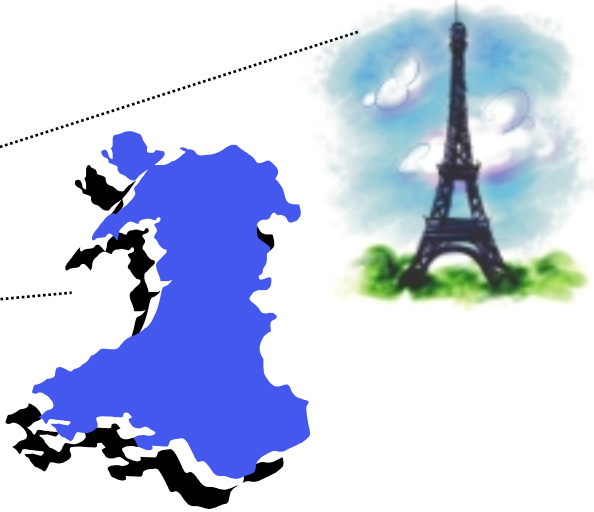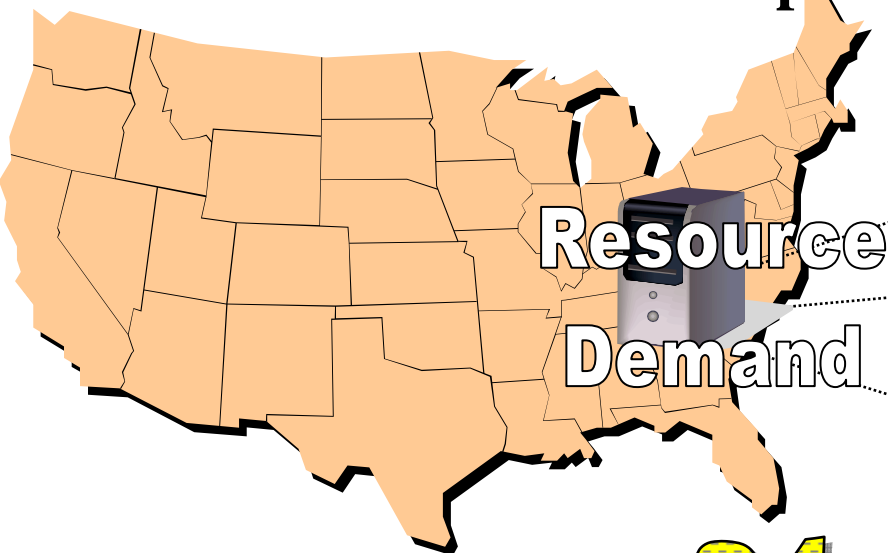# The Modern Mainframe – At the Heart of Your Business

The Mainframe Design Point Fundamentally Better

ODI's New Applications Will Quickly Gain Momentum

European Expansion

Resource Demand

24 x 7

Anticipated Growth

Fundamental Capabilities that make System z a Better Platform

- **Easy to grow capacity rapidly (Scalability)**

- Continuous operations (Availability)

- Flexible use of computing capacity (High Utilization)
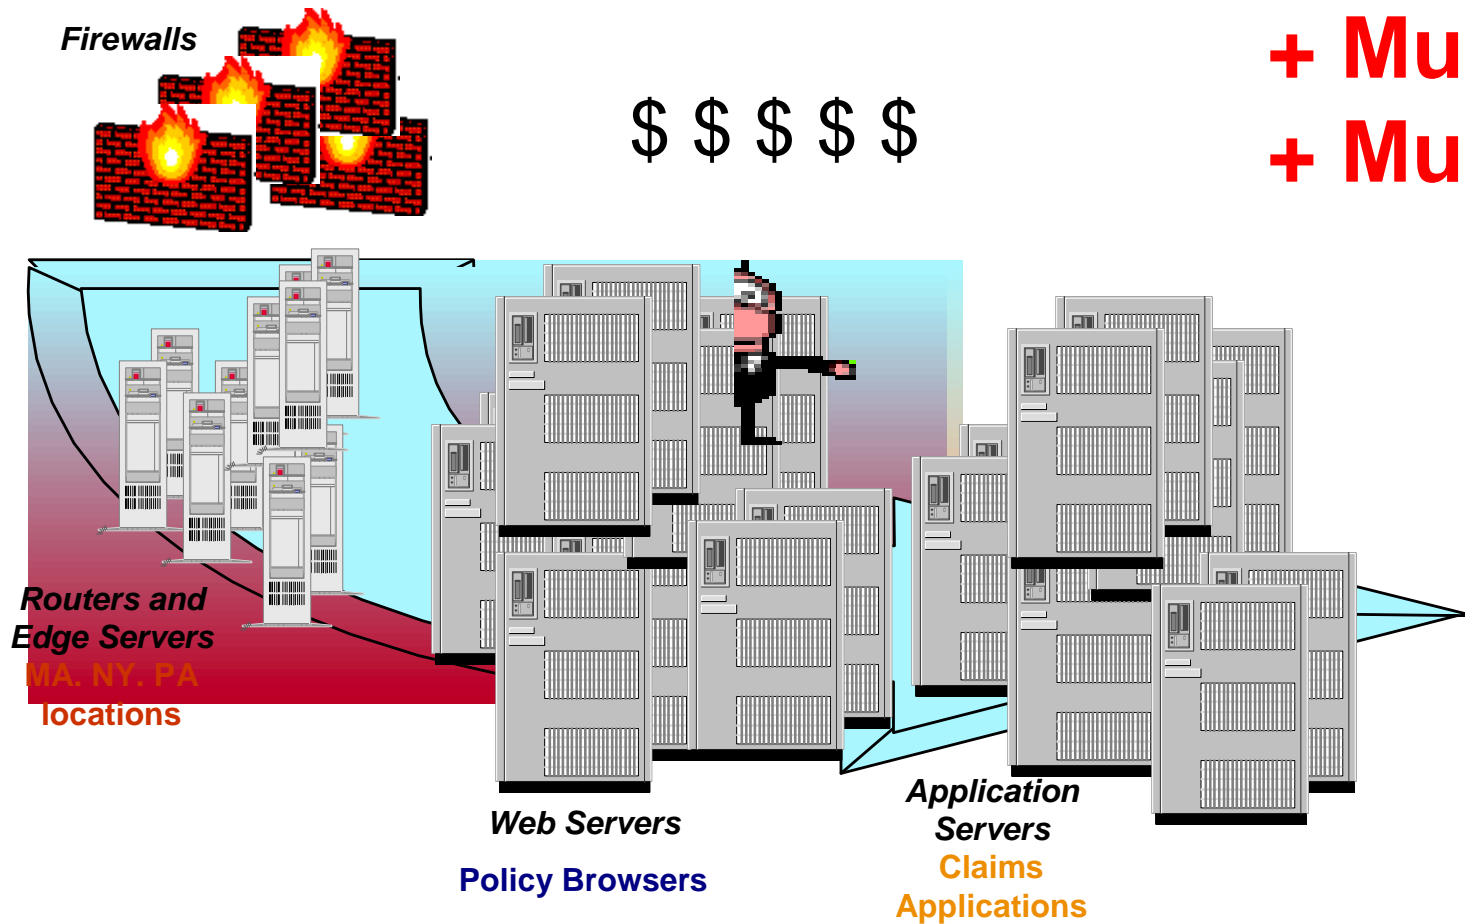
- Ease of management

Approaches to Scalability: Distributed vs. Mainframe

Distributed: scale out

Mainframe: both scale up and scale out

Distributed Server Scale Out

Solution: + Multiply!
+ Multiply!
+ Multiply!

Firewalls

$ $ $ $ $

Routers and Edge Servers
MA. NY. PA
locations

Web Servers
Policy Browsers

Application Servers
Claims Applications

How quickly can more capacity be added?

The Mainframe Solution

Solution: **Multiply!**
+ Multiply!
+ Multiply!

Firewalls

$ $ $ $ $

Routers and Edge Servers
MA. NY. PA locations

Web Servers

Policy Browsers

Application Servers

Claims Applications
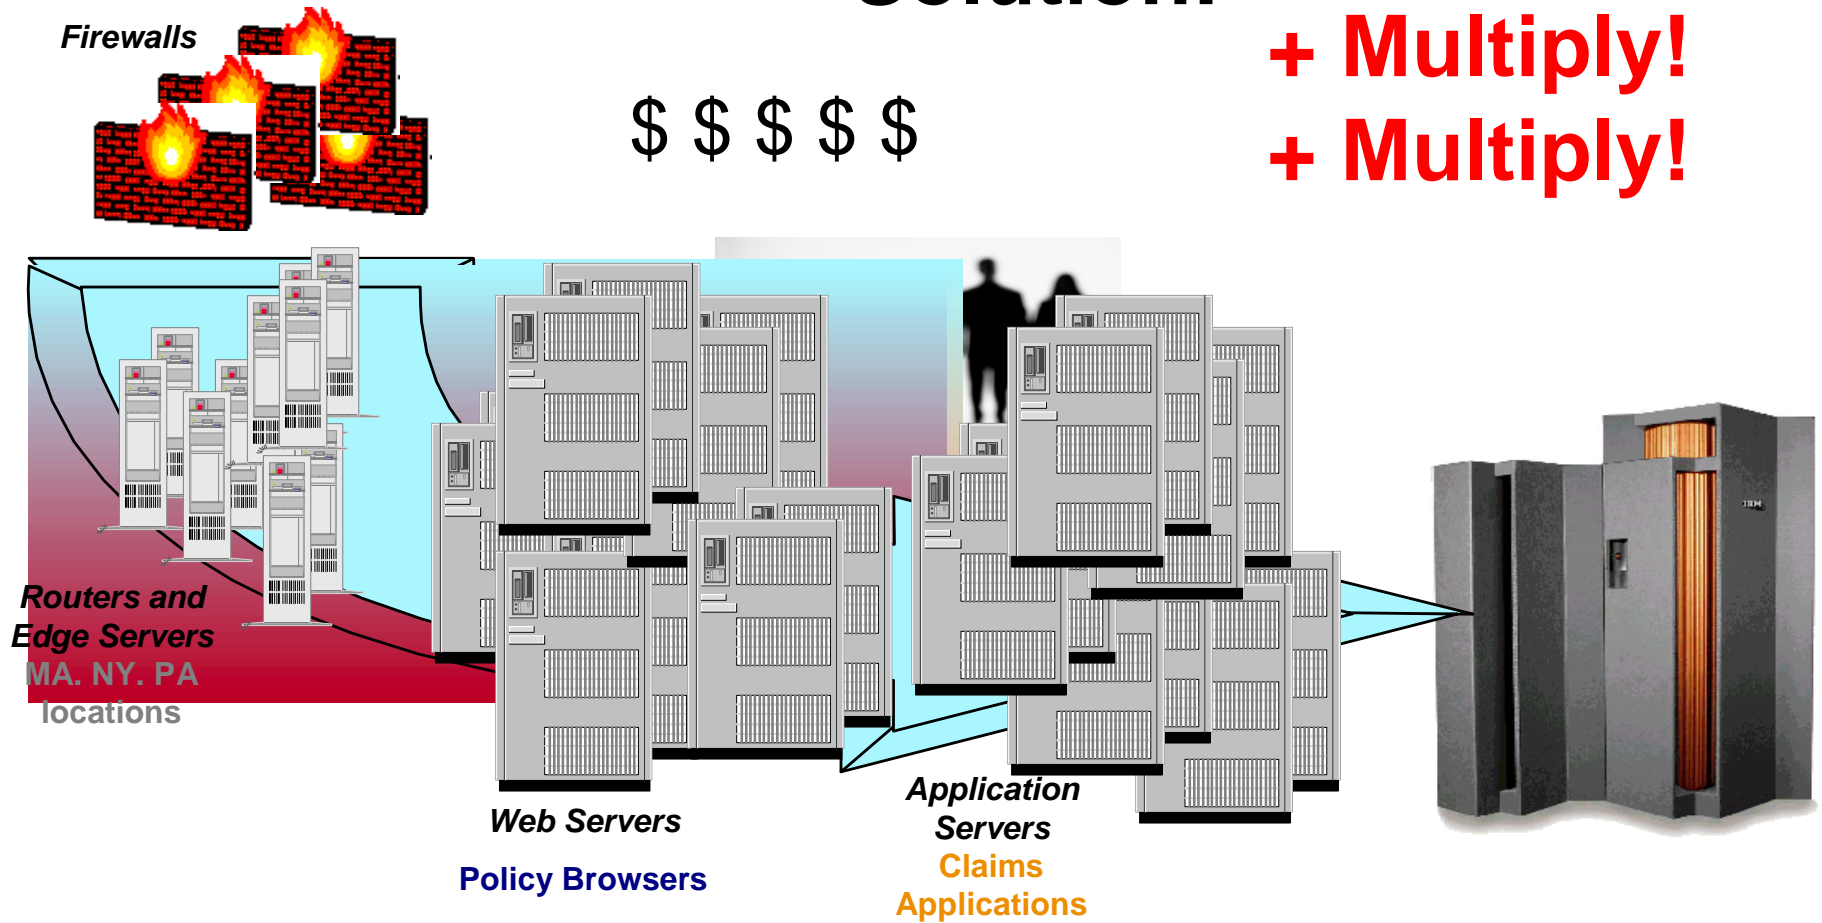
Scaling on System z

- Up to 54-way scaling in single SMP (18,000 MIPS)

- SYSPLEX clustering enables further growth with virtually no limit
  - ▶ Clustering up to 32x32-way

- I/O Bandwidth available to support growth

- Capacity on Demand
  - ▶ Grow as needed by turning on processors
  - ▶ Add processors without a shutdown

- System automatically uses new resources without re-engineering of systems or applications

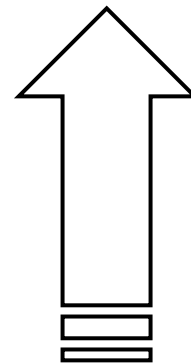- Automatically respond to changing business conditions

CPU Capacity on Demand

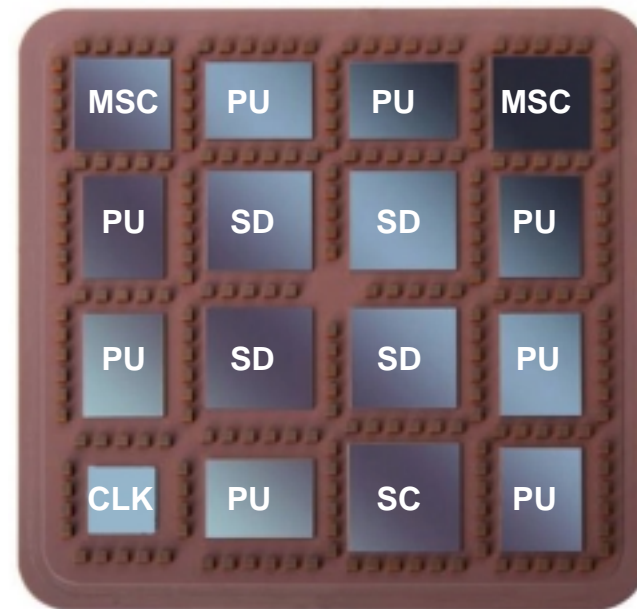**Permanent Growth - Customer Initiated Upgrade (CIU)**
Temporary Growth - Capacity On Demand (COD)
Nondisruptive addition of CPs (Software model upgrade)
ICF, IFL, zAAPs and memory (8 GB memory increments)

"Book"

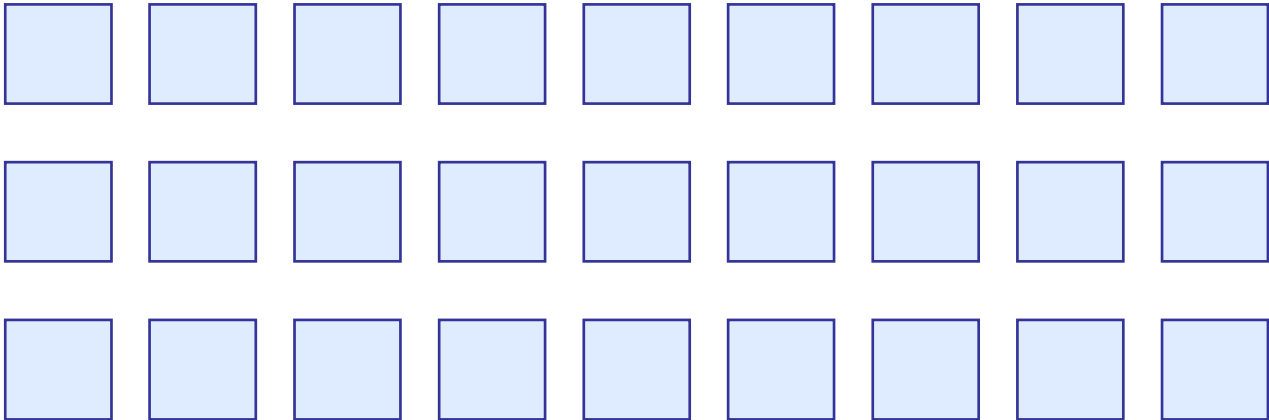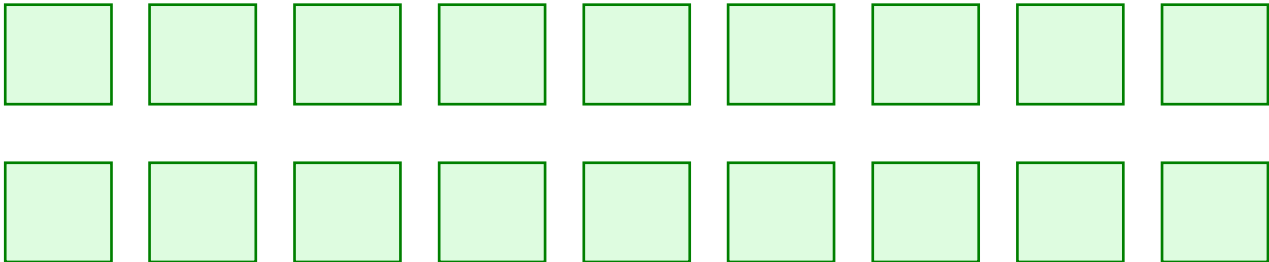| MSC | PU | PU | MSC |
PU	SD	SD	PU
PU	SD	SD	PU
CLK	PU	SC	PU

Capacity on Demand Allows Non-disruptive Growth

54 way SMP

Pay for these general purpose processors "Price per MIP"

Pay for zAAP, zIIP, and IFL processors at a reduced rate (~9% of price per MIP)

Do not pay for capacity on demand processors until used

Mainframe Capacity On Demand

Contract

Use Pre-shipped Hardware

www.ibm.com/servers/resourcelink

NO DISRUPTION

Add more memory too !

Quickly add more capacity

Measurable Benefits using CoD
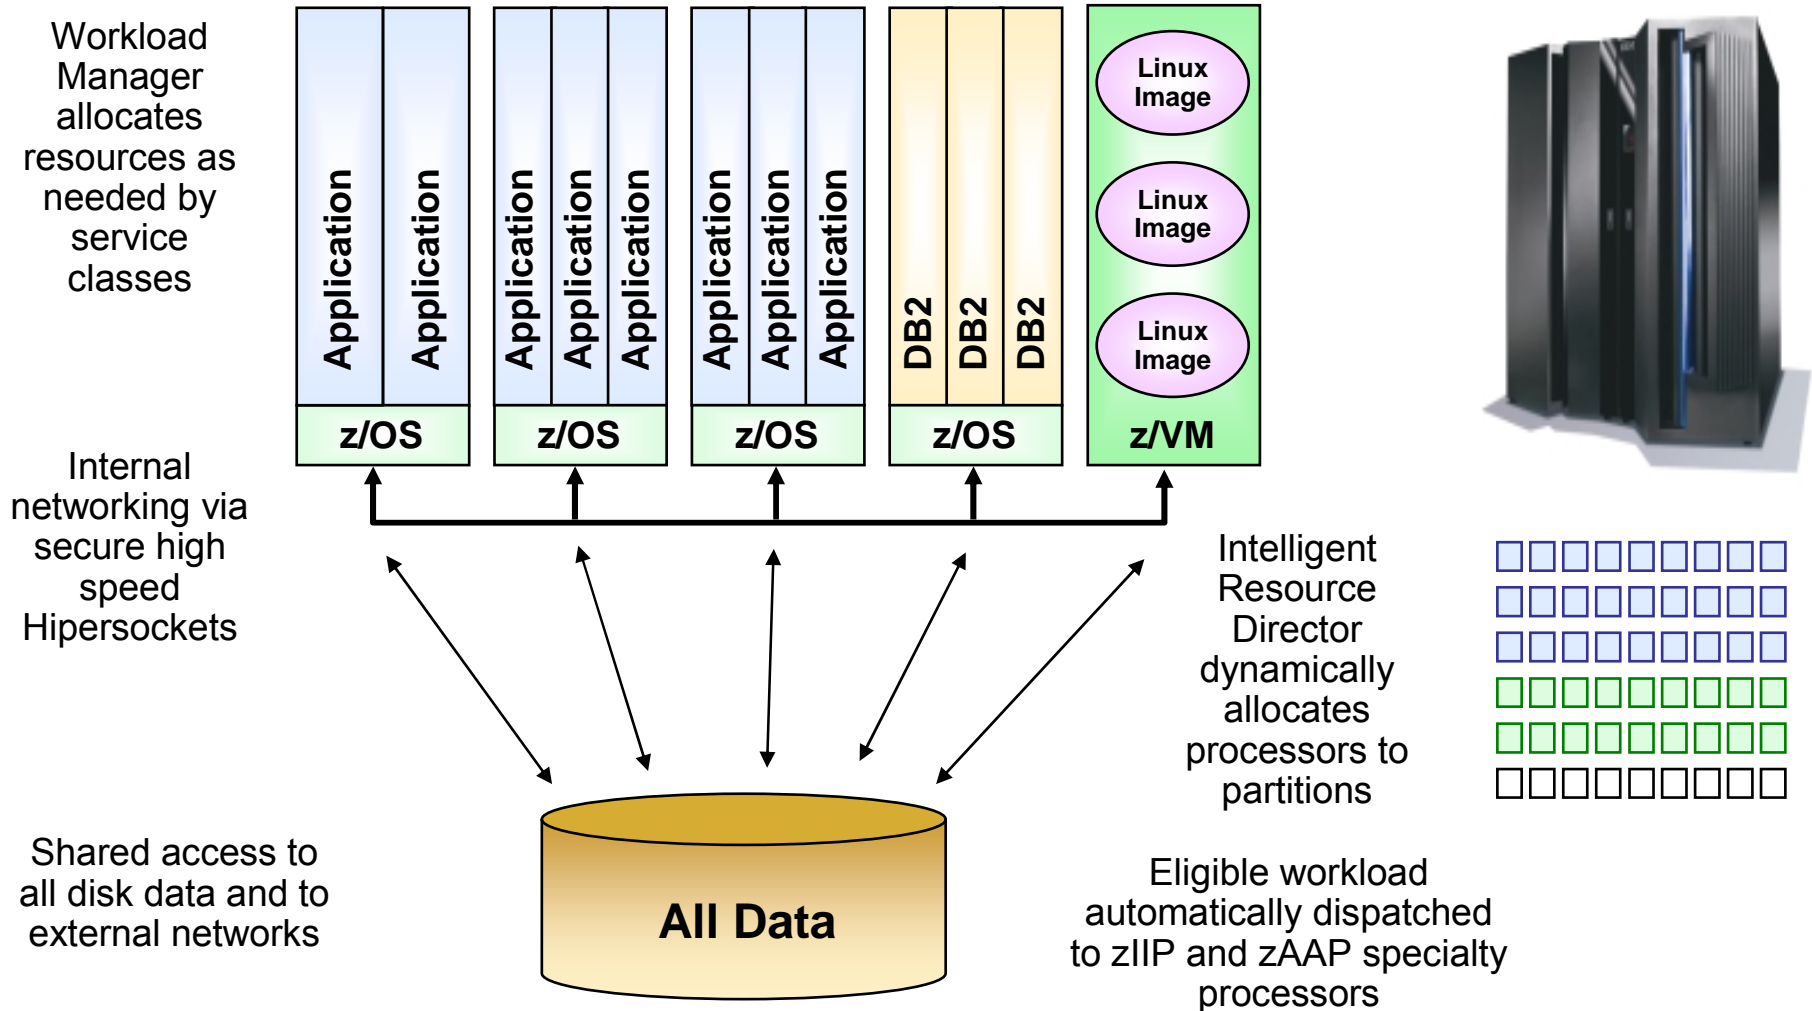
- Provides extra processing power to meet business demands
- Activate dormant processing capacity
- Quickly scale up non-disruptive without adding more servers
- Enables server consolidation to a single platform to reduce operational cost
- Allows investment in future paying only for immediate requirements
- Provides a convenient means of testing resources with minimal investment

Micro Virtualization in System z Enables Immediate Use of Additional Processors

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

Workload Manager allocates resources as needed by service classes

Application	Application		Application	Application	Application		Application	Application	Application		DB2	DB2	DB2		Linux Image / Linux Image / Linux Image
z/OS			z/OS				z/OS				z/OS				z/VM

Internal networking via secure high speed Hipersockets

Intelligent Resource Director dynamically allocates processors to partitions

Shared access to all disk data and to external networks

All Data

Eligible workload automatically dispatched to zIIP and zAAP specialty processors

Parallel Sysplex provides additional Scalability

- Enhanced scalability - system resources from up to 32 machines in the sysplex can be used to handle capacity growth

- Various subsystems (CICS, DB2, MQ etc.) all can take advantage

Coupling Technology

Sysplex Timer

System z9 109

zSeries

ESCON/FICON

Shared data
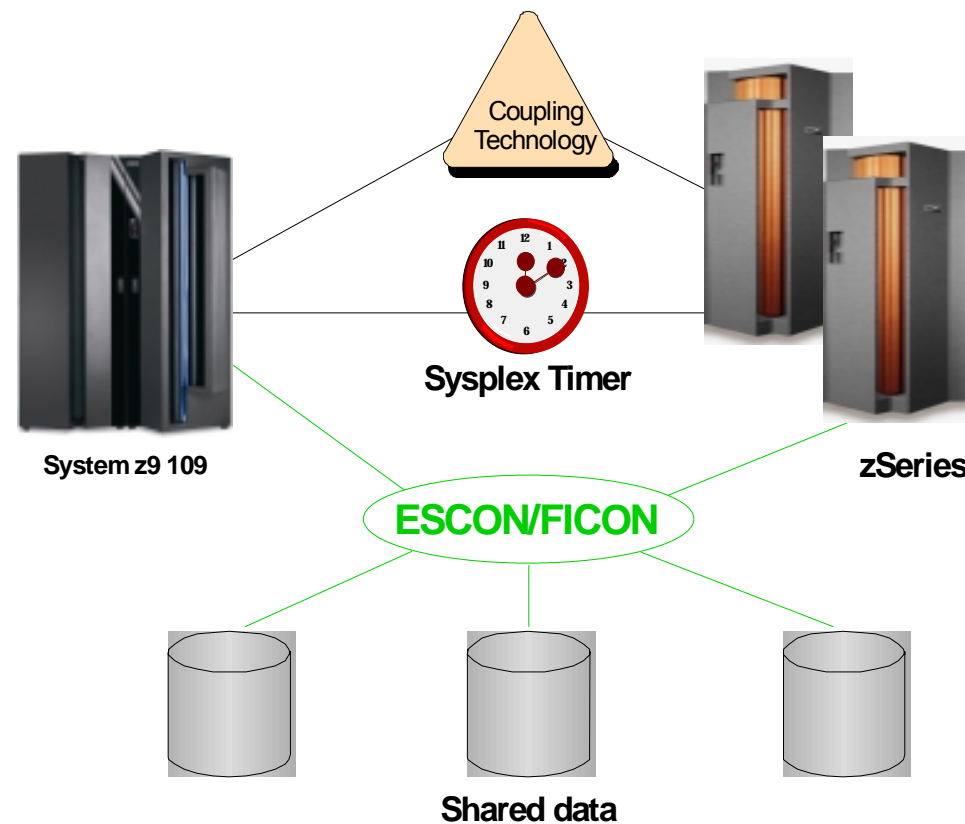
Benchmark for a Large Asian Bank Easily Achieved 8,000+ tps on System z

Database
- 4 DS8300
- 52 TB Storage
- 380 million accounts
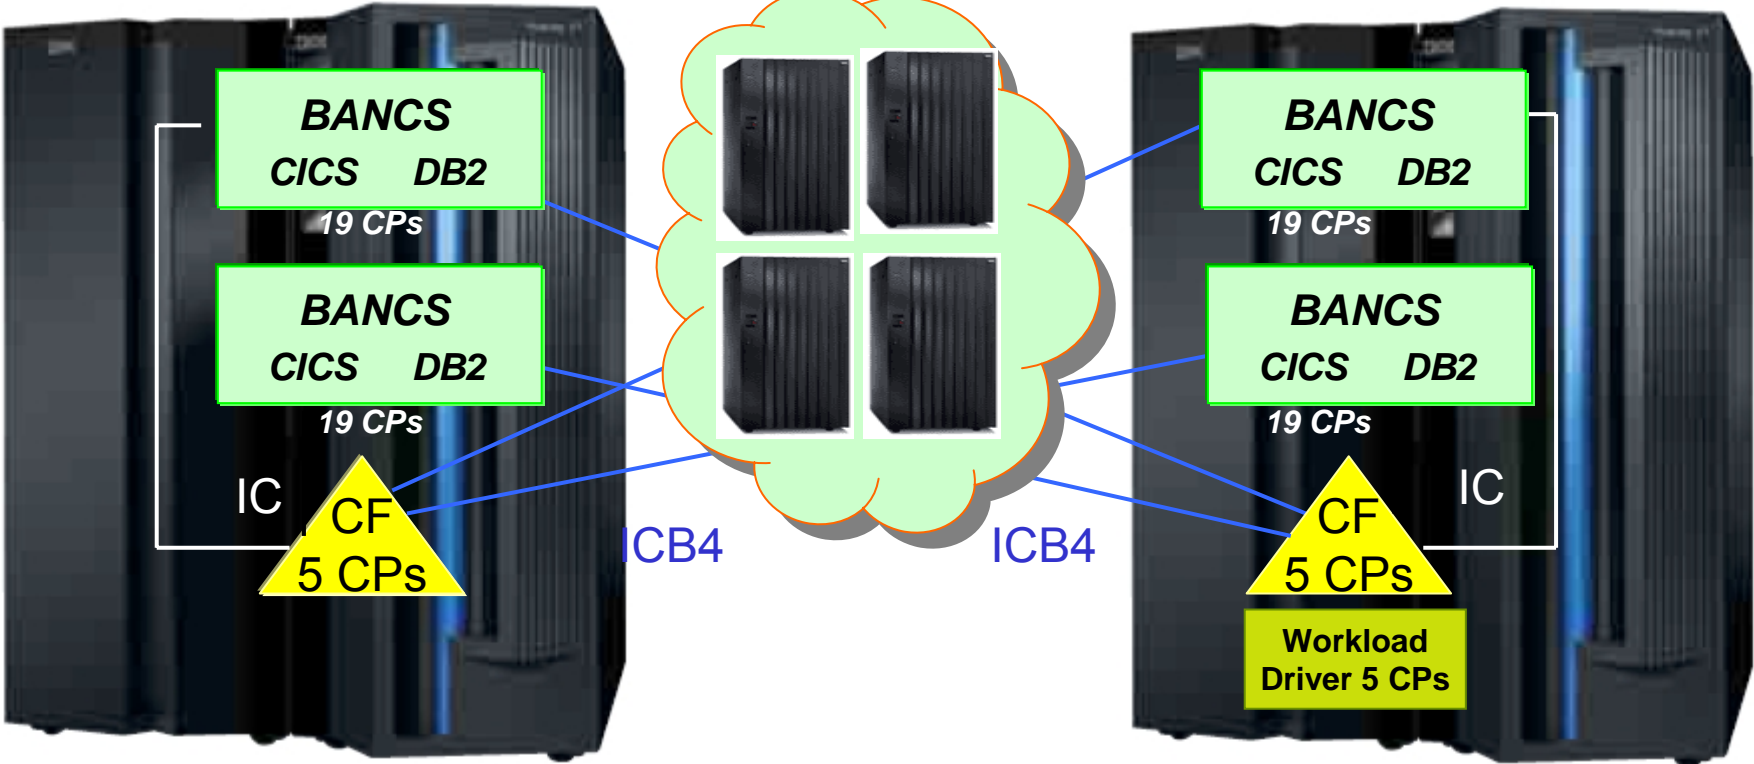
54-way z9

54-way z9

Performance Benchmark for a Large Asian Bank

Achieved: 8024 tps

Goal: 4,100 tps

- 8,024
- 5,723
- 4,665
- 3,120
- 1,589

y-axis: tps (0, 1000, 2000, 3000, 4000, 5000, 6000, 7000, 8000, 9000)

x-axis: MIPS (0, 5000, 10000, 15000, 20000)
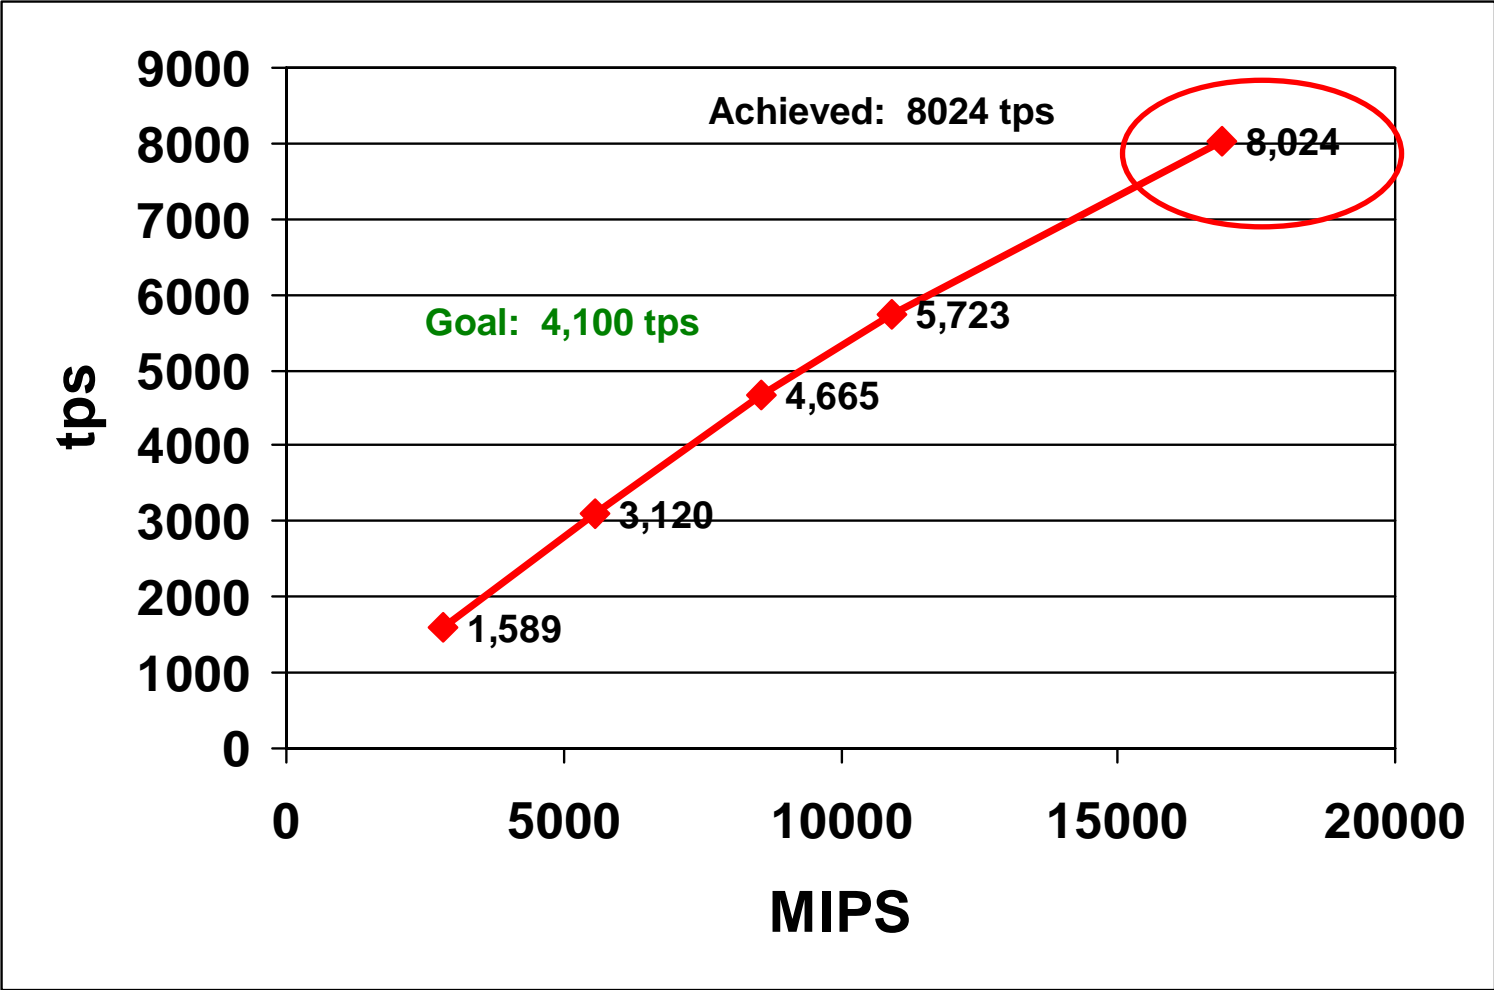
Scalability Requires I/O Capacity
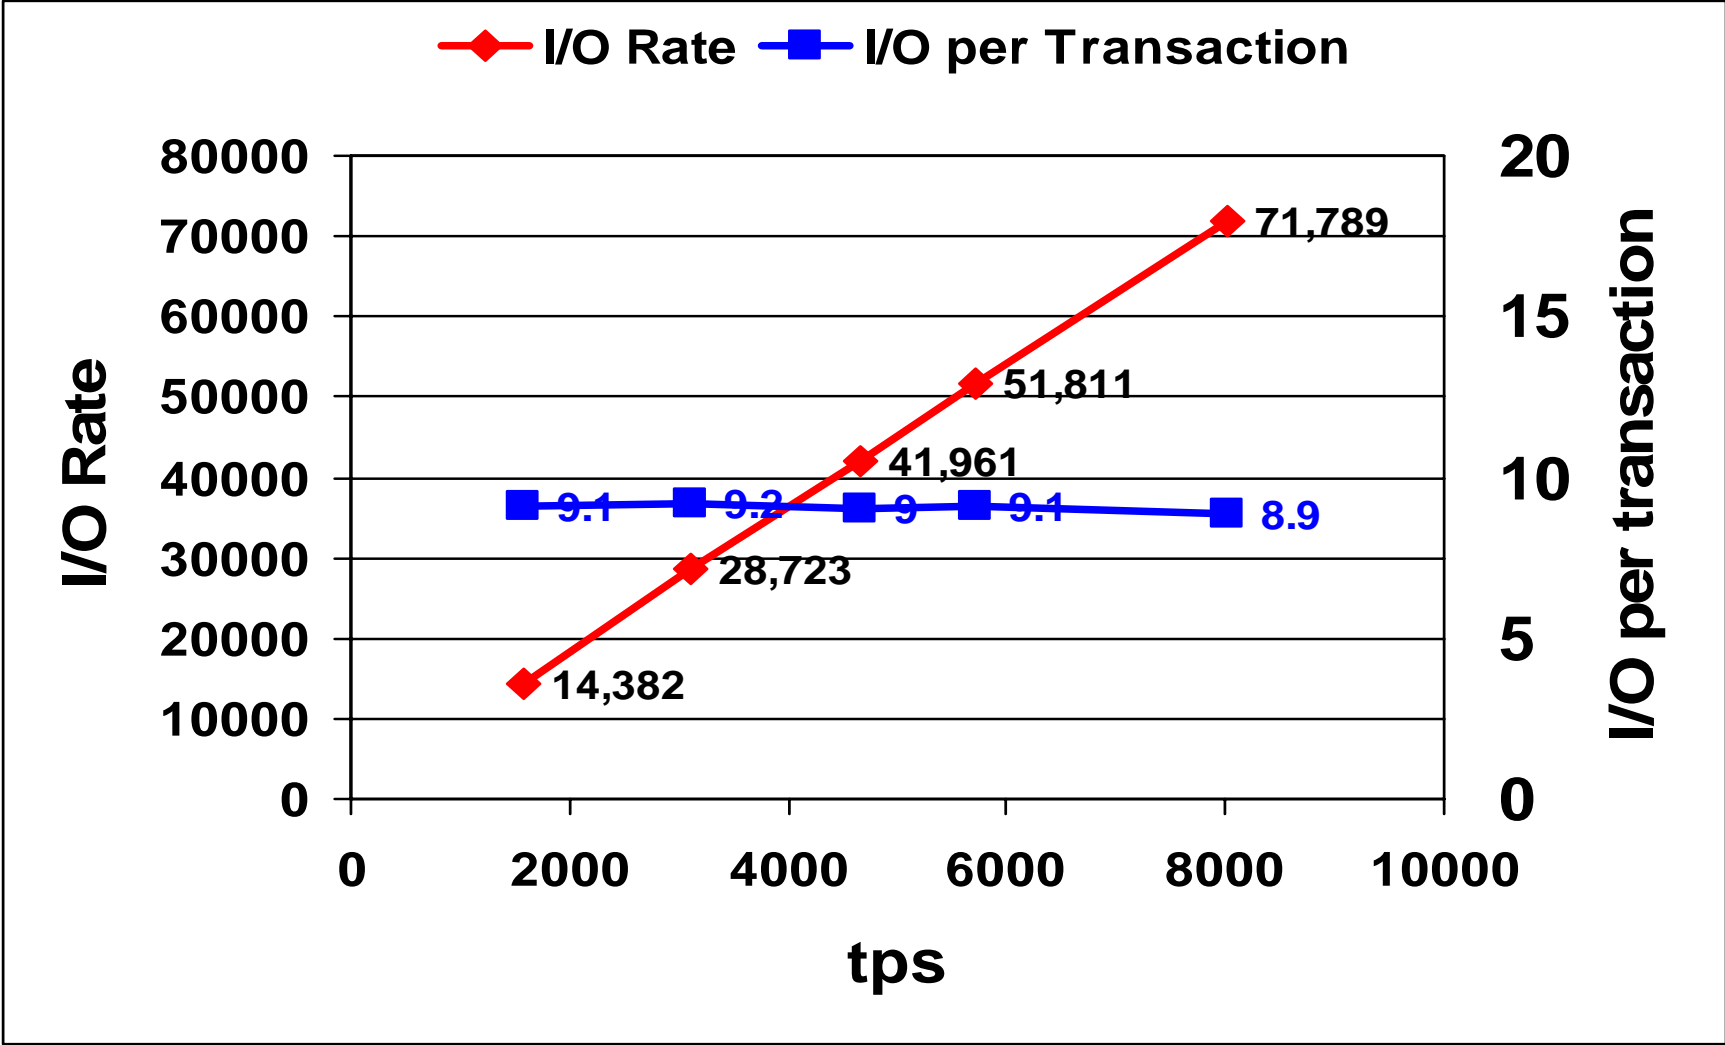
IBM Mainframe Offers Superior I/O Bandwidth Throughout the System

Comparison of mainframe with distributed

Per Processor Capacity
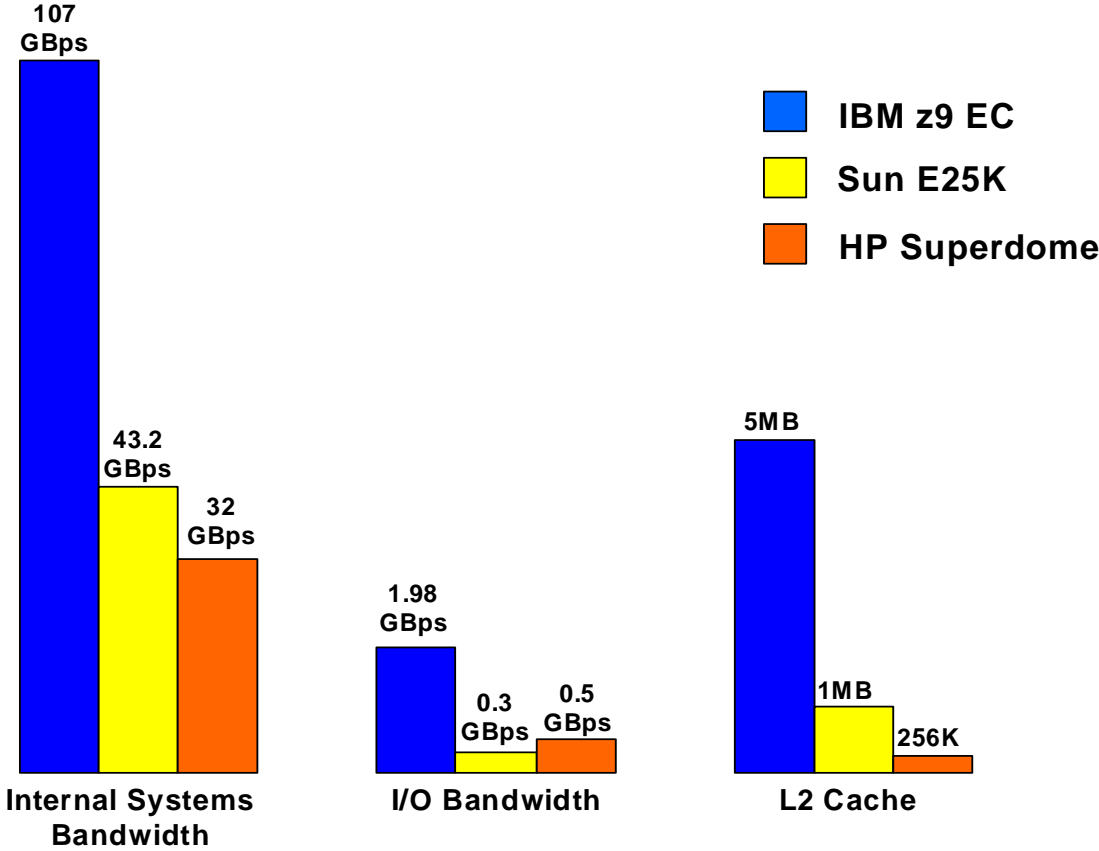
Legend:
- IBM z9 EC
- Sun E25K
- HP Superdome

Internal Systems Bandwidth: 107 GBps, 43.2 GBps, 32 GBps
I/O Bandwidth: 1.98 GBps, 0.3 GBps, 0.5 GBps
L2 Cache: 5MB, 1MB, 256K

Fundamental Capabilities that make System z a Better Platform

Key customers have decided that System z is the best platform for critical business operations. Why?

- Easy to grow capacity rapidly (Scalability)
- Continuous operations (Availability)
- Flexible use of computing capacity (High Utilization)
- Ease of management

How Important is High Availability? Fractional Improvements Make a Difference

Example 1: Financial Services Company

- ▶ $300B assets, 2500+ branches, 15M customers
- ▶ Retail banking, loans, mortgages, wealth management, credit cards
- ▶ CRM System – branches, financial advisors, call centers, internet
- ▶ Number of users – 20,000+

	Unix/ Oracle	zSeries/ DB2
Availability %	99.825%	99.975%
Annual outage	15h 20m	2h 11m
Cost of Downtime	**$45.188M**	**$3.591M**

Sources: ITG Value Proposition for Siebel Enterprise Applications, Business case for IBM eServer zSeries, 2004 & Robert Frances Group, 2005

Financial Impact of Downtime Per Hour

Industry segment	Cost
Energy	$2,818K
Telecommunications	$2,066K
Manufacturing	$1,611K
Financial	$1,495K
Information Technology	$1,345K
Insurance	$1,202K
Retail	$1,107K
Pharmaceuticals	$1,082K
Banking	$997K
Consumer Products	$786K
Chemicals	$704K
Transportation	$669K

Why Does the Mainframe Achieve Better Availability?

Distributed Production

Distributed Development & Test

Distributed Batch
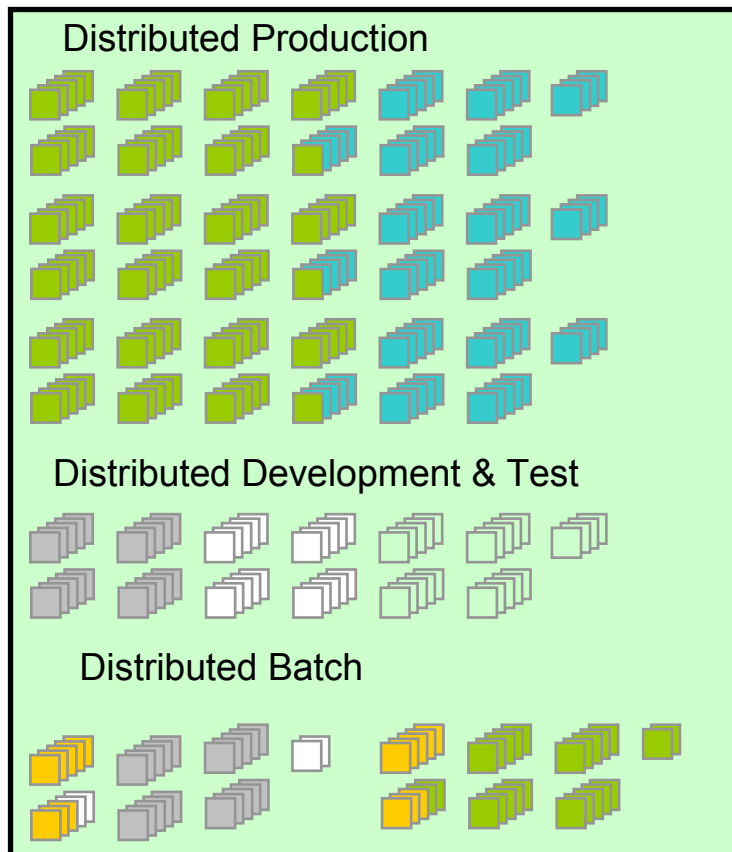
Mainframe

- Too many piece-parts to achieve comprehensive reliable and serviceability (RAS)
- Even harder with 3 different vendors responsible for overall RAS

- Comprehensive approach to RAS
- Software designed to leverage sysplex failover
- Continuous availability is often achieved

A Comprehensive Design for RAS Provides "Five 9's" Availability

- System z comprehensive approach to hardware reliability and serviceability
  - ▶ Measurement and analysis of failure history during design
  - ▶ Error prevention before shipment
  - ▶ Error detection during runtime
  - ▶ Error correction while running
  - ▶ Problem determination to decide if maintenance is required
  - ▶ Ongoing maintenance capability around the world
  - ▶ Change out parts while running

Example of IBM Mainframe Comprehensive Design for Hardware RAS

Error detection
- ▶ Instantaneous
- ▶ Error domain
- ▶ Data capture

Recovery
- ▶ Fence
- ▶ Reset
- ▶ Fault tolerance

Error prevention
- ▶ Technology
- ▶ Design
- ▶ Test

Reliable Operations

Measurements & analysis
- ▶ Product RAS
- ▶ Failure analysis

Problem determination
- ▶ Basic and extended
- ▶ System reference code
- ▶ Problem correlation
- ▶ Problem isolation

Change management
- ▶ Hardware
- ▶ Drivers

Preventive Maintenance
- ▶ Maintenance
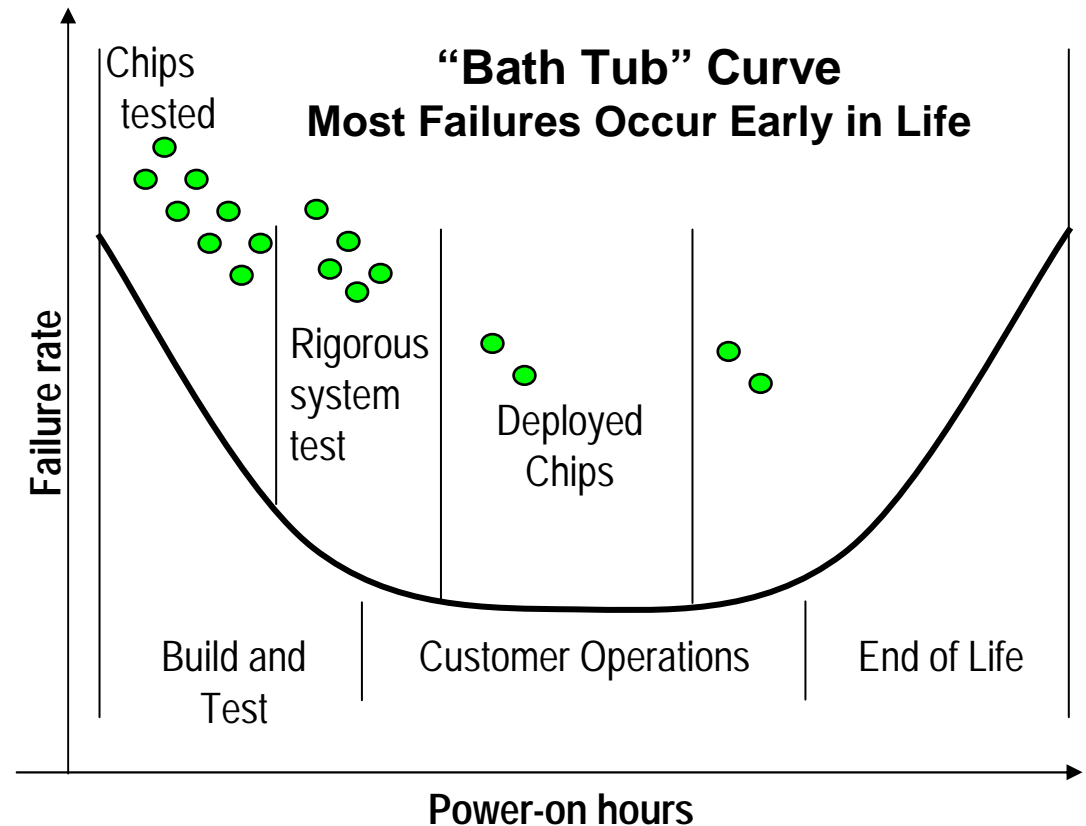- ▶ Service personnel
- ▶ Parts
- ▶ Problem management

Change management examples:

- ■ Replace parts while system is running

Mainframe Chips are Stress Tested to Ensure the Highest Availability

- **All chips and subsystems are stress tested under extreme conditions**
  - ► Millions of test patterns
  - ► Burn in at 150% of the typical voltage

- **Designed to detect and remove defective chips as early as possible in life cycle**

- **Chips have 10X more reliability than other industry standard components when shipped**

"Bath Tub" Curve
Most Failures Occur Early in Life

Chips tested

Rigorous system test

Deployed Chips

Failure rate

Build and Test

Customer Operations

End of Life

Power-on hours

Mainframe Service Processor Monitors Failure Detections

- Real-time monitoring of recoverable errors
  - ▶ Memory
  - ▶ Caches
  - ▶ Internal Bus Parity Errors
  - ▶ Errors are detected and reported before they become "hard" errors
- Redundant Internal Paths and components
- Automatic "Phone Home" capability to alert IBM to potential problems before they cause an outage
- Hot-pluggable replacement of most major components – no downtime required for upgrades and repairs in many cases

Processor	Fault Isolation Register	Channels		
L1 Cache			Service Processor	Alert / Phone Home
L2 Cache			Error Log	
Memory				
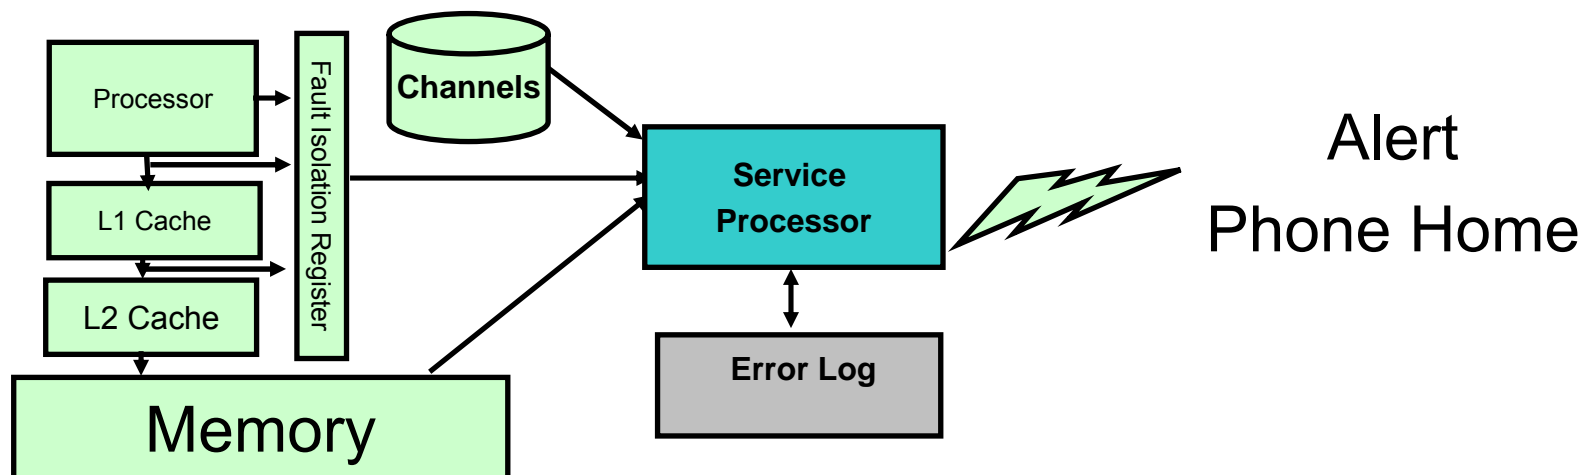
Mainframe Service Processor Replaces Failed Processing Unit with Pre-Installed Spare

- Processing Units (PU) not required
  for a specified configuration remain dormant

- Spare PUs can be utilized for concurrent replacement or for
  concurrent capacity upgrades

- If there is a processor failure, sparing first searches for a spare PU on
  the *same book*; next on an *adjacent* book on the ring; or on *any*
  available spare PU in the server

- If there are no spares, the processor availability feature recovers the
  application to another PU and makes a service call automatically

- Redundant paths to cache and memory

Preventive Maintenance: Backed by People and Parts Capabilities

- IBM provides a comprehensive parts-stocking system to ensure fast access to all mainframe server parts.

- The problem tracking is managed through RETAIN

- The system is managed on the basis of the criticality of each part to server operation, the expected replacement rate of those parts, and geography (location of server types, models, and features).

- Service representatives can **dial in to the failed server** to perform remote diagnostics.

- If the onsite CE requires additional support, the service structure provides 3 additional levels of **24 hours support** - **worldwide access** to error data, logs, and traces as well as to history data and problem repair data.

- The removed parts are returned for failure analysis

Hardware Maintenance While System is Running

Capability	System z9
ECC on Memory Control Circuitry	Transparent
Oscillator Failure	Transparent
Microcode Driver Updates	Replace While Running
Book Replacement	Replace While Running
Memory Replacement	Replace While Running (Book Offline)
Memory Bus Adapter (MBA) Replacement	Replace While Running Connectivity to I/O Domain remains
Self Timed Interface Failure to I/O	Replace While Running Connectivity to I/O Domain remains
Processor Upgrades	Replace While Running
Physical Memory Upgrades	Replace While Running
I/O Upgrades	Replace While Running
Spare CPU's	2 Pre-installed per System

Availability Built in by Design

"The IBM mainframe platform retains industry-leading availability characteristics even for single-system instances.

"For example, standard service availability commitments … suggest that the mainframe is delivering 99.9% to 99.99% scheduled availability versus 99.5% for distributed server platforms in non clustered configurations."

Source: *Forrester, 2005 Mainframe Market Outlook, February 4, 2005*

Parallel Sysplex provides Availability

- **All machines in the sysplex are able to access all the shared disks**
    - ▶ maintains availability if a machine or operating system image goes down

- **Patch/upgrade various subsystems (CICS, DB2..) non-disruptively**
- **Multiple versions of a subsystem can coexist together (e.g.: DB2 7 and 8)**
- **The work of a failed subsystem instance can be taken over by any other instance in the sysplex**
    - ▶ The work of a failed TOR or AOR in CICS can be taken over by any other TOR or AOR in the sysplex
    - ▶ The work of a failed WebSphere MQ instance can be taken over by any other WebSphere MQ instance
    - ▶ If a DB2 instance fails new work is modeled by the others
        - − In flight work is recovered by restarting a failed DB2 instance somewhere else
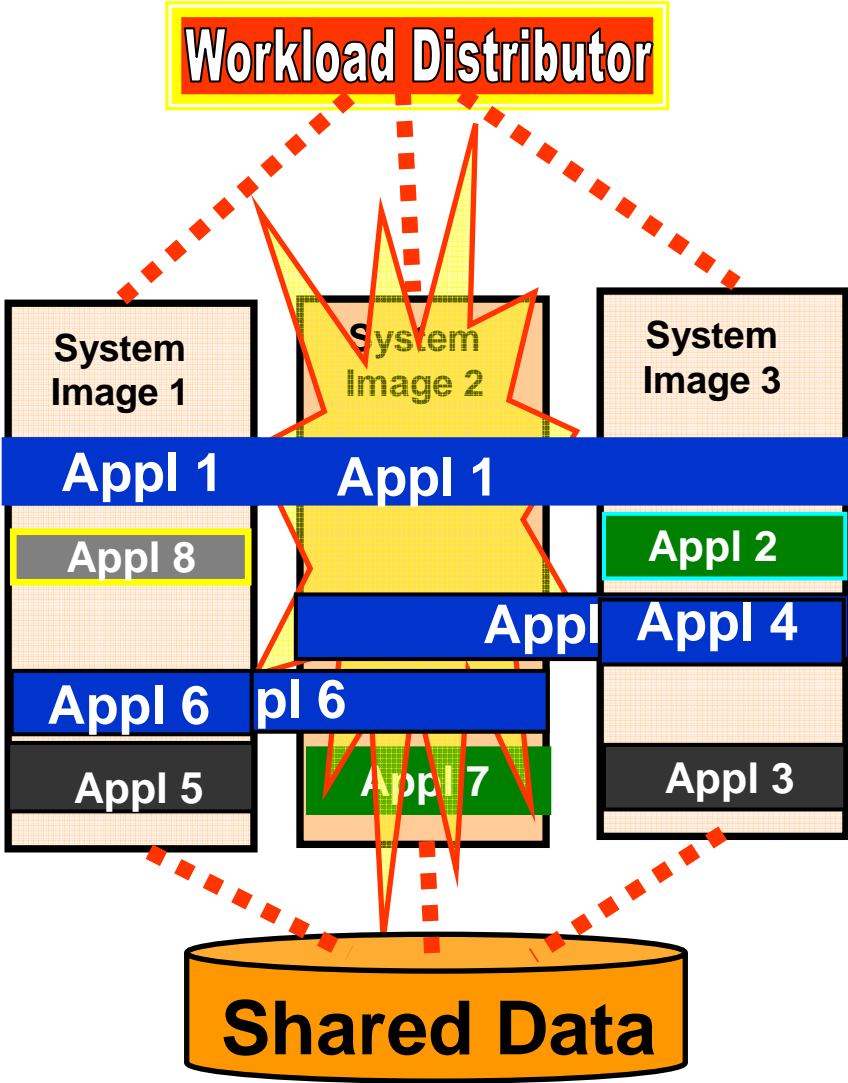
Sysplex Shared Data Enables Failover
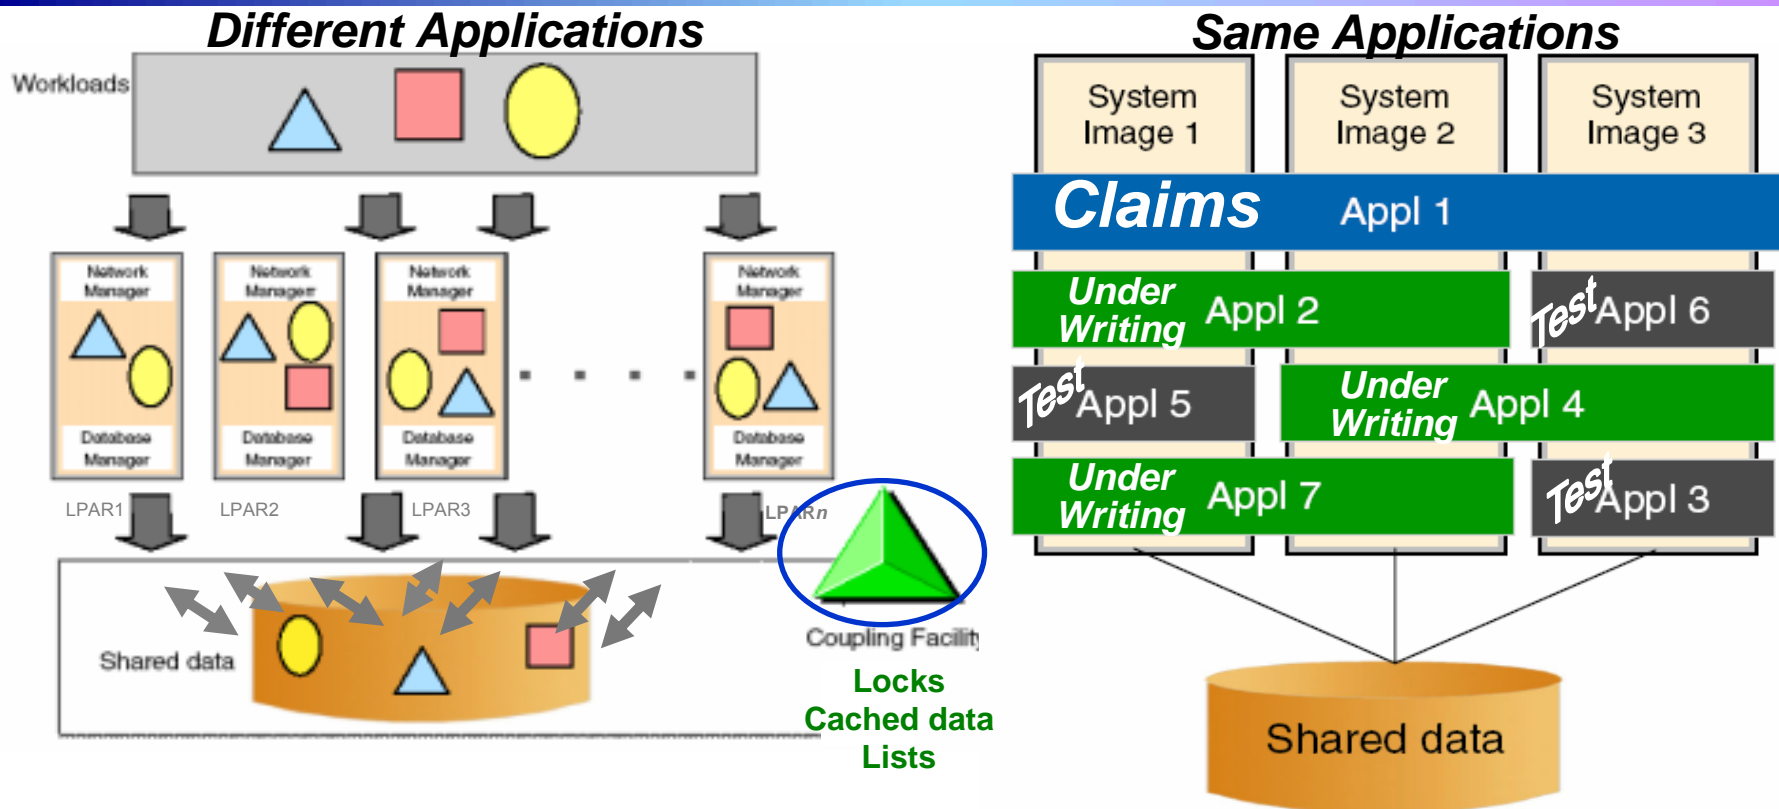
- **The Coupling Facility* provides the technology to shared data concurrently amongst different or similar workloads**
- **Data sharing provides for continuous availability and reliability**
- **CICS, IMS, DB2, MQ, VSAM, exploit the coupling facility for shared data**

Sysplex Failover

Cloned Application

Workload Distributor

System Image 1

System Image 2

System Image 3

Appl 1

Appl 1

Appl 8

Appl 2

Appl

Appl 4

Appl 6

pl 6

Appl 5

Appl 7

Appl 3

Shared Data

Nno move data
Maintenance
Preserve data / fail
Data validity

Continuous Operation is Not Just a Hardware Responsibility

- IBM middleware designed to exploit sysplex clustering for continued operation in the event of a system image failure

- IBM software designed to avoid shutdowns while applying maintenance updates
  - All Parallel Sysplex products including DB2, IMS, IRLM, VSAM, and CICS are designed for rolling update in a cluster

CICS Keeps Operating While Recovering Failed Regions

- AOR – Application Owning Region – runs business logic code
- TOR – Terminal Owning Region – runs presentation code
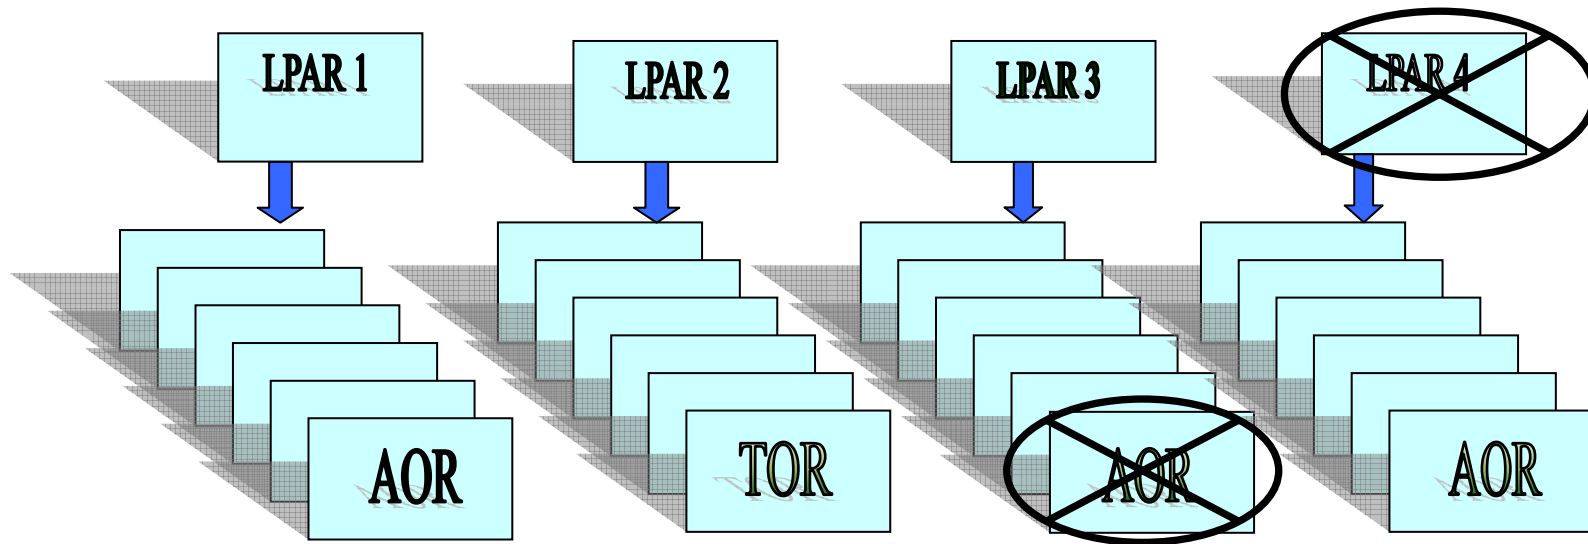
Large Asian Bank Benchmark Demonstrates CICS Resiliency from Failures

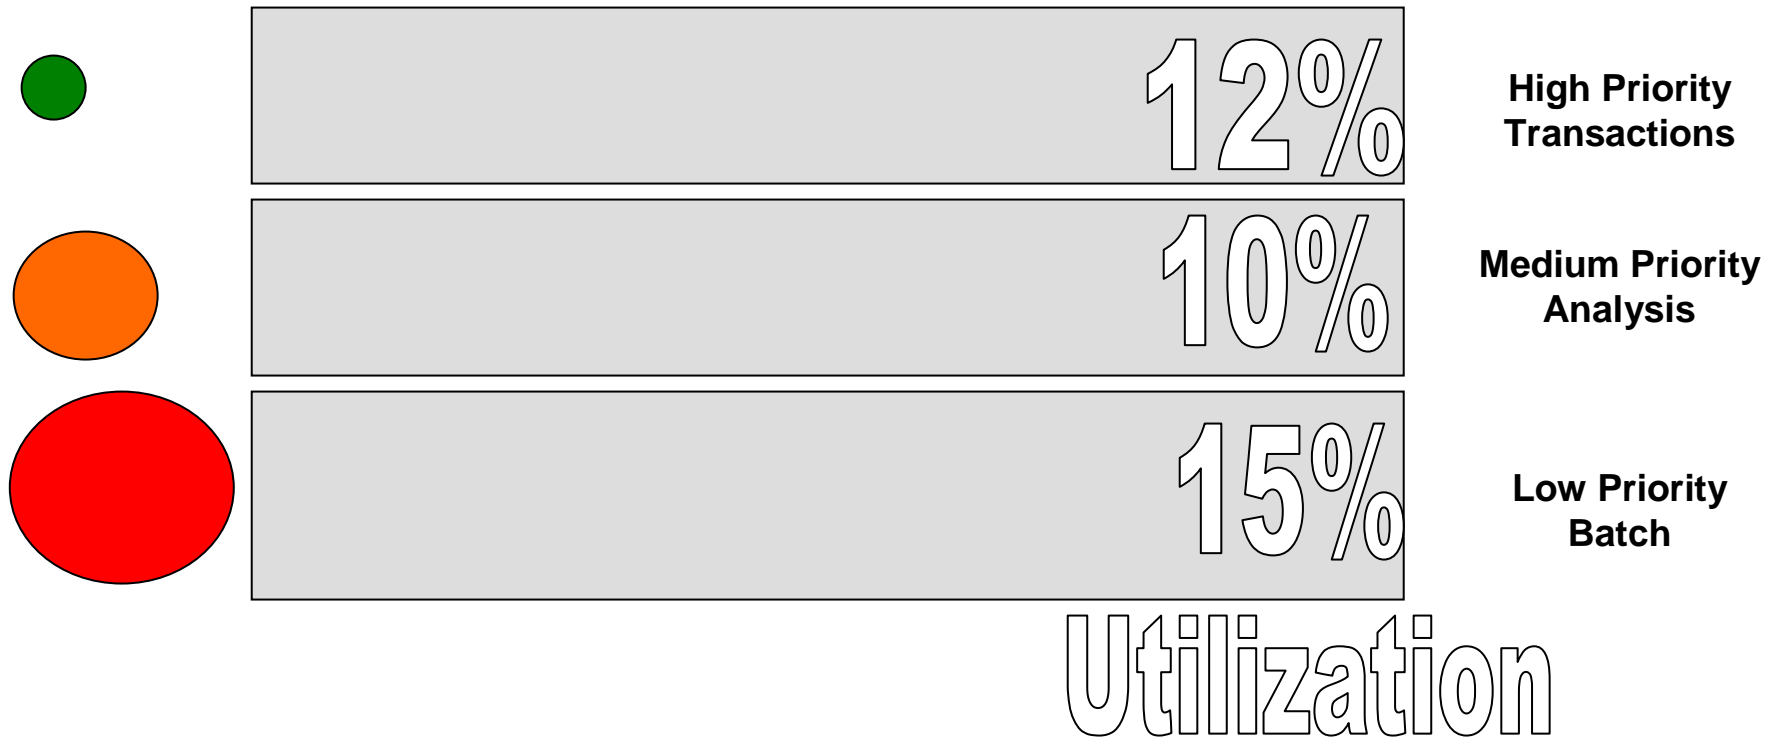|	INDUCED AOR FAILURE	INDUCED TOR FAILURE	INDUCED LPAR FAILURE
DOWNTIME	0	0	0
Steady State tps	6250	5000	6271
DB2 Recovery Time	N/A	N/A	8mins
Temporary increase in Response Time	64ms	70ms	76ms

Fundamental Capabilities that make System z a Better Platform

Key customers have decided that System z is the best platform for critical business operations. Why?

- Easy to grow capacity rapidly (Scalability)

- Continuous operations (Availability)

- Flexible use of computing capacity (High Utilization)

- Ease of management

Distributed Servers

Purchased New Machines each running separate workload

🟢	**12%**	**High Priority Transactions**
🟠	**10%**	**Medium Priority Analysis**
🔴	**15%**	**Low Priority Batch**

Utilization

X **Purchasing separate machines had higher administration**

X **Along with lower utilization**

X **Support problems did not go away**

Flexible to Respond to Change

- Mainframe Workload Management
  - ► Policy Based (Classify Work by business value)
  - ► All resources pooled
  - ► Dynamic changes can be made flexibly

Online Transactions | **Batch Processing**

Resources easily shift to changing workloads
e.g. online vs batch

Balanced

Transaction "heavy"

Data "heavy"
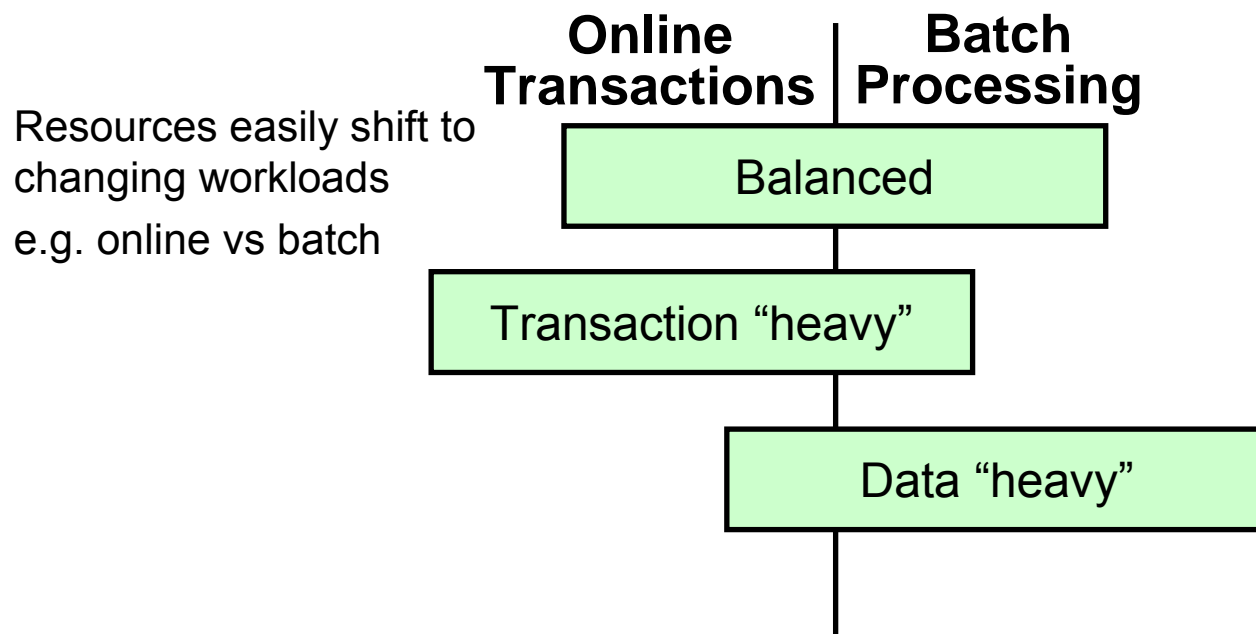
Mainframe Policy Driven Workload Management

Casual customer Low priority

Gold customer High Priority

Business Partner High Priority

Data Analysis (Best can do)

Customer care High Priority (Business Hours)

Classification Rules

Response Time
80% trx complete with
0.5 seconds – (24x7)
(Claims Trx Class)

Average Response Time
All trx must complete witin
4 seconds – (18x5)
(Policy Browser)

Average Response Time
All trx must complete
within 2 seconds
Loc: MA, NY, PA USA

A. Discretionary Response Time
between 7:00 – 16:00
B. Response Time percentile
50% trx complete within 3 mins
between 16:00 – 7:00

Response Time (8 x 5)
A. South East = 0.8 Secs
B. South West = 1.5 Secs.
C. North West = 2.5 Seconds
D Nort East = 2.5 Seconds

Service Class

Claims

Policy

Browser

B P

UnderWriter

CustCare

Mainframe Workload Management

- Monitoring the workloads of various users and applications

- Monitoring system-wide resources to determine whether they are fully utilized

- Inhibiting the creation of new user workloads when certain shortages of resources exist

- Dynamically adjusting resources to meet service level objectives

- Change the priority of users automatically to adjust the consumption of system resources

- Selects the resources to be allocated, if a choice of resources exist to balance the executing workloads
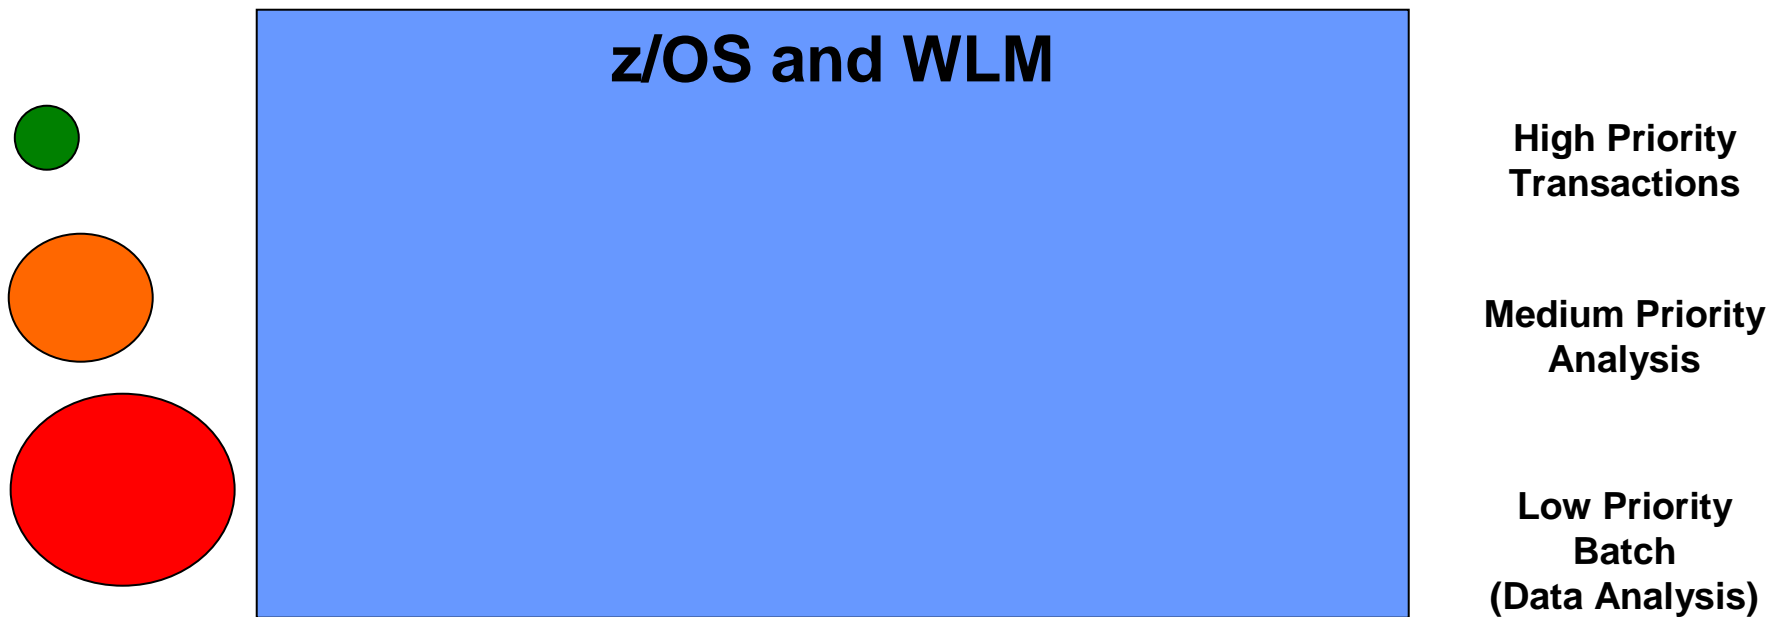
Mainframe Workload Management

Workloads deployed to z/OS can be differentiated and prioritized based on business policy, and managed to meet Service Level Agreements

Resources are automatically allocated, adjusted and reallocated to meet objectives

WLM will manage LPARs, CPUs, channels, I/O subsystems and DASD, TCP/IP connections, servers, etc.

Enables 100% utilization of capacity

z/OS and WLM

High Priority Transactions

Medium Priority Analysis

Low Priority Batch (Data Analysis)

Transaction type:
- Web "buy" vs "browse"
- B2B
- Batch payroll
- Test

User/user type:
- Top 100 clients
- Typical clients
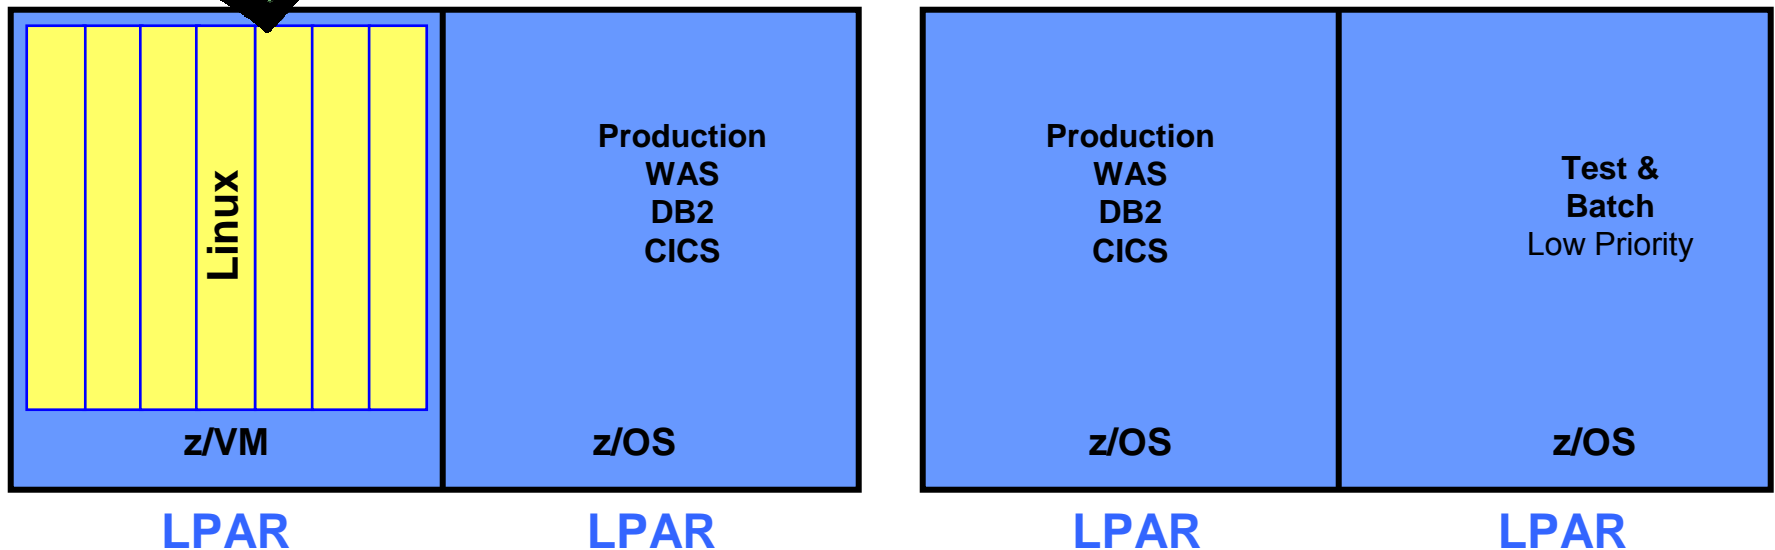- Executive
- Design team

Time periods:
- 1AM - 4AM
- Mon - Fri
- Weekends
- End of quarter

Workload Management Prioritizing Across Images in a Server

"Intelligent Resource Director (IRD)" further differentiates z/OS with its ability to manage resources across multiple partitions in a server

PR/SM, IRD and WLM work together to ensure that the resources of the server are correctly balanced to enable work to complete within stated policy goals

MORE

CF

Linz/VM	Production WAS DB2 CICS	Production WAS DB2 CICS	Test & Batch Low Priority
z/VM	**z/OS**	**z/OS**	**z/OS**
LPAR	**LPAR**	**LPAR**	**LPAR**

Processor resources, data bandwidth and I/O queueing decisions are perfectly balanced across the server to manage diverse workloads within the parameters of stated business goals
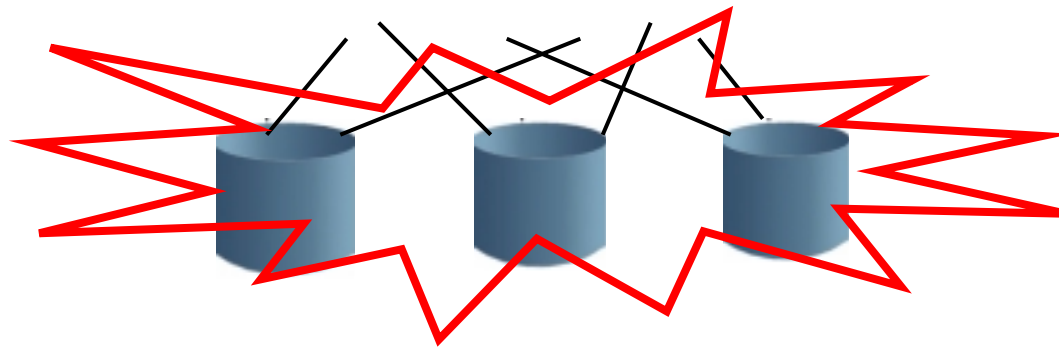
Fundamental Capabilities that make System z a Better Platform

Key customers have decided that System z is the best platform for critical business operations. Why?

- Easy to grow capacity rapidly (Scalability)

- Continuous operations (Availability)

- Flexible use of computing capacity (High Utilization)

- Ease of management

Service Levels Manage Storage

- **What performance objectives are required by data**

- **When and how to backup data**

- **Whether datasets should be kept available for use during backup or copy**

- **How to manage backup copies kept for disaster recovery**

- **What to do with the data that is obsolete or seldom used**

Rule-based Policy Management
to Manage the Storage Process

Gold customer
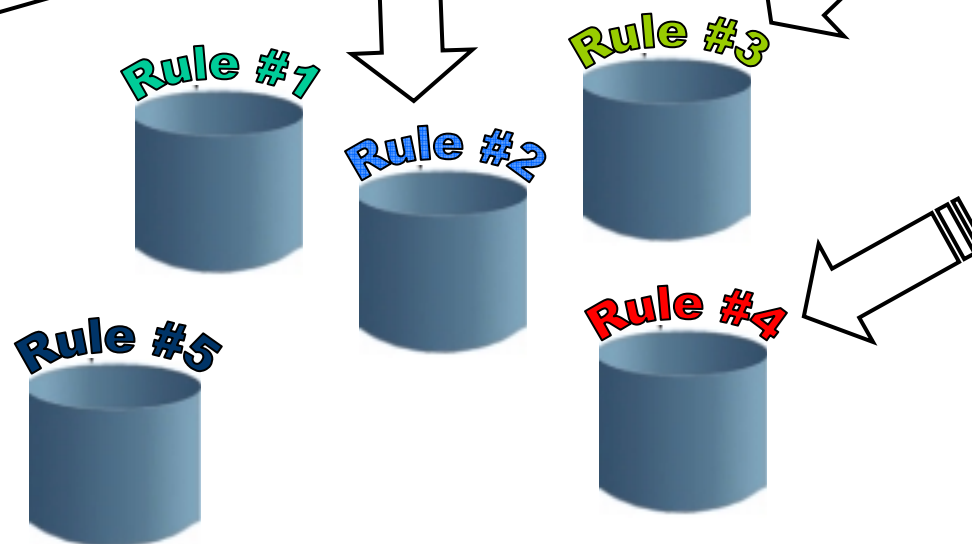High Priority

Casual customer
Low priorty

Business Partner
High Priority

Customer care
High Priority
(Business Hours)

Data
Analysis
(Best can do)

Rule #1

Rule #2

Rule #3

Rule #4

Rule #5

Rule-based Policy Management to Manage Backup/Restore Automatically via Hierarchical Storage

Gold customer High Priority

CritSit !

Business Partner High Priority

Data Analysis (Best can do)

Rule #1

Rule #2

Rule #4

Secondary Storage

Secondary Storage

Secondary Storage

Third Storage Media

Tape
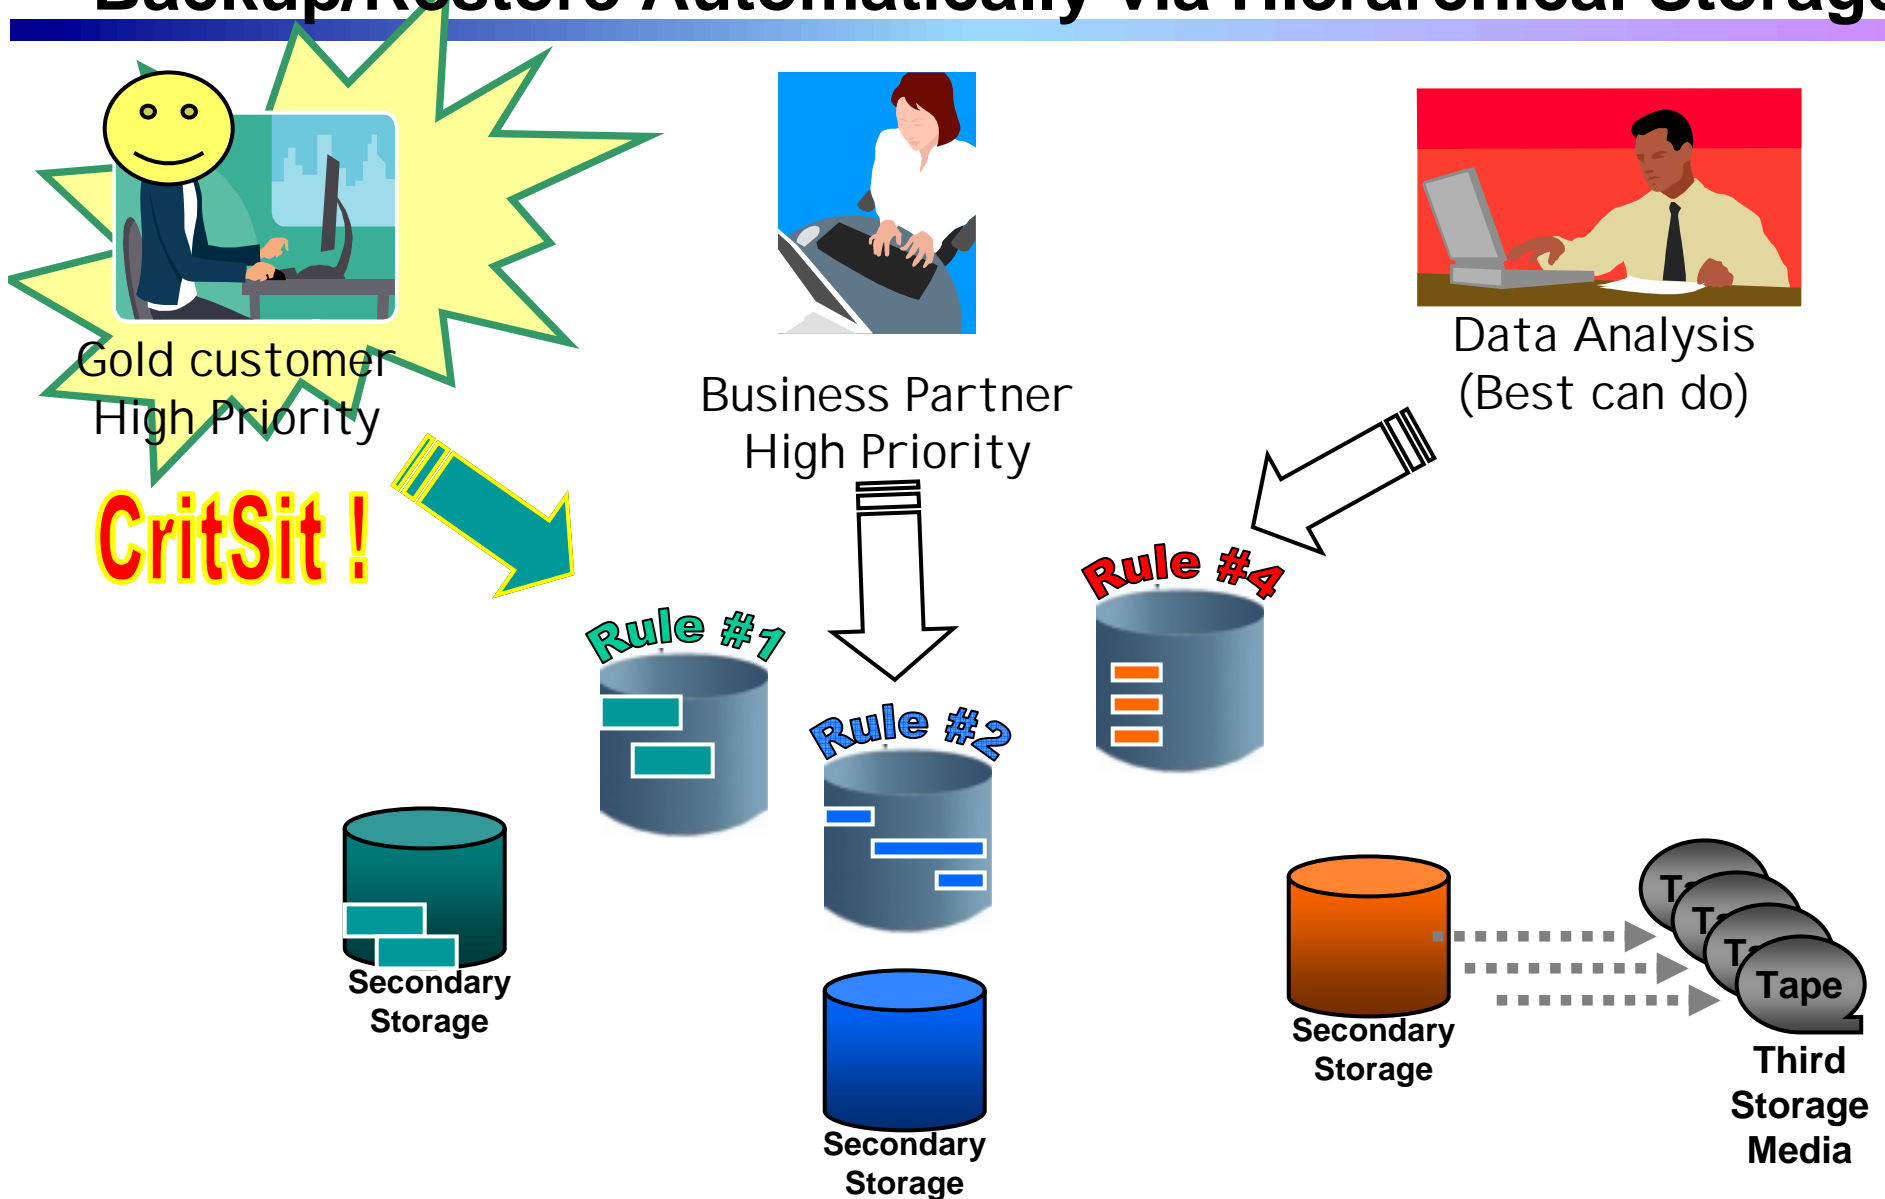
Simply The Best Qualities of Service!

ODI's CIO

Fundamental Capabilities that make System z a Better Platform

Key customers have decided that System z is the best platform for critical business operations. Why?

- Easy to grow capacity rapidly (Scalability)

- Continuous operations (Availability)

- Flexible use of computing capacity (High Utilization)

- Ease of management