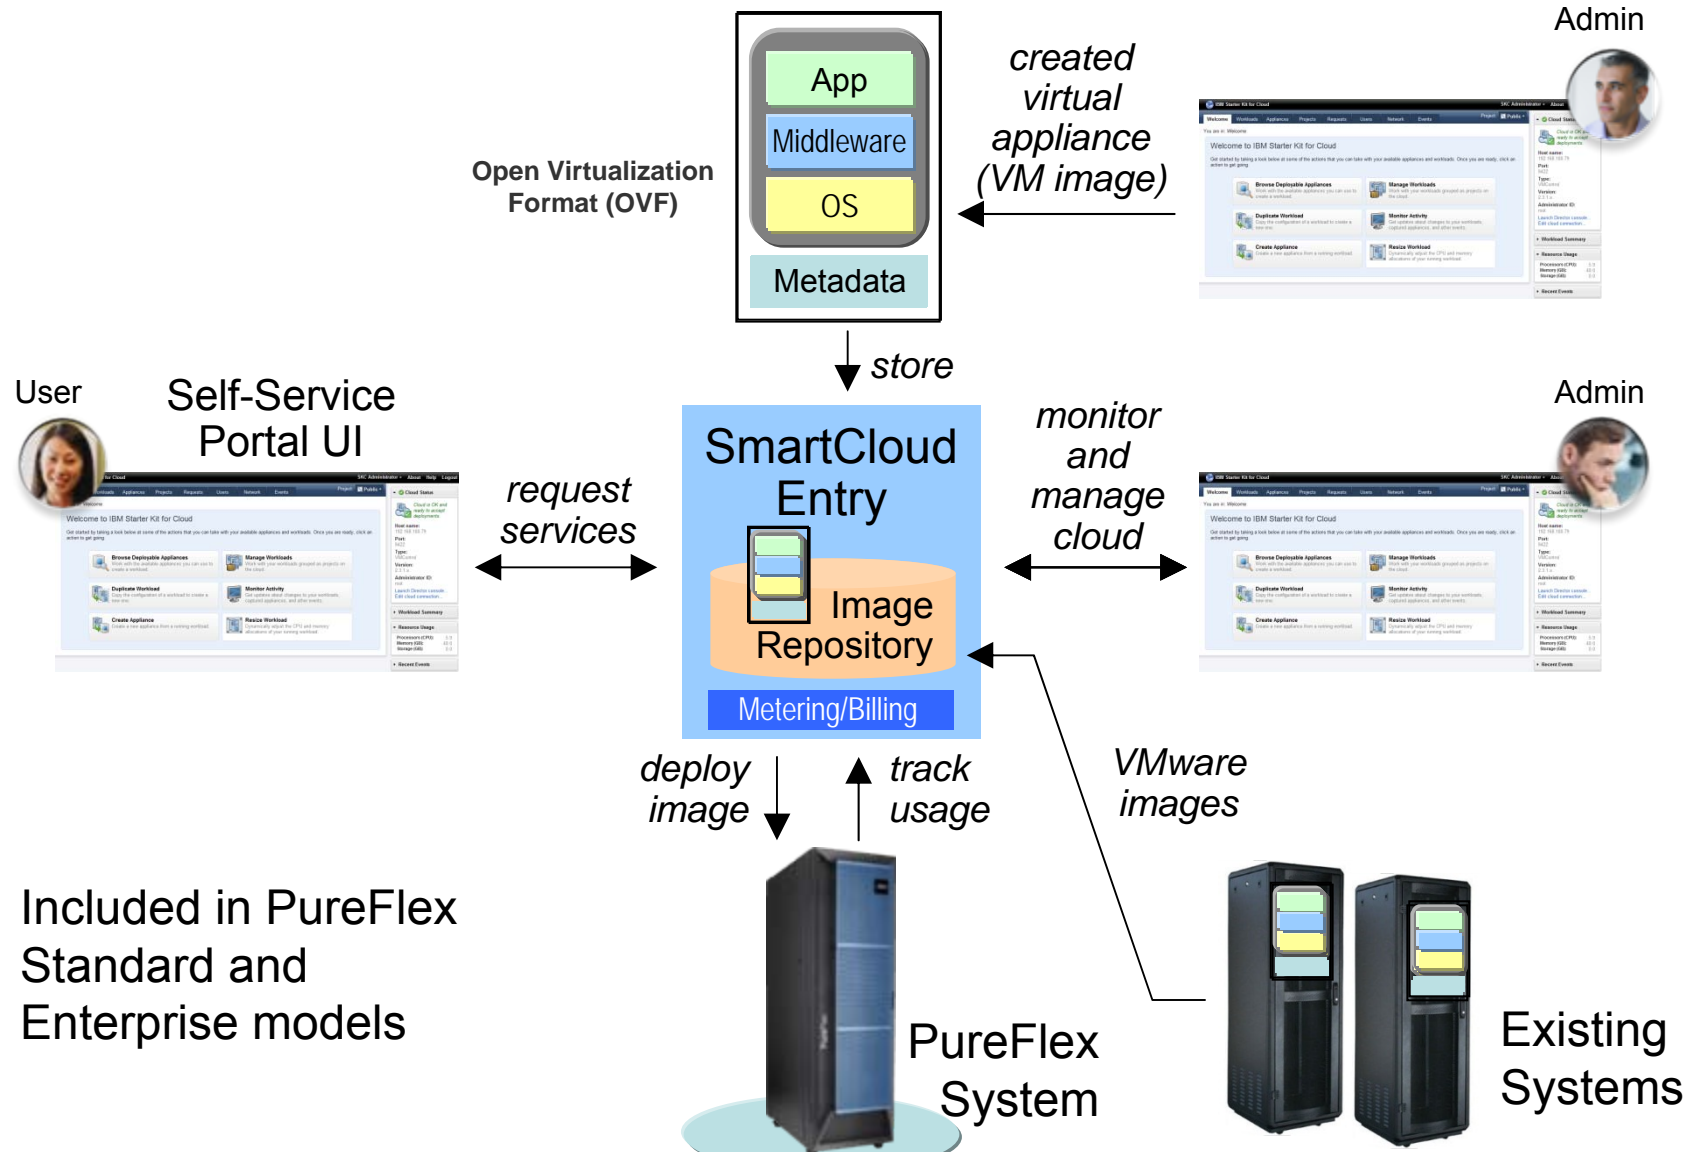
“Single Pane of Glass” for System Management



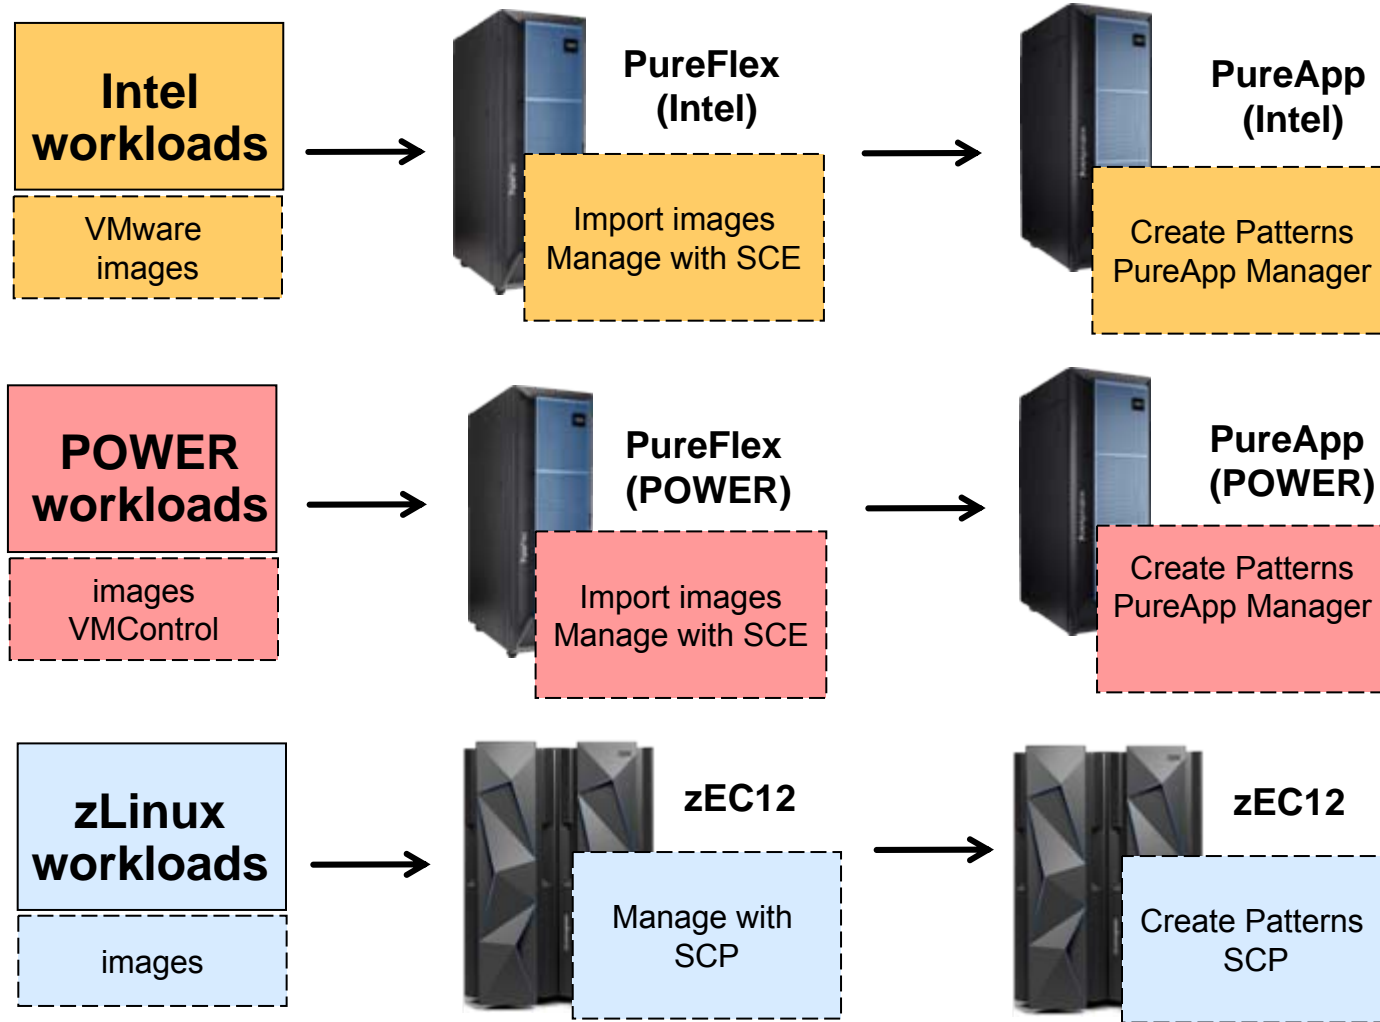
Flex System Manager



PureFlex Includes SmartCloud Entry For Basic Cloud Management



Integrated Systems For PaaS – Benefits



More Labor Efficiency with pattern deployments

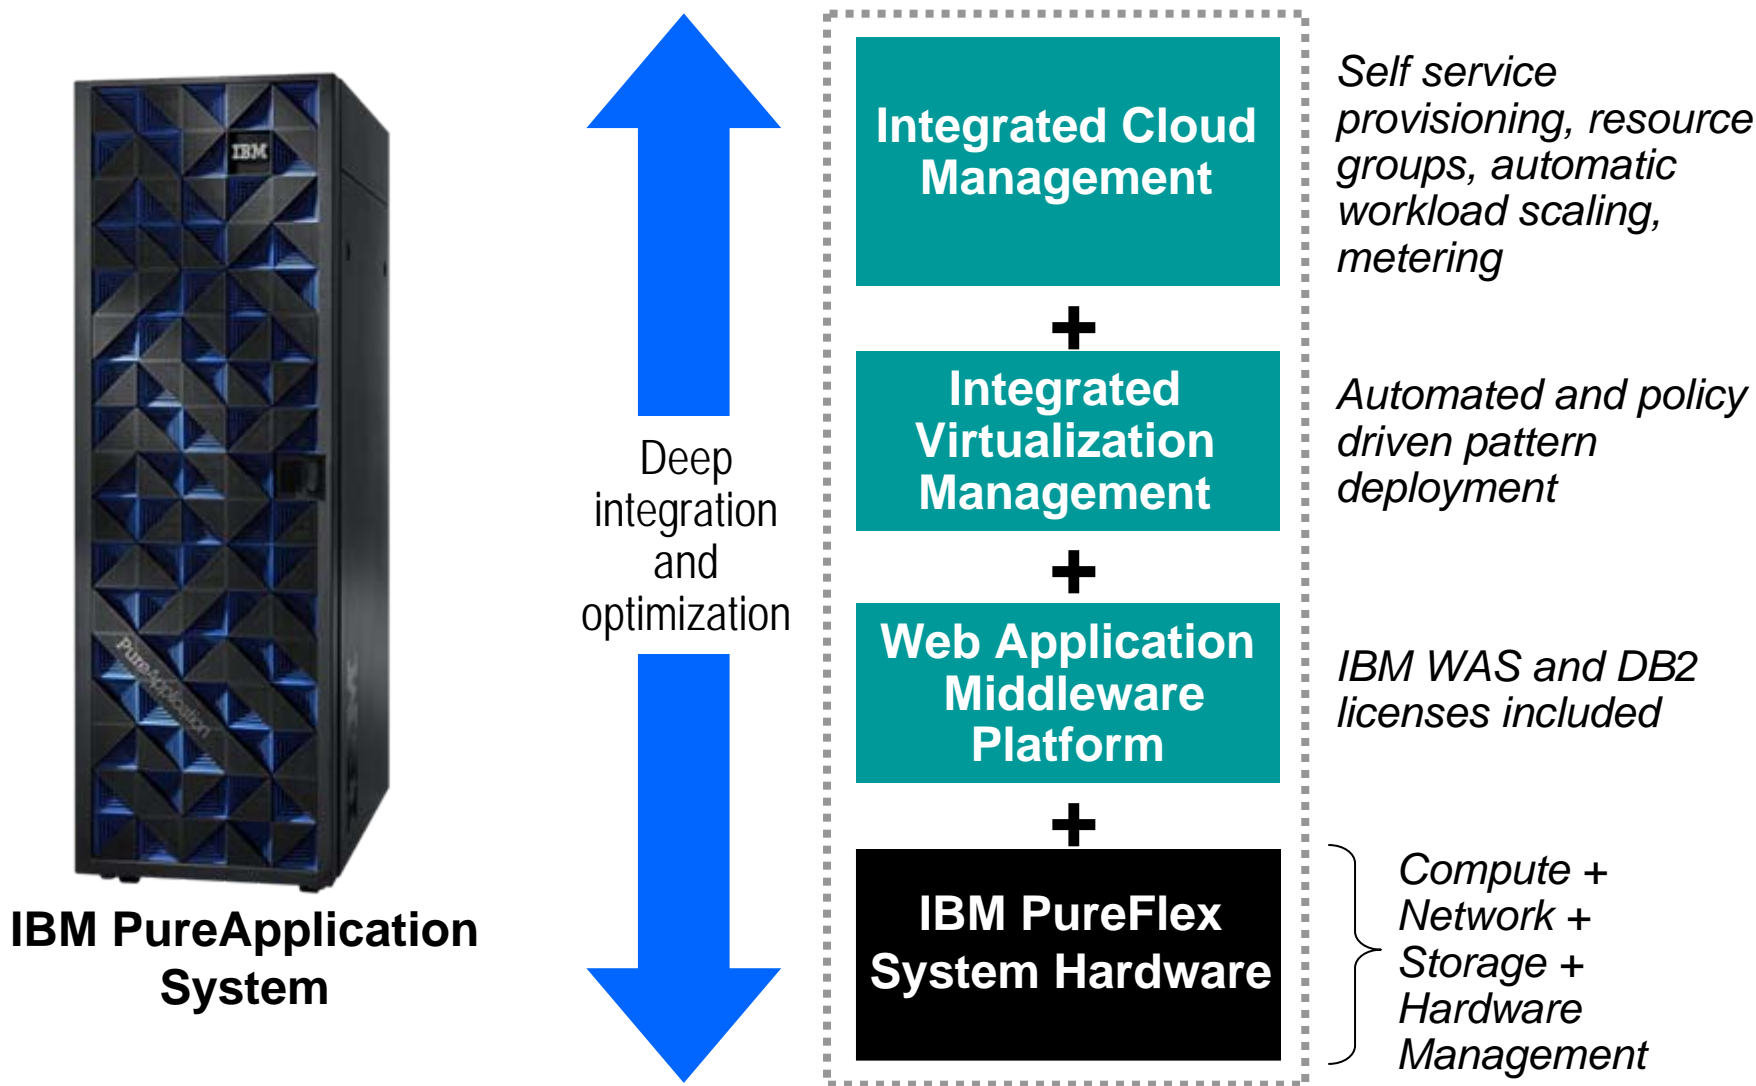
More Agility with automatic load balancing and scaling

Existing workloads and platforms

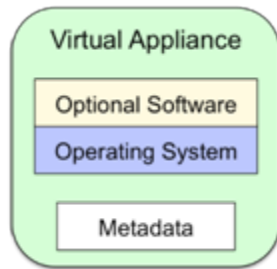
**Simplification
Standardization
Optimized Price/Performance**

**Labor efficiency
Agility**

IBM PureApplication System - *Optimized For Speed, Simplification, And Less Customer Labor*

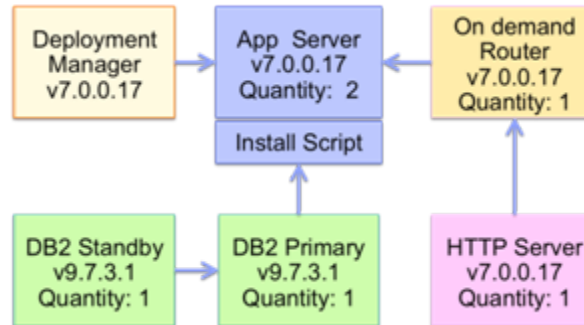


Easier Deployment With Built-In Expertise For Multiple Workload Pattern Types



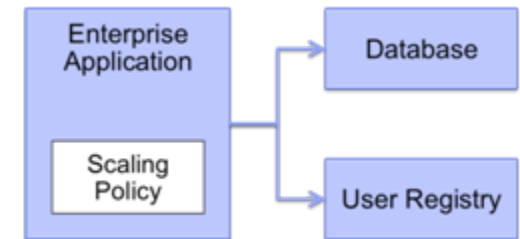
Virtual Appliance

- Standard software installation and configuration on OS
- Images created through extend/capture
- Traditional administration and management model
- Single server



Virtual System

- Flexibility and control over the middleware topology
- Script packages and administrator customizations
- Topology awareness
- Administrator-driven scaling
- Multi-server



Virtual Application

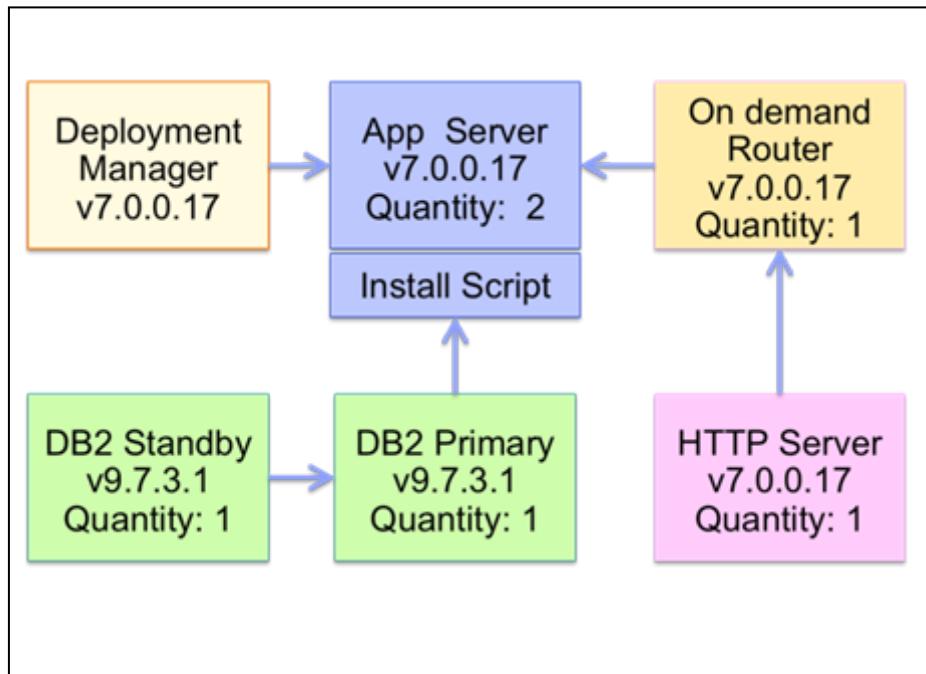
- Client describes application and IBM creates the topology for the deployed application
- Application awareness
- Policy-driven scaling
- Multi-server

Configuration flexibility
Skills required
More labor

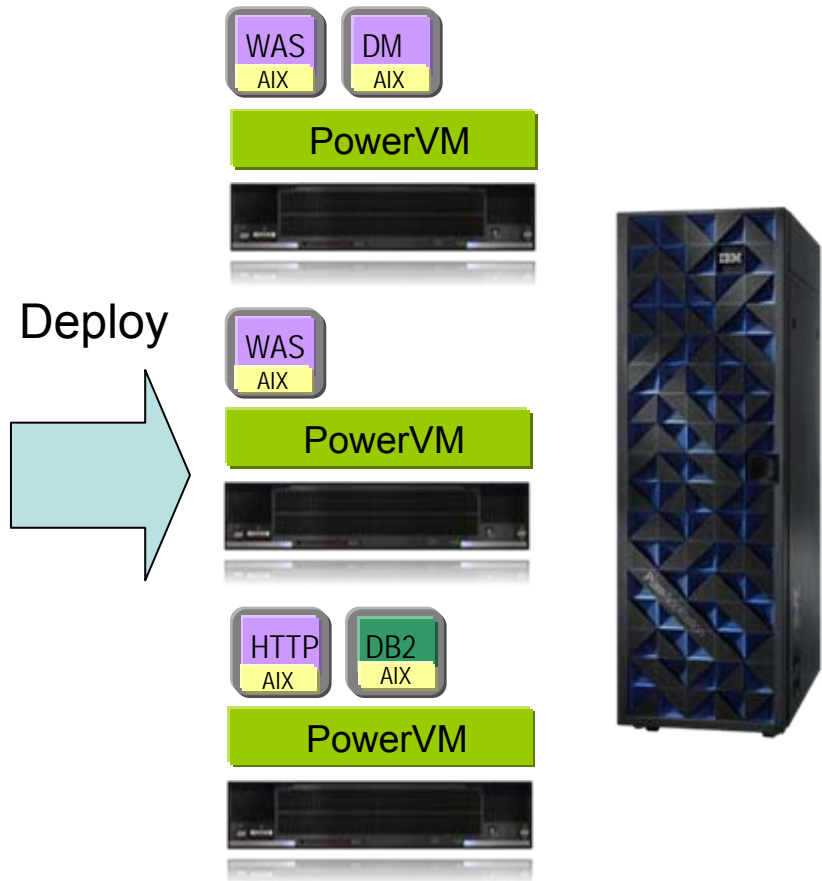
Easiest to use
Less skills required
Administrator productivity

Virtual System Patterns Speed Up Workload Deployment

Select Virtual System Pattern



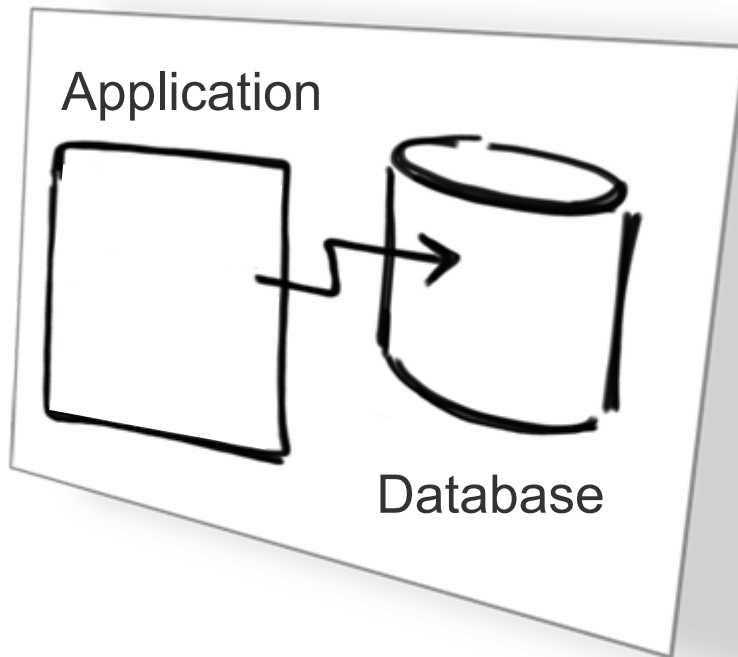
Describes virtual machines,
network connections, software
stacks and configurations



PureApplication Manager
deploys virtual machine images

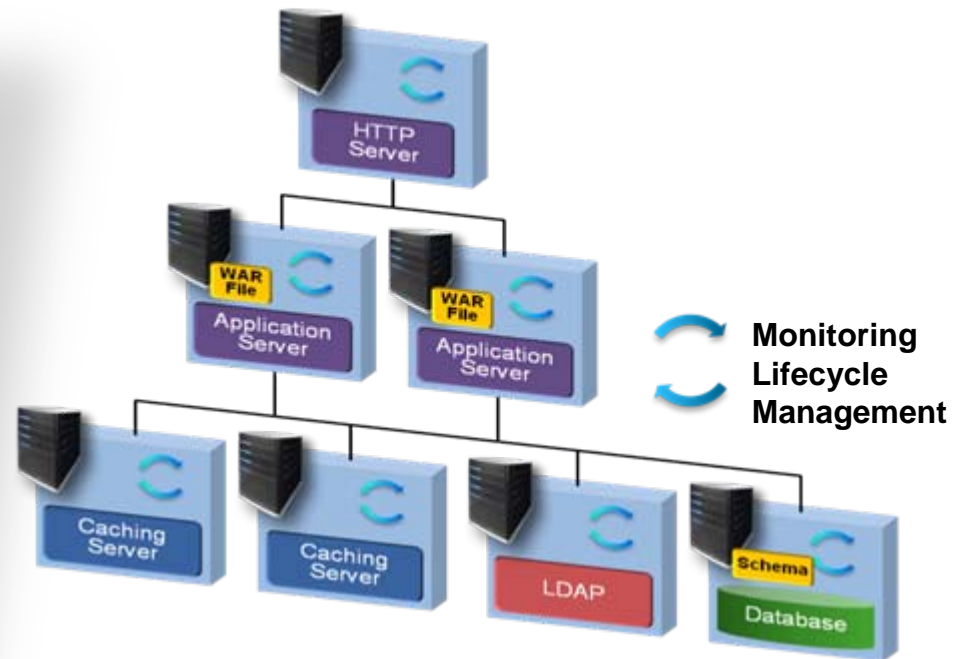
Virtual Application Pattern Further Simplifies Workload Management

What the business wants...



Just provide application code, DDL,
and specify policies

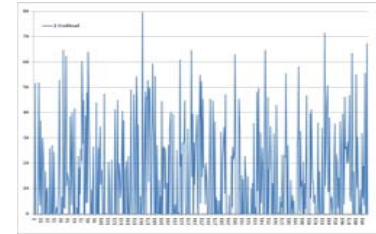
What's required...



PureApplication Manager constructs and
deploys this pattern

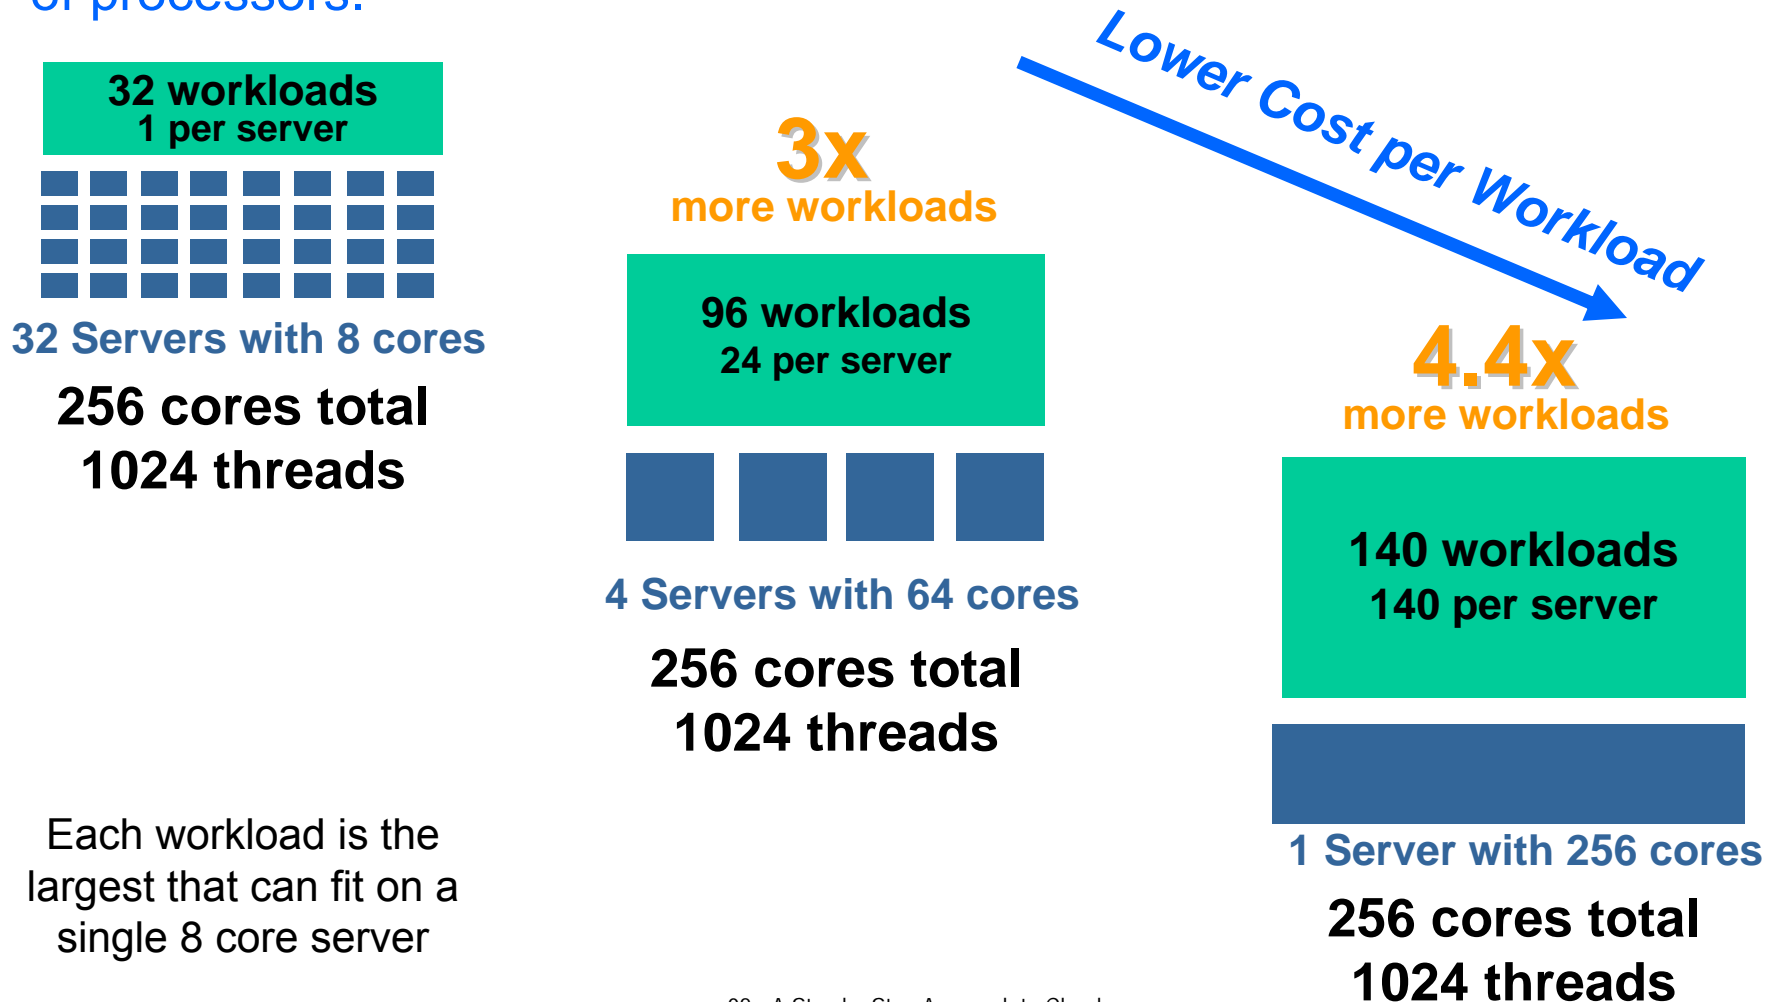
Larger Servers With More Resources Make Effective Virtualization Platforms

- Most workloads experience variability in demand
- When you consolidate workloads with variability on a virtualized server, the variability of the sum is less (statistical multiplexing)
- The more workloads you can consolidate, the smaller is the variability of the sum
- Consequently, bigger servers with capacity to run more workloads can be driven to higher average utilization levels without violating service level agreements, thereby reducing the cost per workload



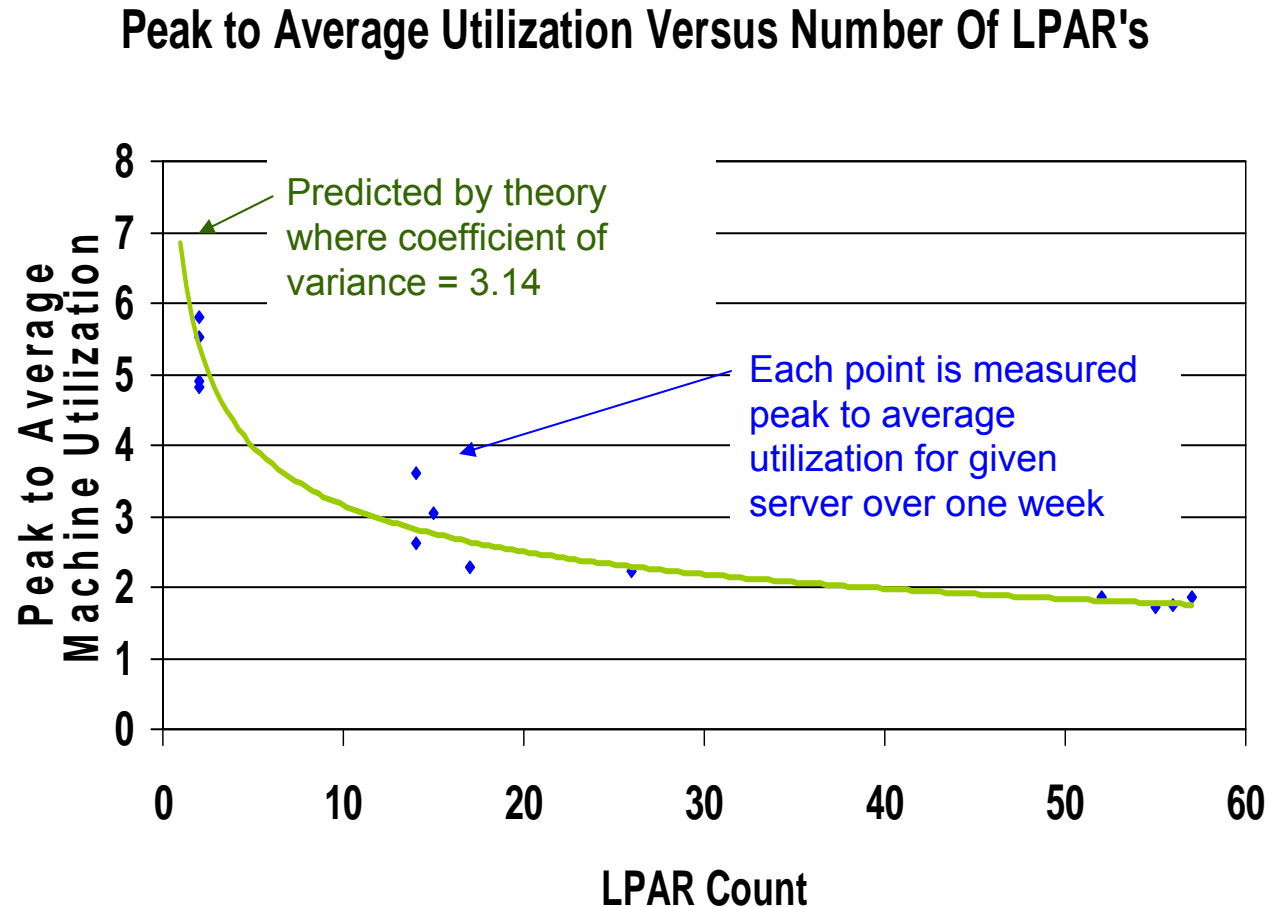
Theoretical Benefit Of Statistical Multiplexing

A single virtualized server with a large pool of shared processors can run more workloads than several smaller servers with the same total number of processors.



Surveys Of Customer Data Confirm Theory

- Customer utilization data from all 13 production servers for a week **matches** theoretical statistical multiplexing model
- Servers with more LPARs have less variance in their utilization

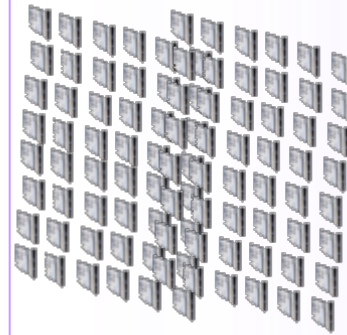


Statistical Multiplexing Helps Large Insurance Company Cut Costs (Upgrade Decision)

Older Power 595 and 780 servers (2,880 cores)



Upgrade to small new servers



414 Power 740 servers
2 socket servers
(13,248 cores)

\$116M 3 year TCA

Upgrade to large new servers

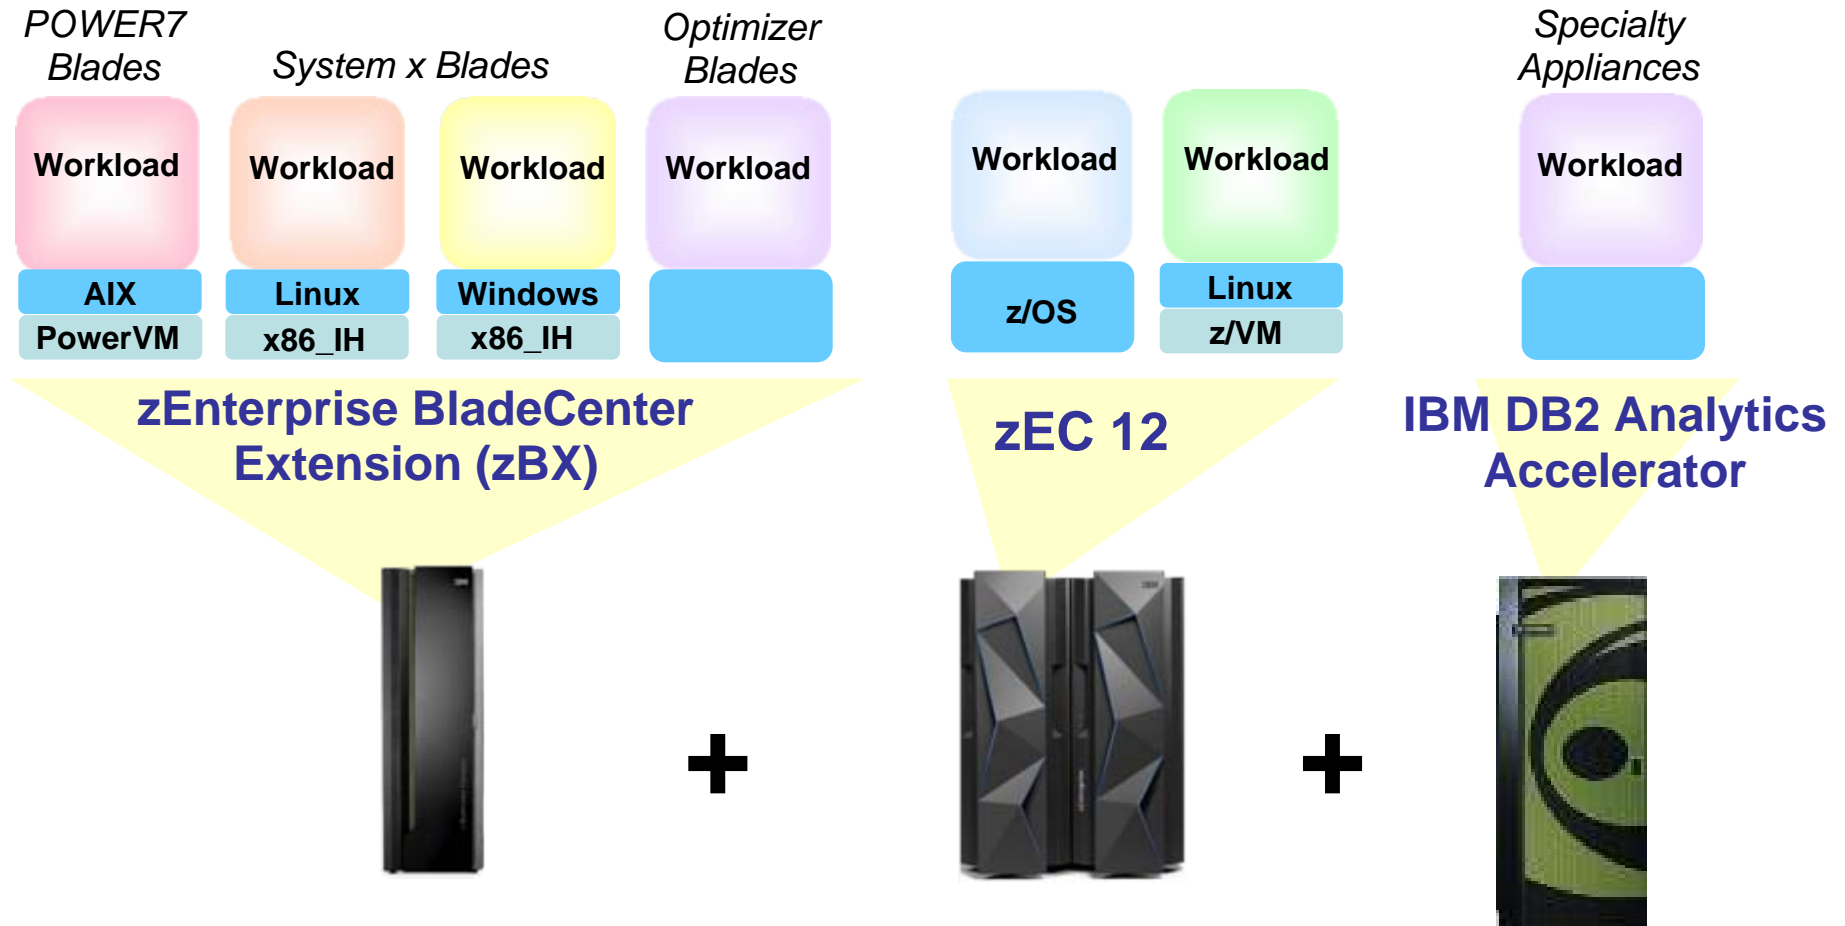


6 IBM Power 795
(1,280 cores)

\$54.6M 3 year TCA

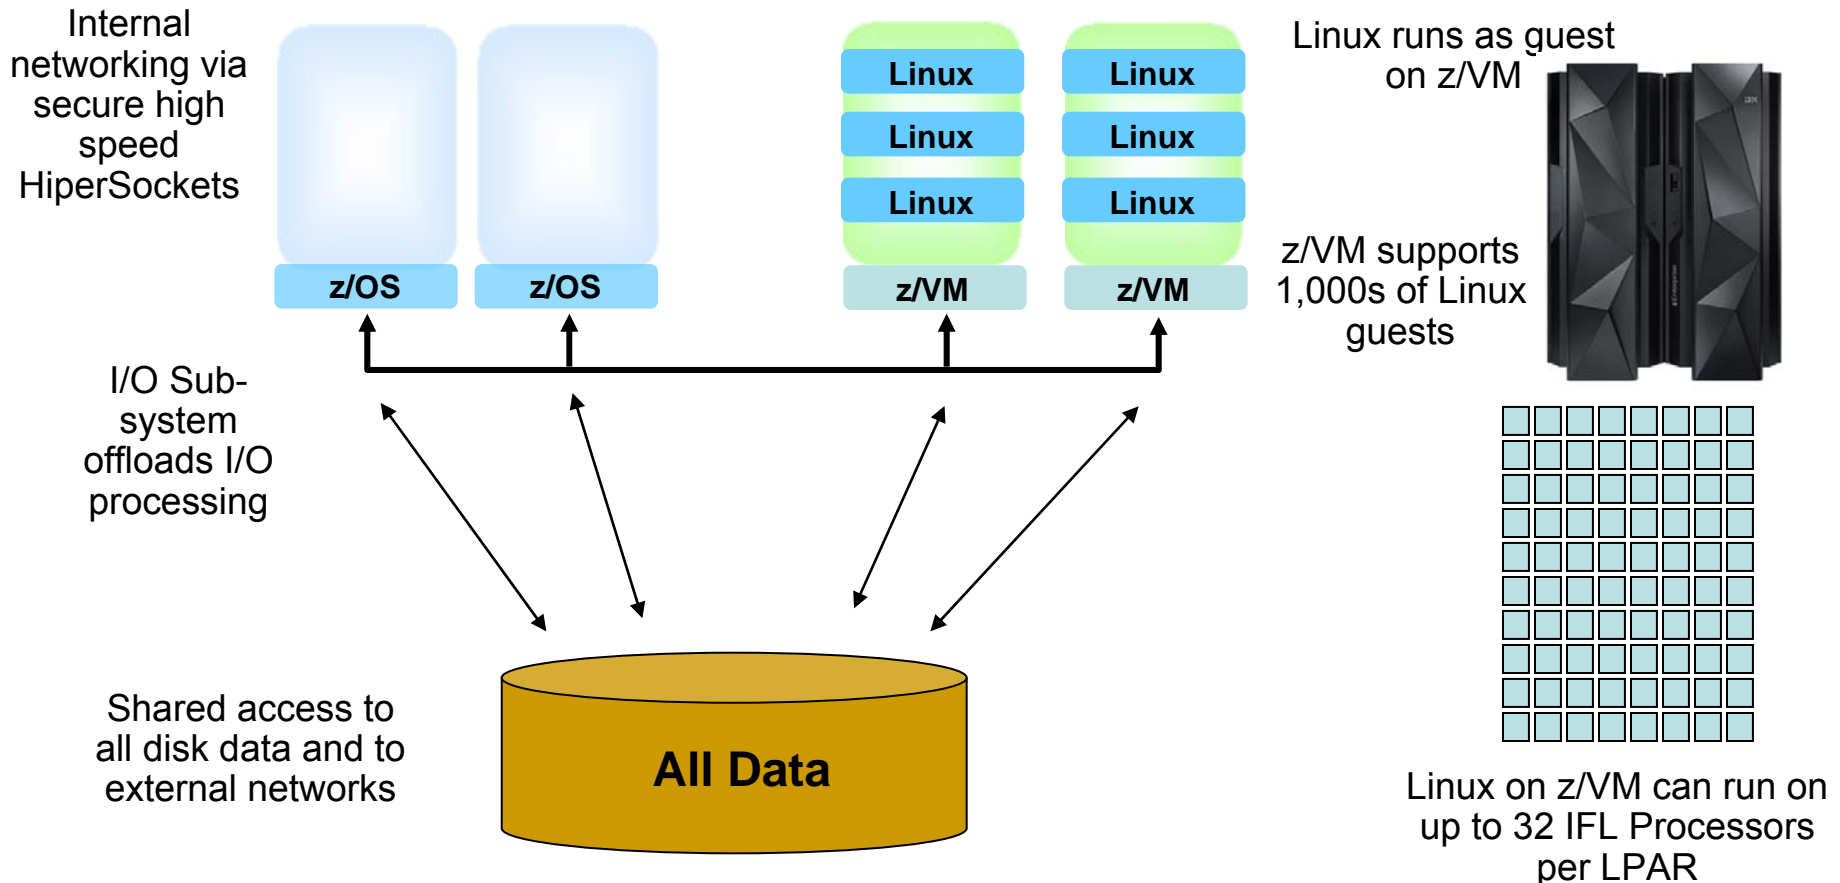
zEnterprise Provides Multi-Architectural Support For Workloads With z Affinity

Use a Best Fit Strategy for Workload Assignment



A Closer Look At Linux On z/VM

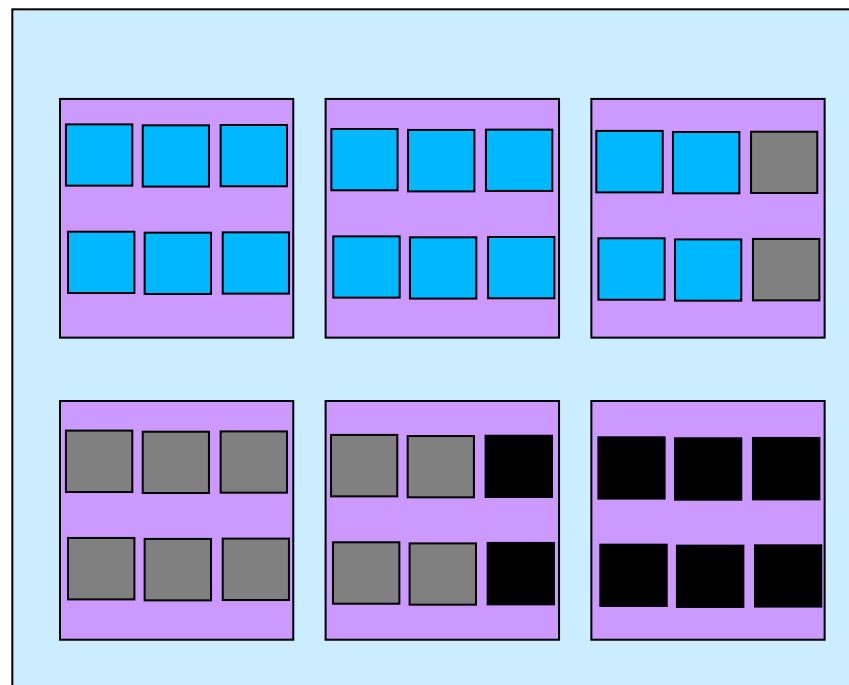
Logical Partitions Share Processors, Common Cache Structures, and I/O



System z Capacity On Demand Provides Elasticity To Handle Unexpected Peaks

- Capacity on Demand
 - ▶ “Books” are shipped fully populated
 - ▶ Activate dormant processors as needed
 - ▶ Use for temporary or permanent capacity
 - ▶ Self-managed on/off
- New capacity is immediately available for work without service disruption

One Book with 36 Processors



Active processors – pay full price



Inactive processors (On/Off CoD) – pay only 2% of full price



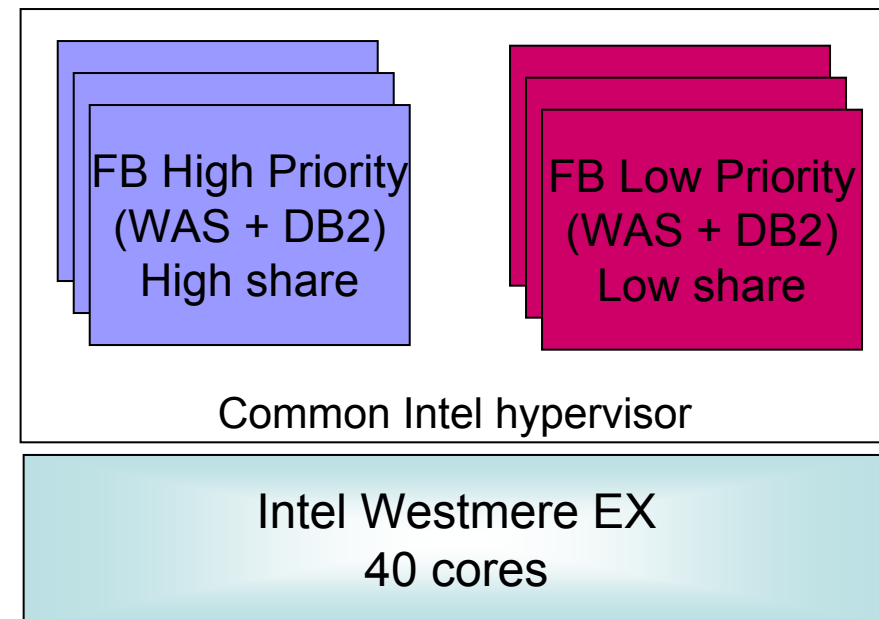
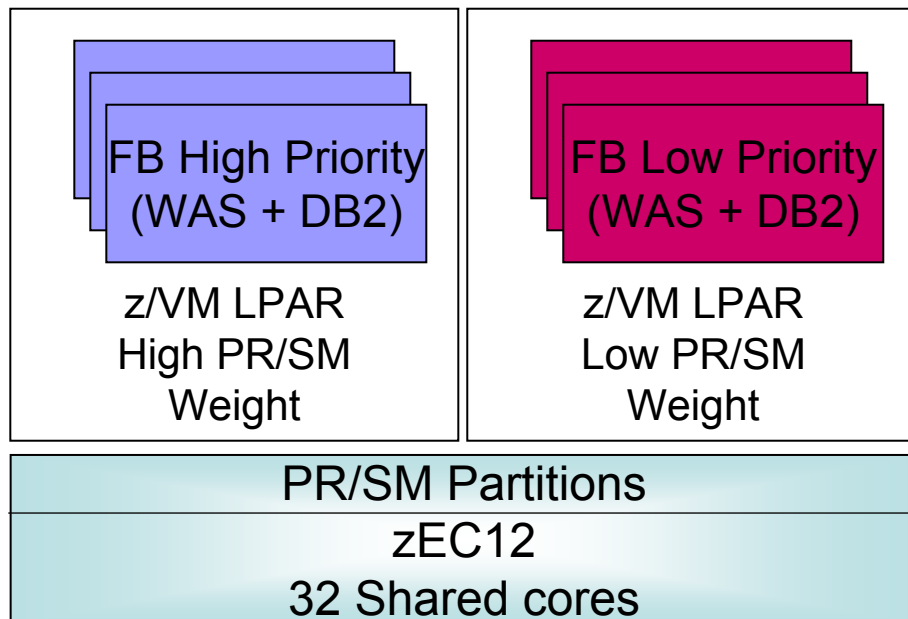
Dark processors (unused) – no charge

Workload Management In A Private Cloud Environment

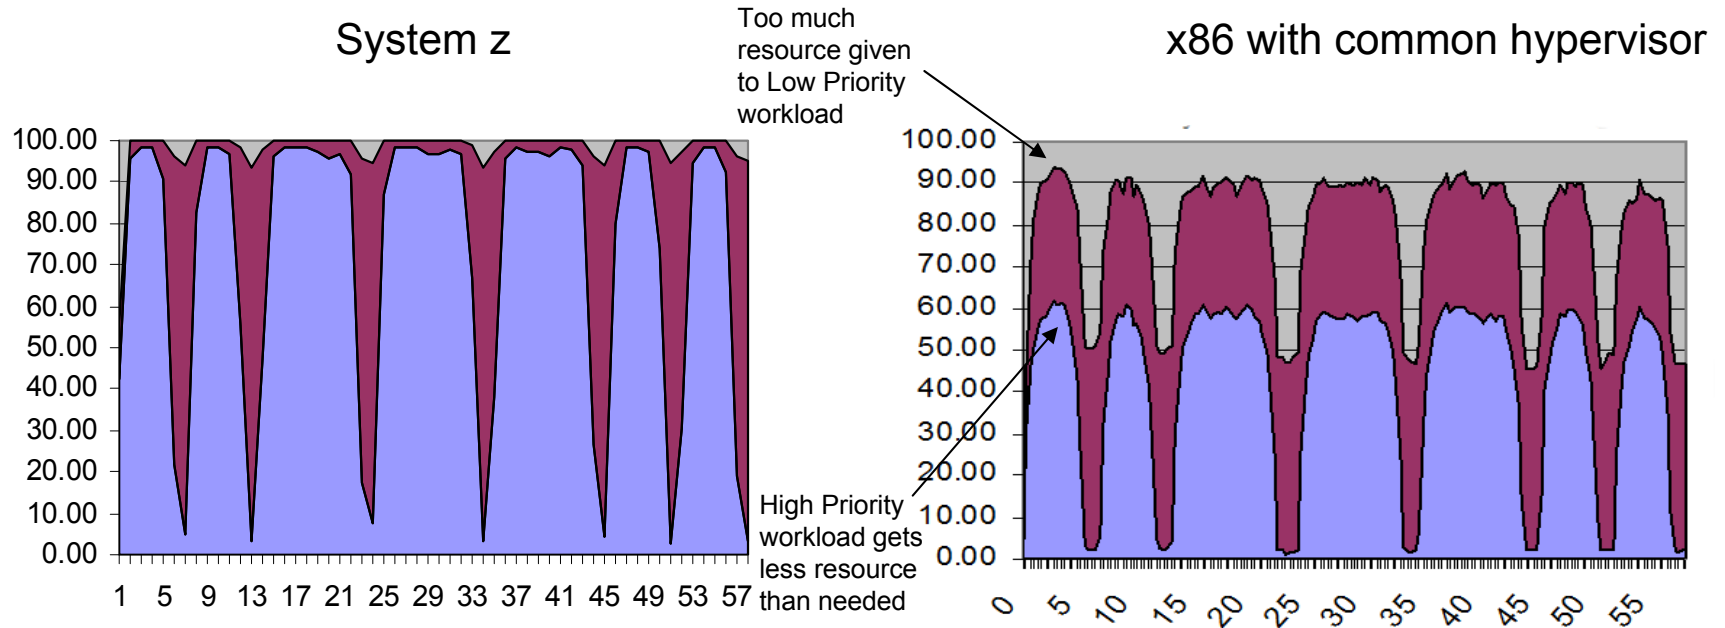
- Hosting platforms must be able to support high priority and low priority workloads together when sharing resources
 - ▶ Enables maximum utilization of the hosting platform
- Particularly relevant in a Private cloud environment
 - ▶ Multiple tenants with different priorities
- Desired behavior when mixing workloads
 - ▶ Low priority workloads “give up” resources to high priority workloads when required, soak up unused resources when available
 - ▶ High priority workload performance must not degrade

Comparison of System z PR/SM To Intel Common Virtualization Environments

- High Priority web workload has defined demand over time
- SLA requires that response time does not degrade
- Low Priority web workload has unlimited demand
- It “soaks up” unused CPU minutes



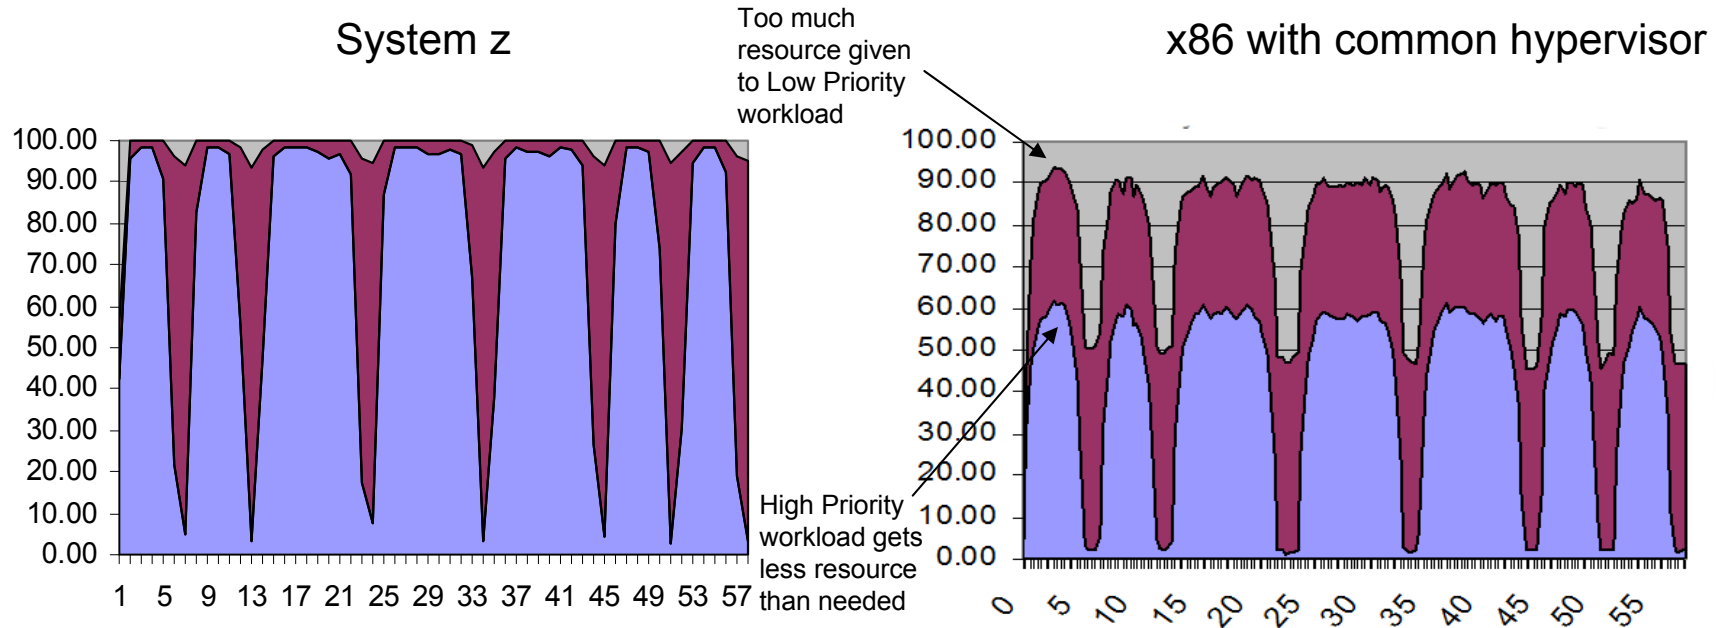
Workload Management Comparison Results



- Priority Workload
 - ▶ No throughput reduction
 - ▶ No response time increase
- Low Priority Workload
 - ▶ Soaks up remaining CPU minutes
- Unused CPU minutes 1.9%

- Priority Workload
 - ▶ 31% throughput reduction
 - ▶ 45% response time increase
- Low Priority Workload
 - ▶ Soaks up more CPU minutes
- Unused CPU minutes 21.9%

System z Virtualization Enables Mixing Of High And Low Priority Workloads Without Penalty

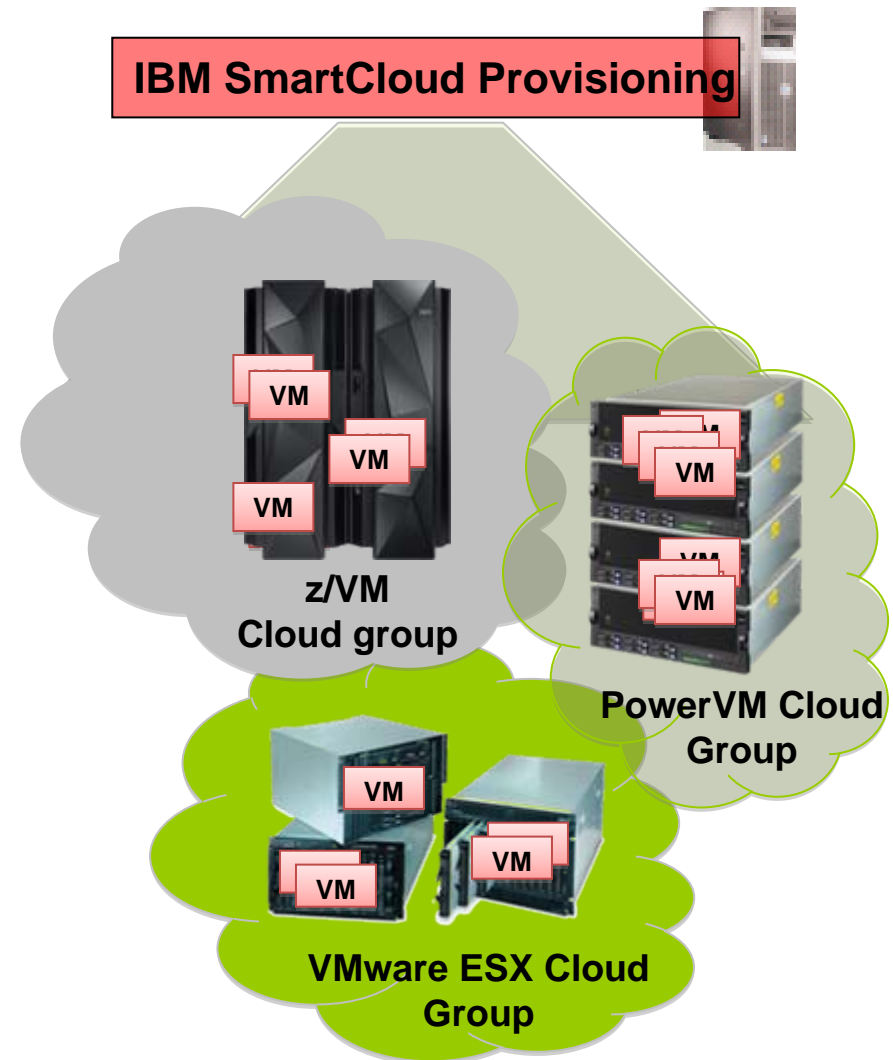


- Perfect workload management
- Consolidate workloads of different priorities on the same platform
- Full use of available processing resource (high utilization)

- Imperfect workload management
- Forces workloads to be segregated on different servers
- More servers are required (low utilization)

IBM SmartCloud Provisioning For Virtual Cloud Environments

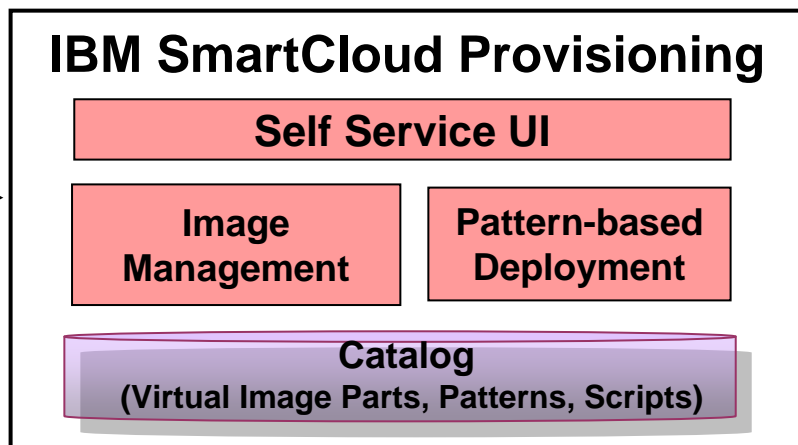
- Self-service automated provisioning of virtual systems and application into a pool/cloud of external virtualized hardware running a supported hypervisor
 - ▶ Supports various hypervisors, including, zVM, PowerVM, VMware ESX
- Standardized IBM middleware virtual images and patterns



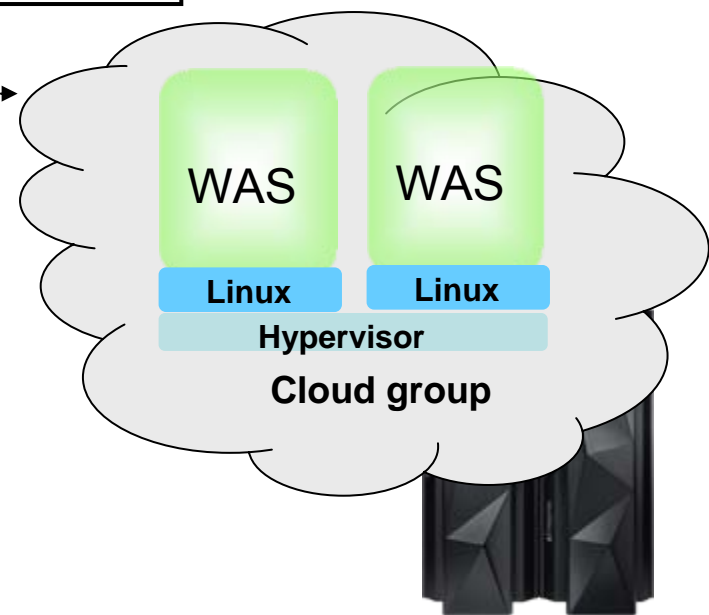
IBM SmartCloud Provisioning



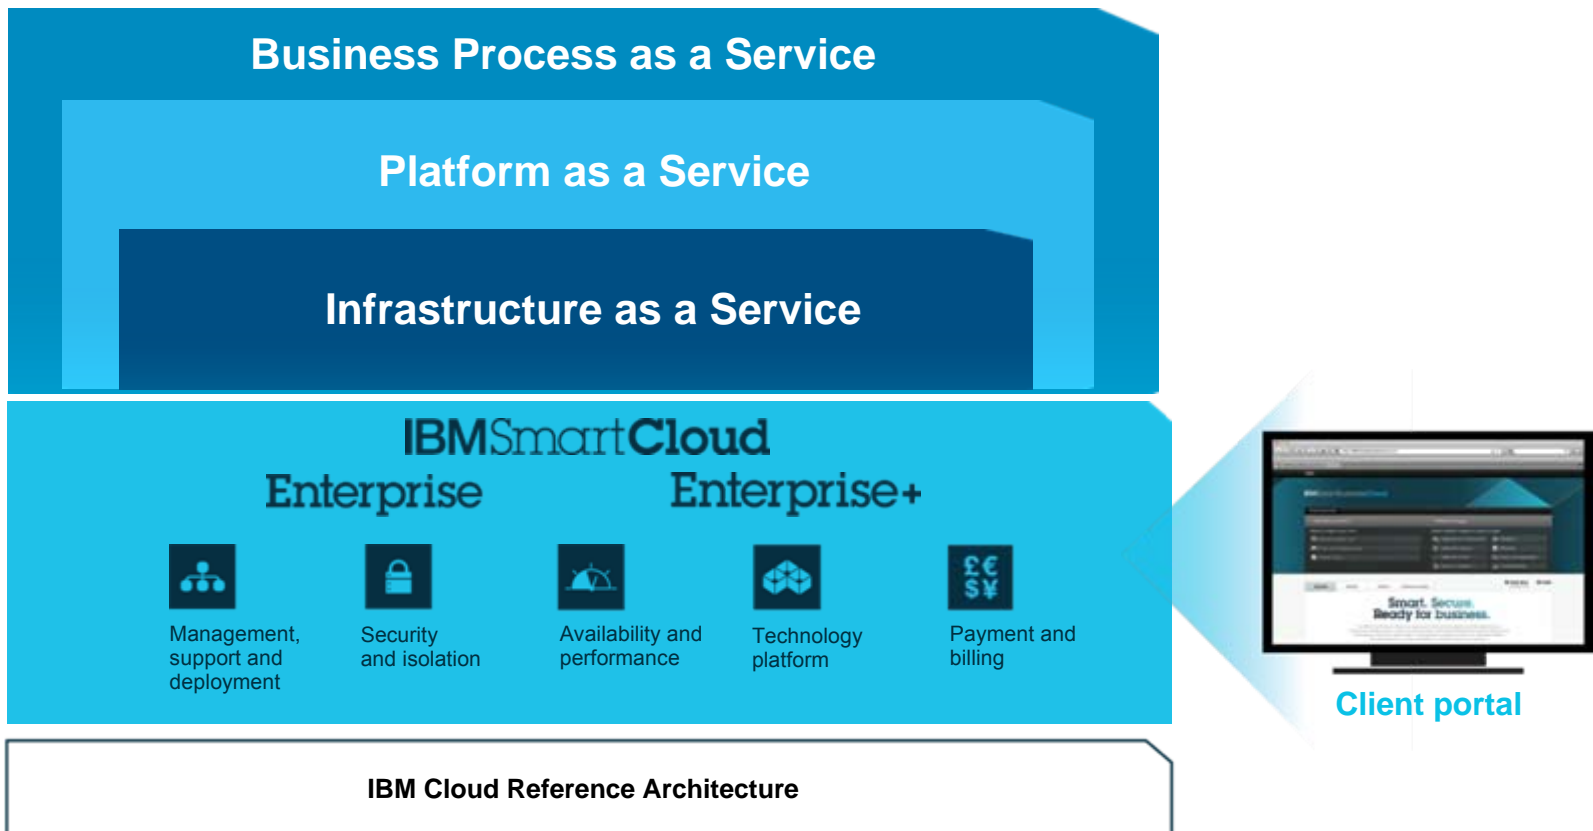
User creates and deploys a pattern



- Self-service console for users
- Catalog stores virtual image parts, patterns, scripts for quick-starts
- Drag and drop tooling for creating and deploying virtual applications using catalog
- Intelligent placement algorithm optimize resource utilization based on cloud activity



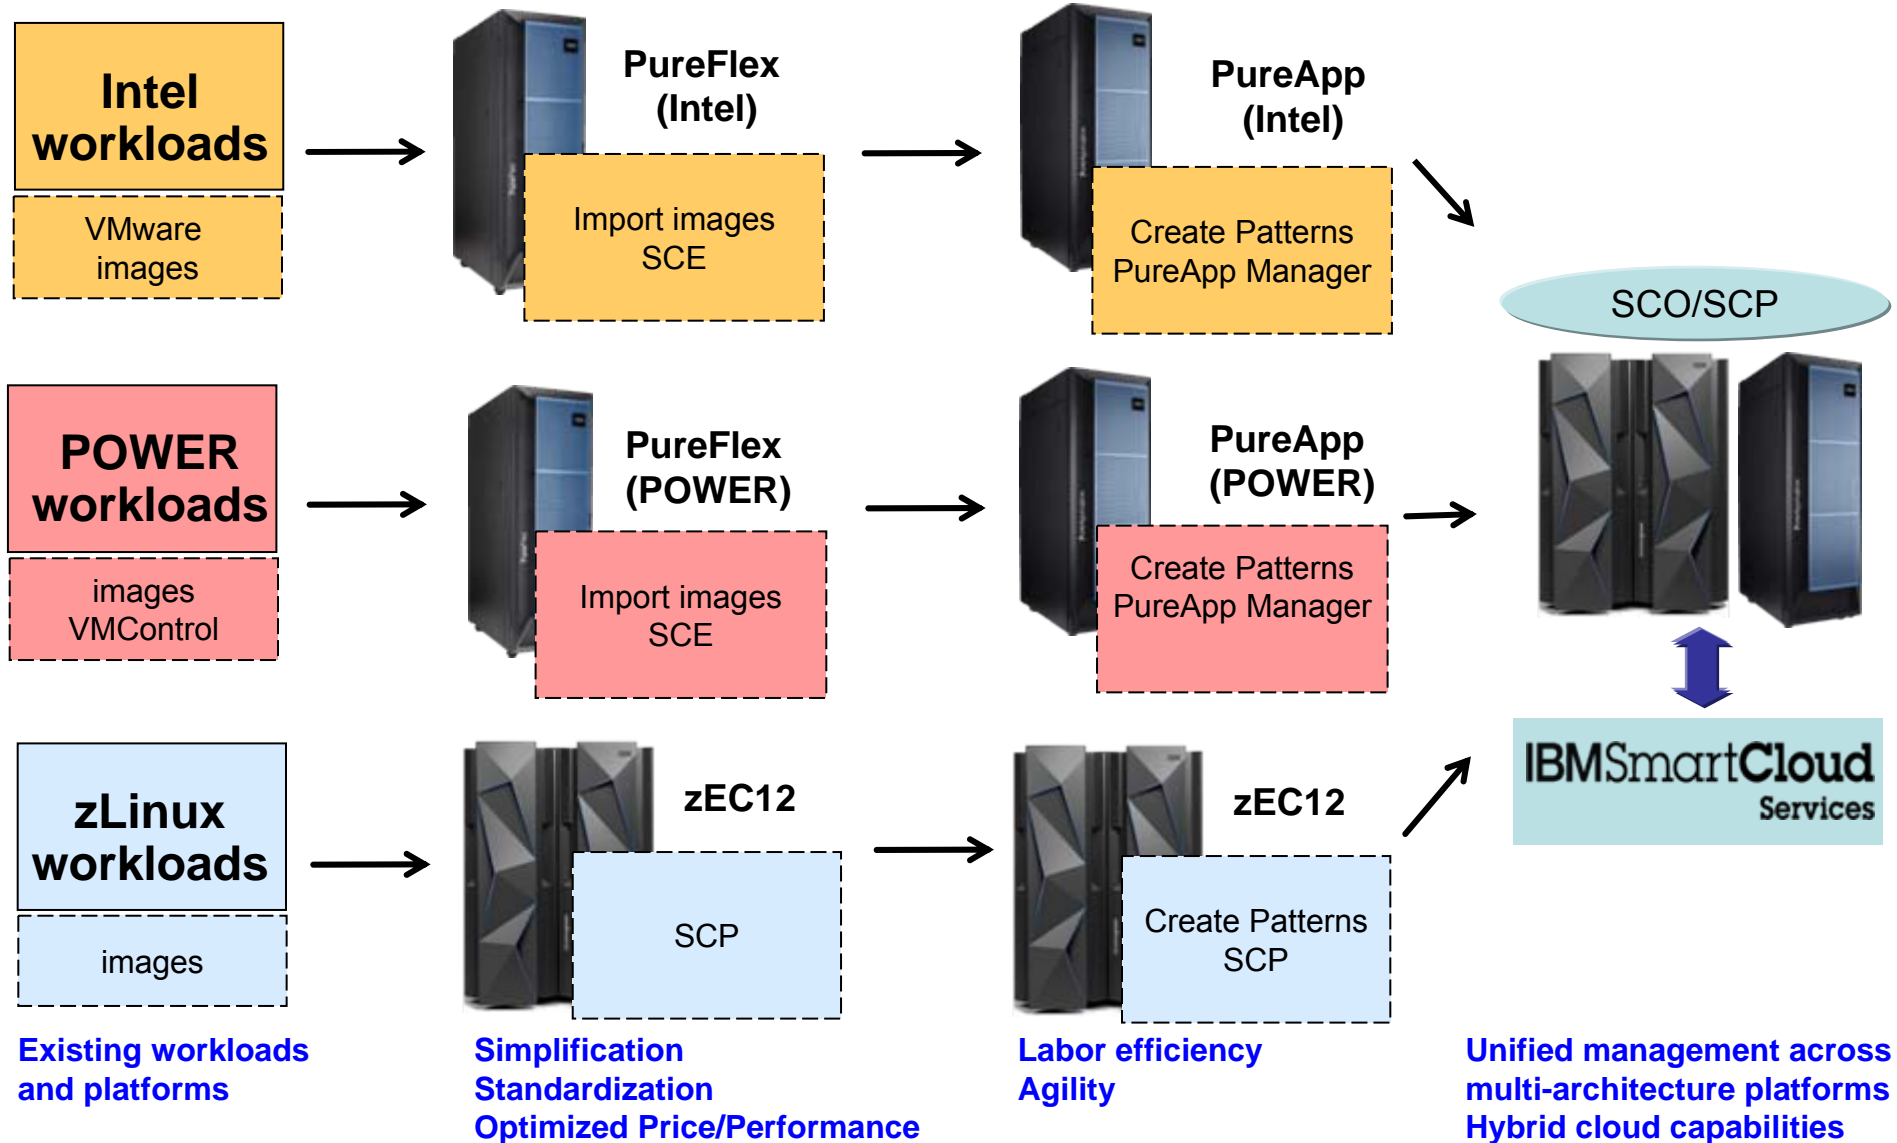
IBM's Public Cloud Offerings



Comparison of Public Cloud Offerings

	SmartCloud Enterprise	SmartCloud Enterprise+
Core Technology	KVM, x86 only	VMWare, AIX (x86 and Power Platform)
Application Applicability	Emerging application architectures that build resiliency into the application. Dev and test environments	Production, with full stack enablement for existing application designs
Service Scope	Unmanaged, cost optimized to compete with public IaaS providers	Managed Hosting service scope, 30%+ more efficient than traditional managed hosting deployments
Service Activation Performance	Self-service, ungoverned	~ 1 day, including full service activation. Shorter than hosted BAU today (2-4 weeks)
Guest Management	Client responsibility	Complete ITIL process support
Client ITSM Integration	Client responsibility	Integrated across client's non-cloud and cloud environments
Compliance	IBM internal security standards	ISEC above hypervisor, IBM internal below
SLA	99.5% for portal or VMs	98.5-99.9% for VM availability;
Service management	Below hypervisor	Above and below the hypervisor
Backup and recovery	Client Responsibility	IBM Provides as Integrated Service

Unified Data Center Cloud – Benefits



Strategic Direction – Unified Management For The Entire Data Center Cloud

- Creation of **common, standards based architecture** to avoid vendor lock in
- Entry points and offering packages are **modular** serving different markets, buyers, and price points
- **Clean upgrade paths** with **progression** to advance along the maturity curve

Bundle Option

SmartCloud Foundation

Factory Integrated &
Optimized by IBM

SmartCloud Orchestrator

Orchestrate Services across multiple environments and domains

SmartCloud
Provisioning

SmartCloud
Provisioning

Automate Optimized
Workloads

SmartCloud
Provisioning

SmartCloud Entry

SmartCloud Entry

High End Servers

Customer integrated
hardware

PureFlex System

PureApplication System

High End POWER
and
zEnterprise System

Blueprint For Cloud Success

- A practical cloud strategy should try to preserve and re-use existing image investments
- Survey workloads and identify candidates for a private cloud
- Assess best fit platform for each workload
 - ▶ Architectural match
 - ▶ Workload optimized systems
 - ▶ Individual workload size
 - ▶ Quality of Service requirements
- Define step-by-step projects
 - ▶ Immediate benefits for each step to maintain project enthusiasm
 - Price/performance
 - Labor
 - Agility
- Free IBM Eagle TCO service can help you assess fit for purpose and cost per workload