System z Enables Solutions For A Smarter Planet

The Rule Of Three

Quiz

- What is the Rule of Three?

  a) A form of government in the Soviet Union

  b) Two's company, three's a crowd

  c) The Three Stooges in charge

  d) A rule to help you recognize when a smarter planet solution will cost less to deploy on the mainframe

The True TCO Of The Mainframe

Every smarter planet solution you talked about today costs less to deploy on System z. Is that for real?

Service Oriented Finance CEO

Yes and we can give you a rule of thumb to help recognize when a smarter planet solution will cost less on System z.

The Rule Of Three

- The cost of deploying a new application will usually be less on a mainframe if:

1. **It is an incremental workload on an existing mainframe**
2. **It can make use of a specialty processor**
3. **Disaster recovery is required**

A Short Primer On Key Mainframe Concepts

- Incremental workloads
- Specialty processors
- Disaster recovery

The System z10 Frame Contains Many Processors

64 way SMP

PLUS: 11 Dedicated System Assist Processors

PLUS: 2 Spare Processors

PLUS:
Max 336 FICON I/O Driver Channels
Or Max 1024 ESCON Driver Channels

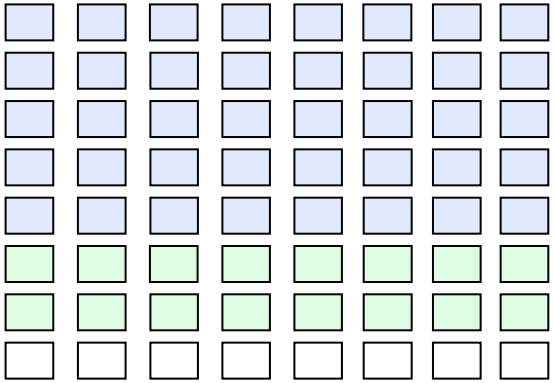

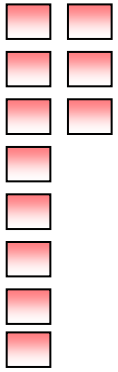

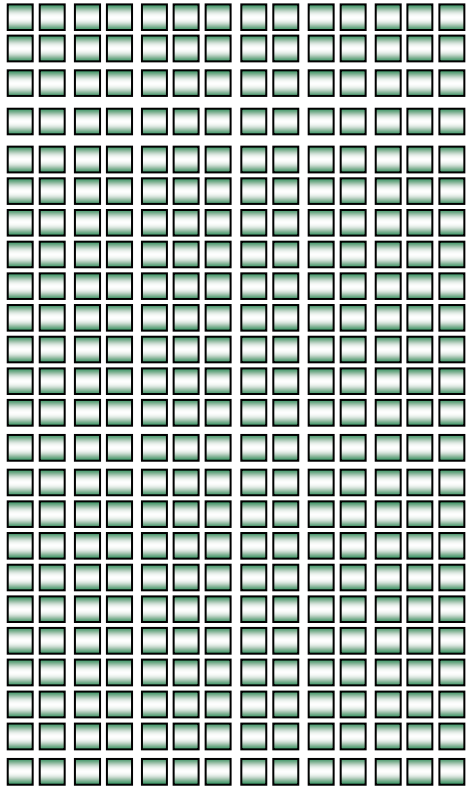

Application Execution

I/O Scheduling

Dedicated I/O Operations

- **Decimal Floating Point Accelerator**
  - Implemented in hardware - one per core
- **CP Assist for Cryptographic Function (CPACF)**
  - Two cores share a CPACF

These Are Used To Run Several Workloads Concurrently

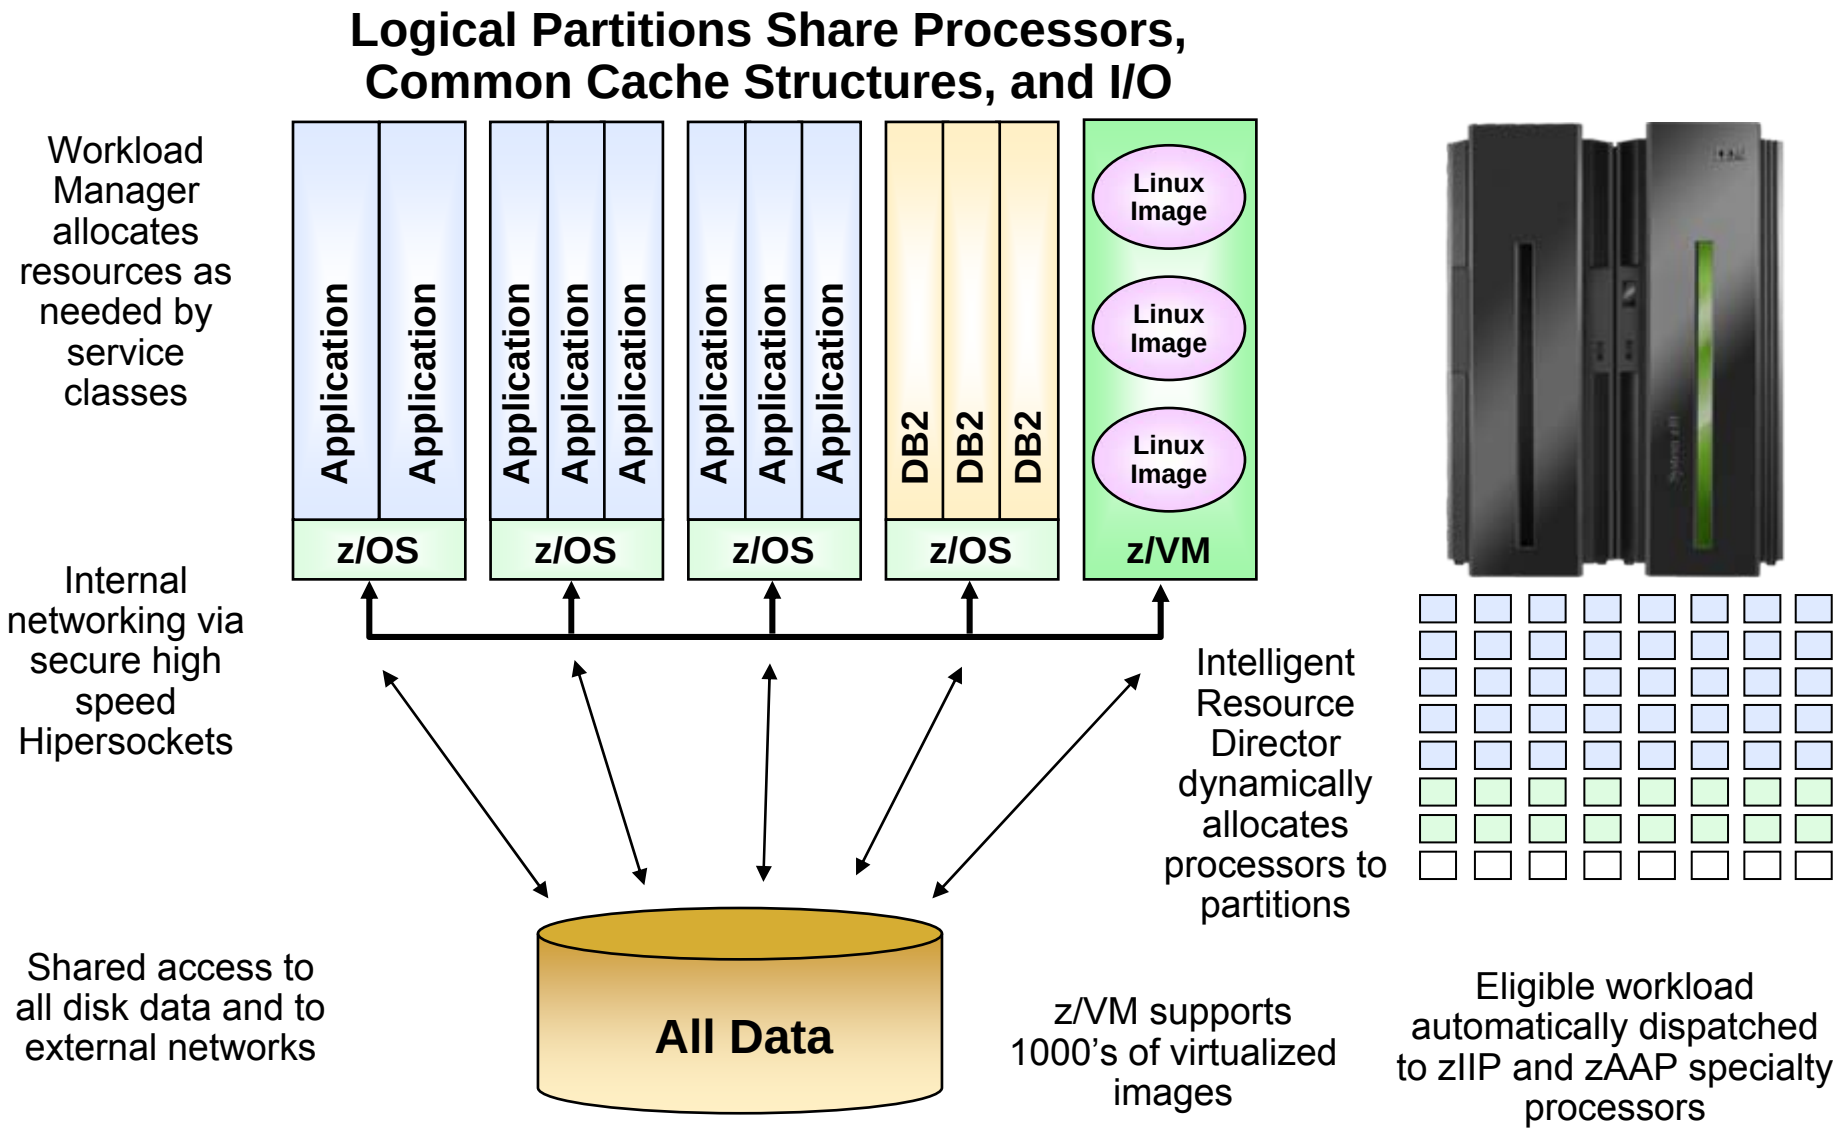

This Is What It Looks Like In Operation

New workloads are incremental to the existing workloads

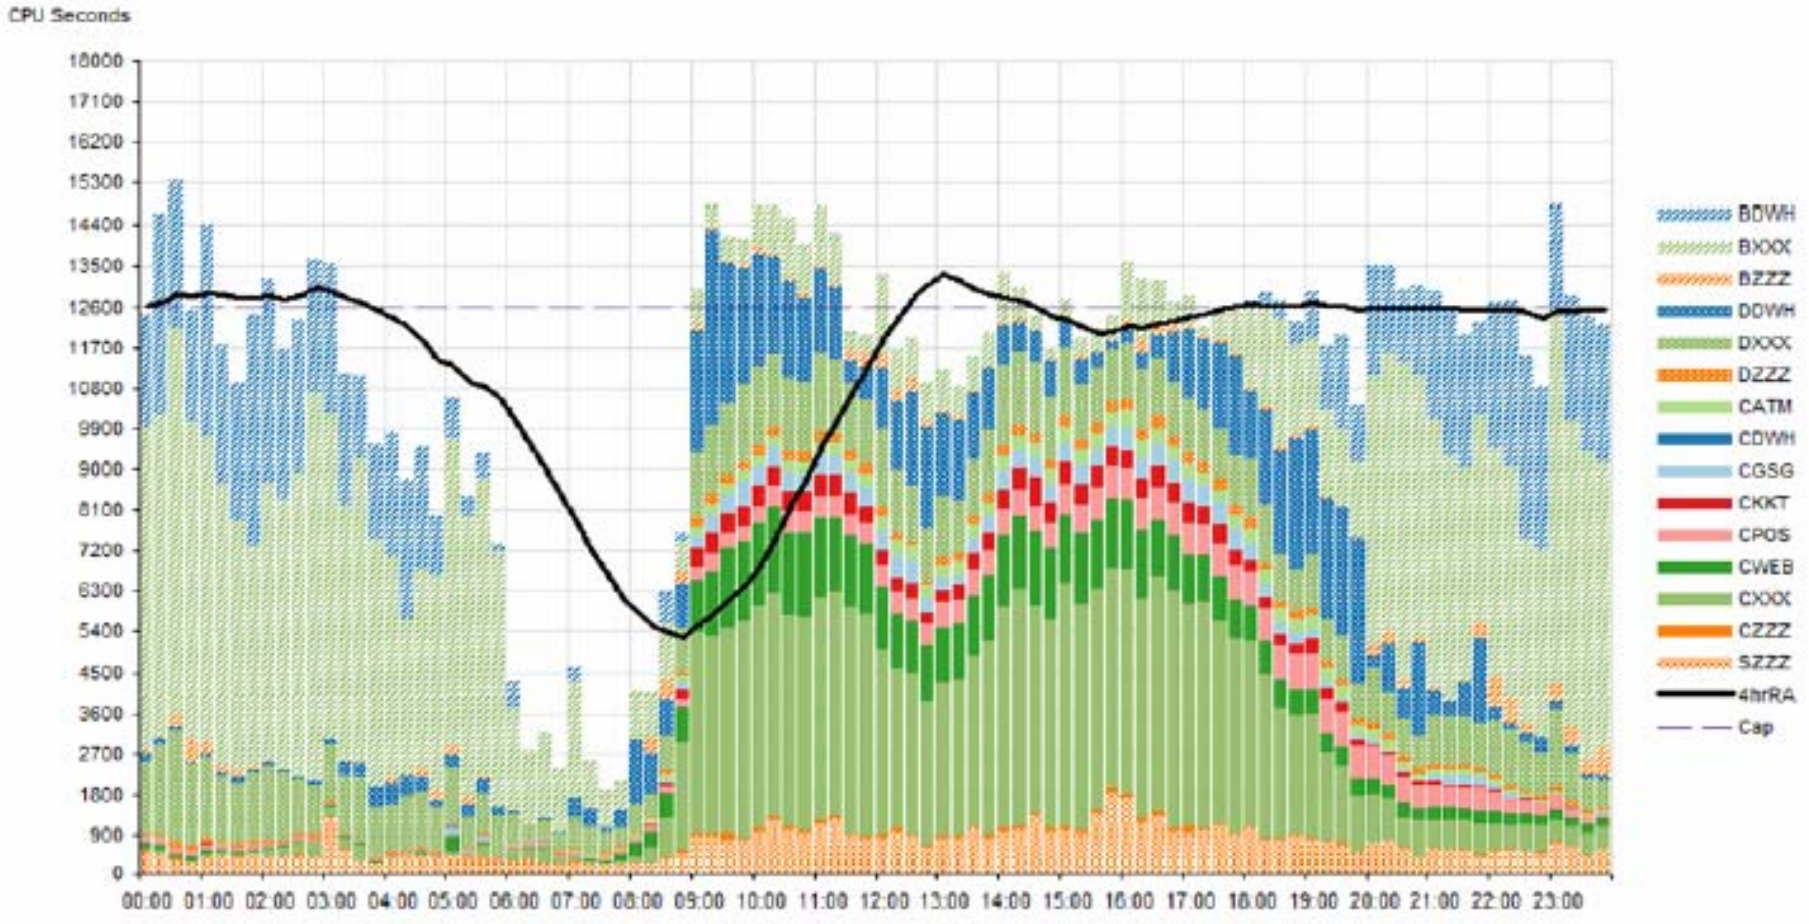

“Specialty Engines” Reduce Cost For New Workloads

- Special assist processors for System z
  - For Java workloads (**zAAP**)
  - For selected DB2 workloads (**zIIP**)
  - For Linux workloads (**IFL**)
- Attractive pricing
  - $125K for a 920 MIP processor (90% discount)
  - No charge for IBM software running on zAAP/zIIP
  - IBM software running on IFL costs 120 PVU’s
  - Free upgrade to next generation!

- Requirements
  - Max number of zAAP =< number of general purpose processors
  - Max number of zIIP =< number of general purpose processors
  - No limit on the number of IFL’s

Disaster Recovery – Fast Failover For Less

Primary Site
64 way SMP

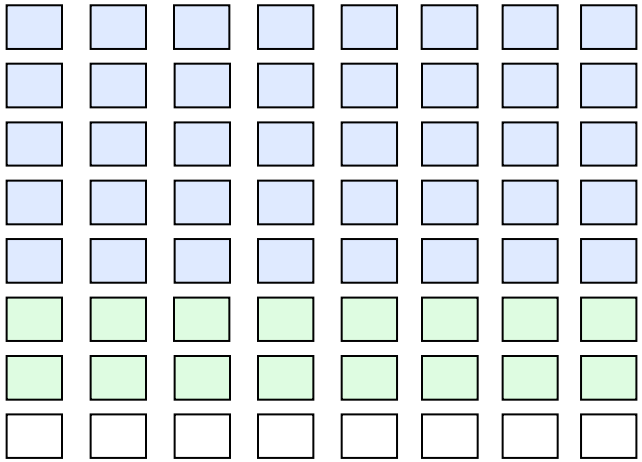

Capacity back up on demand
Pay regular price for one active processor

All other dormant processors discounted 98% ($7K per year for general purpose, $2K per year for specialty engines

Alternative Site
64 way SMP

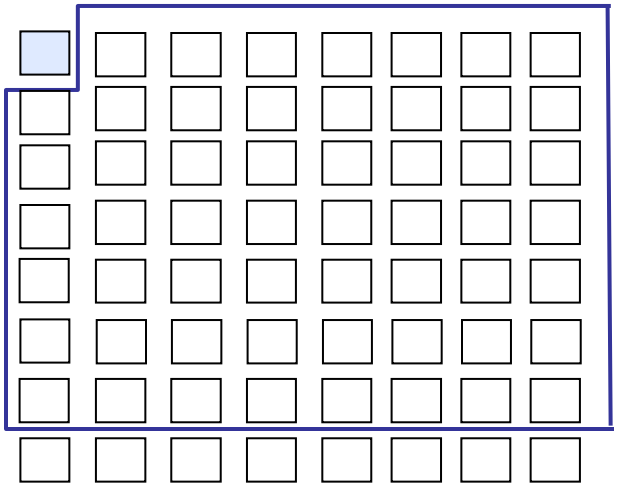

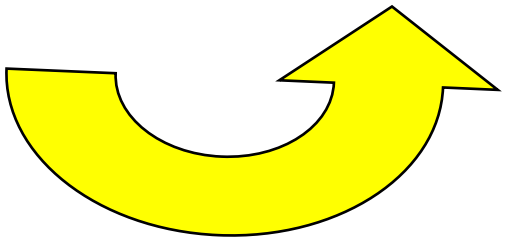

Site Failover With GDPS

Pay regular price for frame

TCO Top Down Methodology

IBM

TCO Top Down Methodology

1. Establish Equivalent Configurations

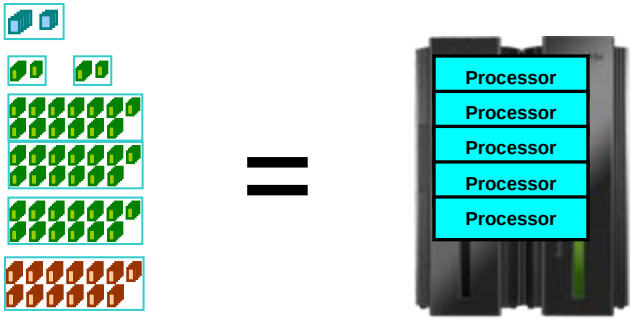

2. Price out Total Cost of Acquisition

3. Add cost of labor and environmentals

TCO Top Down Methodology

1. Establish Equivalent Configurations

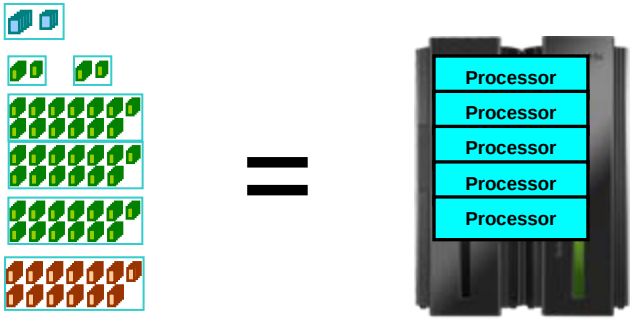

2. Price out Total Cost of Acquisition

3. Add cost of labor and environmentals

Banking Benchmark Comparison

- **Asian Bank**
  - IBM System z9 and DB2
  - TCS BaNCS
  - **15,353** Transactions/second
  - **50** Million Accounts
  - IBM benchmark for customer

- **Bank of China** **
  - IBM System z9 and DB2
  - TCS BaNCS
  - **9,445*** ** Transactions/second
  - **380** Million Accounts
  - IBM benchmark for customer

- **HP/Temenos** *
  - HP Itanium
  - **Temenos T24**
  - **2,153** Transactions/second
  - **13** Million Accounts
  - Largest banking benchmark performance claimed by HP

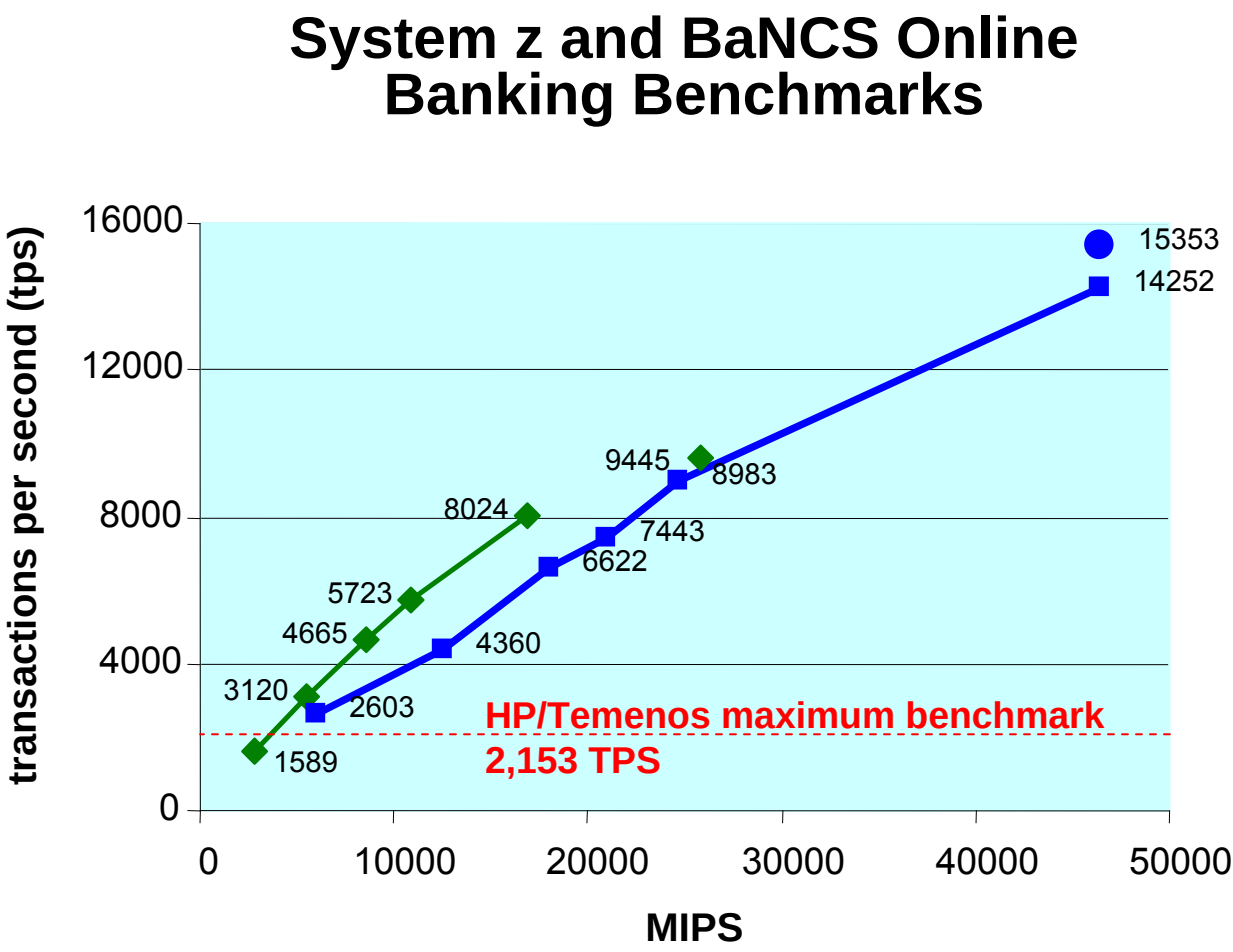

* **SOURCE**: TEMENOS BENCHMARKS; http://h71028.www7.hp.com/enterprise/downloads/TemenosBenchmark.pdf
** **SOURCE**:http://www.enterprisenetworksandservers.com/monthly/art.php?2976 **Source**: InfoSizing FNS BANCS Scalability on IBM System z – Report Date: September 20, 2006
*** Standard benchmark configuration reached 8024 tps, a modified prototype reached 9445 tps

Compare The Processors Needed To Achieve 2,200 Transactions Per Second

Online Injector: 1 x HP RX7620

Temenos T24 Servers:
2 x HP RX7620
3 x HP 9000 Superdome

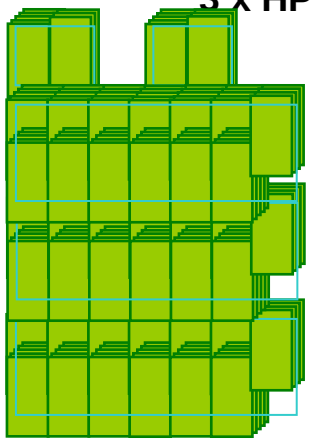

Oracle 10g: 1 x HP 9000 Superdome

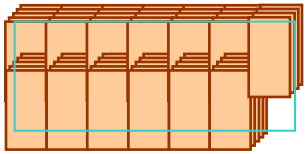

5 processors

(3,906 MIPS)

280 processors

(457,762 Performance Units)

$26.0M
TCA (3yr)

TCS BaNCS and DB2
1x z10 2097-705

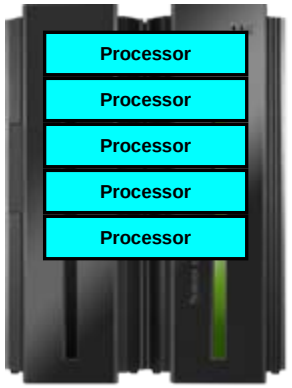

$18.9M
TCA (3yr)

117 Performance Units per MIP

HP Integrity rx7620 - (10U) 1.5GHz 6MB (8ch/8co)

HP 9000 Superdomes - 32W 1GHz 32MB (32ch/64co)

Compare The Processors Needed To Achieve 2,200 Transactions Per Second (with Dev/QA)

Online Injector: 2 x HP RX7620

Temenos T24 Servers:
4 x HP RX7620
6 x HP 9000 Superdome

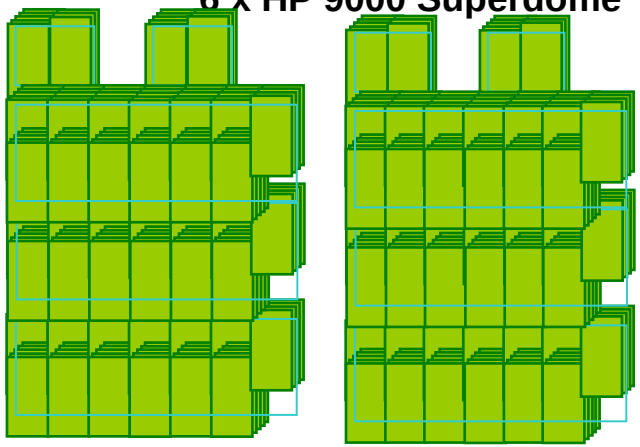

Oracle 10g: 2 x HP 9000 Superdome

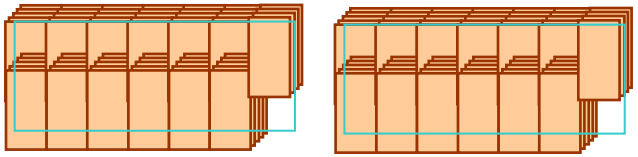

7 processors

(4,906 MIPS)

560 processors

(915,524 Performance Units)

$59.2M
TCA (3yr)

TCS BaNCS and DB2
1x z10 2097-707

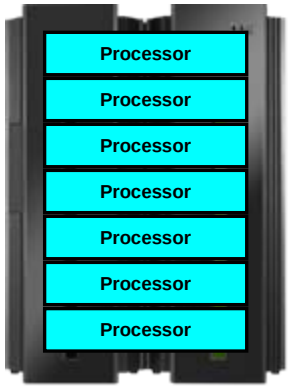

$22.7M
TCA (3yr)

187 Performance Units per MIP

NOTE: Double Distributed Servers, add 1000 MIPS to System z for Dev/QA

HP Integrity rx7620 - (10U) 1.5GHz 6MB (8ch/8co)

HP 9000 Superdomes - 32W 1GHz 32MB (32ch/64co)

Another Customer Case: European Financial Services Offload

- 2x 24-way Production / Dev / Test / Education
  Application, DB, Security, Print and Monitoring
- 4x 1-way Admin / Provisioning / Batch Scheduling

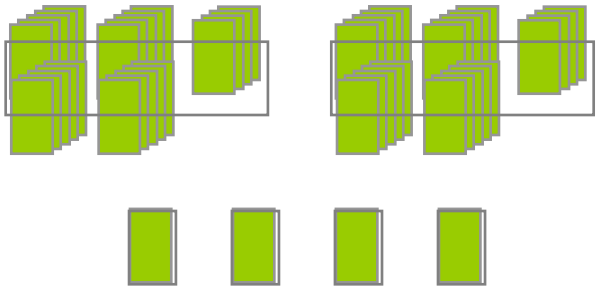

$17.9M TCA (4yr)

52 Unix processors
(222,292 Performance Units)

z890 2-way Production / Dev / Test / Education
App, DB, Security, Print, Admin & Monitoring

2 processors
(332 MIPS)

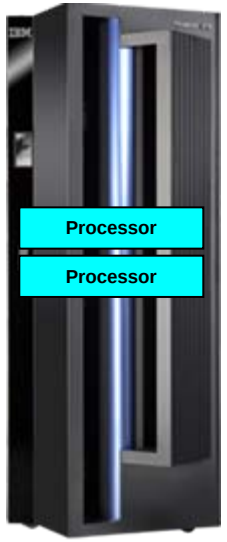

$4.9M TCA (4yr)

Plus:
2x HP SAN Servers (existing)
Many (existing) Windows servers

670 Performance Units per MIP
Disaster recovery not included

Lesson Learned

- It usually takes far more processor cores to deploy on an HP distributed platform
  - Relative Performance Units per MIP have ranged from 87 to 670
  - A typical number is 122
- Performance Unit Capacity for various distributed servers can be found in the Server Consolidation Analysis Report from Ideas International

TCO Top Down Methodology

1. Establish Equivalent Configurations

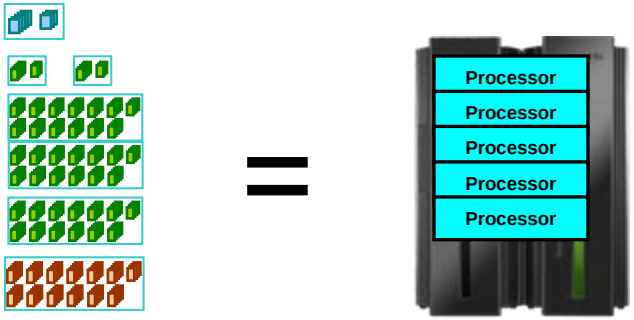

2. Price out Total Cost of Acquisition

3. Add cost of labor and environmentals

Facts To Consider When Pricing Out The Cost Of The New Workload On The Mainframe

- Calculate new workload as an incremental cost to an existing System z
  - LPARs and sub-capacity pricing isolate the incremental cost
- Specialty processors are deeply discounted
- Disaster recovery capacity is deeply discounted
- Incremental costs get cheaper as system grows
- New workload pricing
- DB2 compression advantage
  - Reduces cost of incremental storage
- Technology refresh
  - Don’t pay for existing MIPs
  - No charge to upgrade specialty processors
- Sub capacity pricing vs. co-location
- Storage virtualization is included

On-Line Banking Benchmark Demonstrates Performance Advantages Of Co-Location

Separate Machines

150 tps

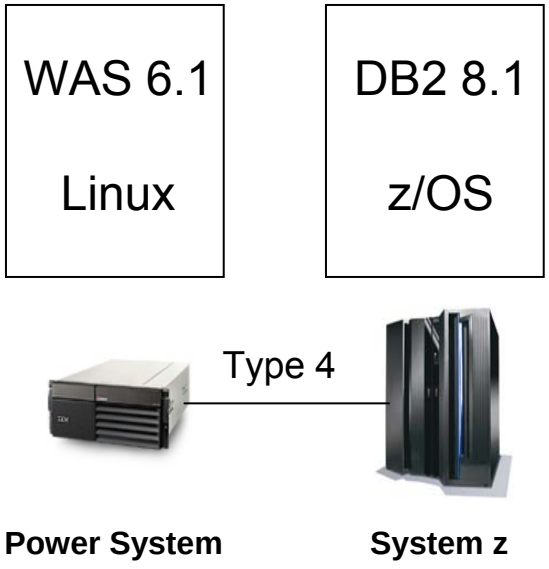

Separate LPARs

160 tps

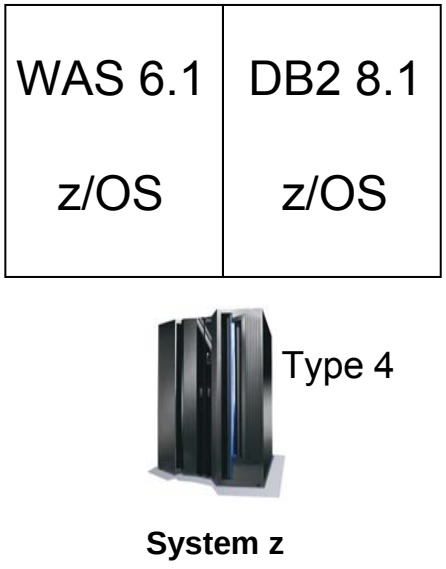

$5.6M (3 years)

Same LPAR

243 tps

$6.5M (+15%)
52% more throughput

Facts To Consider When Pricing Cost On A Distributed System

- Make sure you have estimated core proliferation
- Make sure you have estimated storage proliferation
- Separate production, development, quality assurance servers, fail-over
- Disaster recovery servers
- Infrastructure servers – systems management, networking, security/directory, workload distribution, firewalls, data staging…
- Distributed hardware needs to be repurchased when refreshed
- Migration cost, and loss of agility during the process
  - Dual environments during migration
- Provision for peaks and growth
- Language expansion (CICS/COBOL path lengths are highly optimized)
- Oracle RAC scaling inefficiencies compared to DB2
- Ensure batch can complete in the batch window. If not, then what?
- . . .

Deploy WAS Application On Mainframe z/OS vs. HP Servers

Existing Mainframe

Existing z10:
2 GP 1,720 MIPS
DB2 and utilities
With 20TB storage

Existing Disaster Recovery Site

Existing:
1 GP processor for hot disaster switch-over
1 "dark" DR processor
With 20TB storage

Add 1 LPAR for New Web Application w 1.28TB storage

Prod

1,624 MIPS additional workload

Incremental:

1 zAAP	920 MIPS WAS (85%)
1 GP	541 MIPS DB2
	163 MIPS WAS (15%)

2 GB memory

And Add Disaster Recovery w 1.28TB storage

Prod

Capacity Backup:
1 GP
1 zAAP

3 year cost of acquisition $3.05M

Or Add HP Integrity Superdome 9140 Server w 1.67TB storage

Prod

201,977 Performance Units*

And Add 1 server for Disaster Recovery, Development & QA w 1.67TB storage

Prod

201,977 Performance Units*

3 year cost of acquisition $8.53M

**Production Performance Units required = 1,624 x 122 = 198,128*

European Bank Study Shows WAS On Sun Costs 2.3X More Than zLinux

- Currently 9 distributed Sun servers running WebSphere workload
- Compare running the same workload on IBM System z10 with multiple IFLs
- Scope
  1. Cost – HW, SW, Power and Cooling, and Floor Space
  2. Discipline – Development, Test, Production, DR (on Dev/Test machines)
  3. Five Year TCO with HW acquisition in 1st and 4th year
  4. Migration labor costs are included for Linux on System z
  5. A projected annual capacity growth of 0%

Distributed TCO is €1,054,019 (2.3X) more expensive than z10 over 5 years

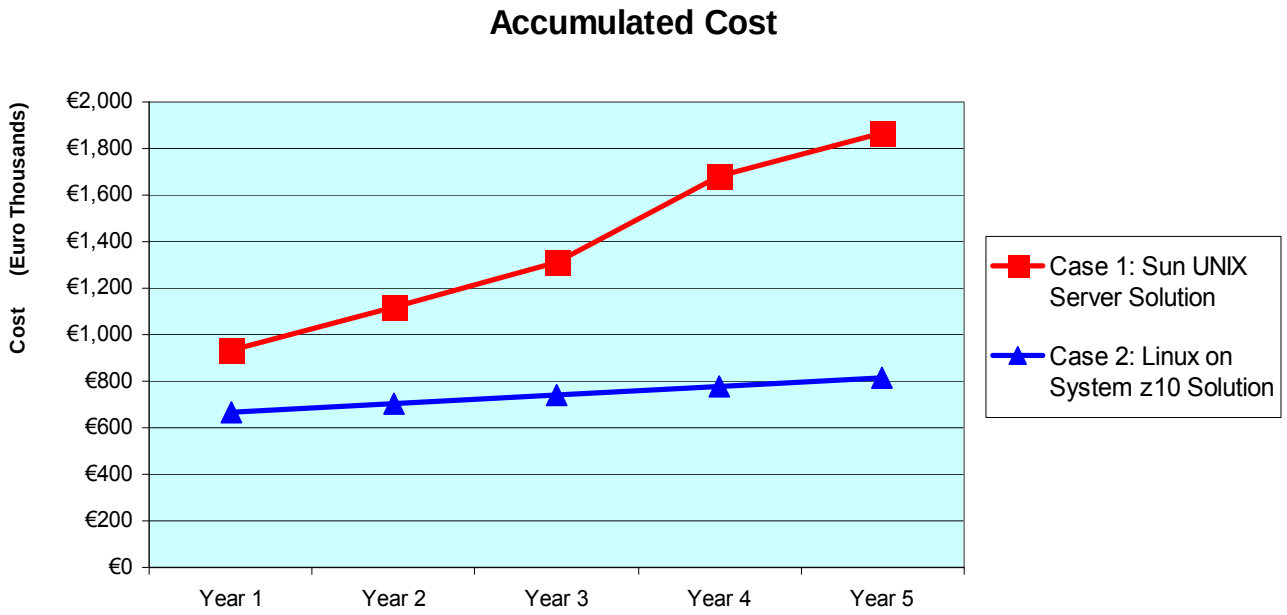

US Bank Study Shows WebSphere Process Server On Sun Costs 5.8X More Than System z

- Currently 3 distributed Sun servers running WebSphere workload
- Compare running same workload on IBM System z10 using zLinux or z/OS
- Scope
  1. Cost – HW, SW, Power, and Floor Space, but NOT labor
  2. Discipline – Production, QA, Development/Test, and DR
  3. Five Year TCO including HW acquisition in 1st and 4th year
  4. 3033 MIPS of workload on z/OS
  5. 3791 MIPS of workload on Linux for System z

Distributed TCO is $21,214,907 (3.8X) more expensive than z/OS over 5 years

Distributed TCO is $23,802,441 (5.8X) more expensive than Linux for System z over 5 years

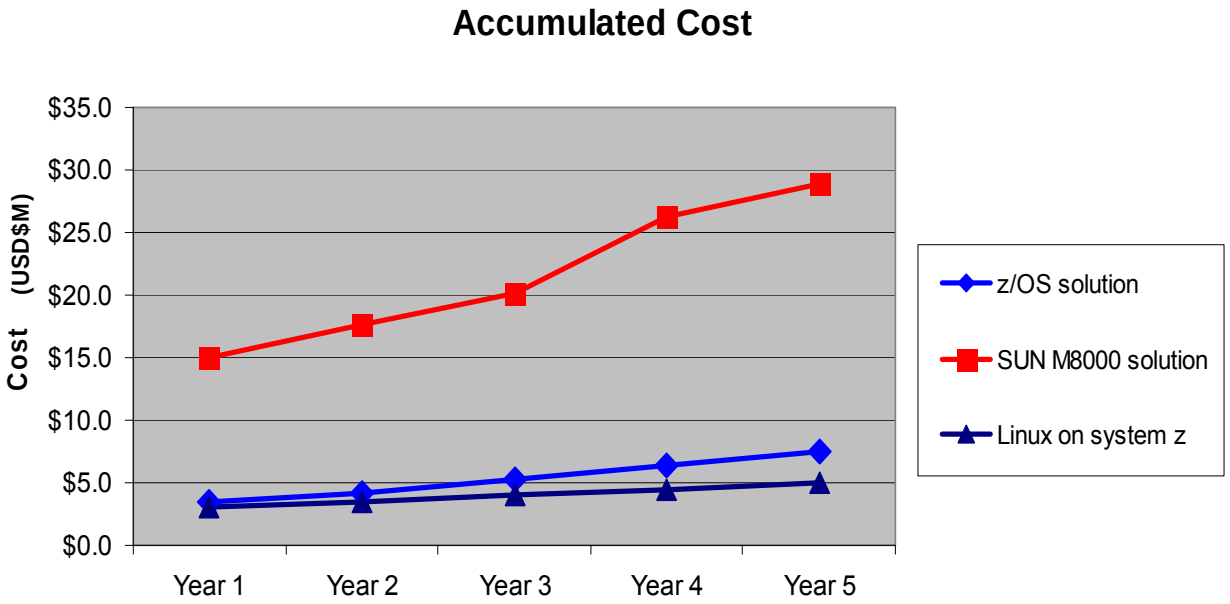

L.A. Bank Study Shows WebSphere Message Broker On Distributed Costs 2.1X More Than z

- Compare running same workload on IBM System z10 using zLinux or z/OS
- Scope
  1. Cost – HW, SW, Power, and Floor Space
  2. Discipline – Production, QA, Development/Test, and DR
  3. Five Year TCO including HW acquisition in 1st and 4th year
  4. +120 MIPS & 2 zAAPs of workload on z/OS
  5. 1 IFL for WMB production workload, 4 IFLs for dev/test etc. on Linux for system z
  6. 1 server for WMB production workload, 2 servers for dev/test etc. on distributed

Distributed TCO is $2,527,463 (1.9X) more expensive than z/OS over 5 years

Distributed TCO is $2,757,439 (2.1X) more expensive than Linux for system z over 5 years

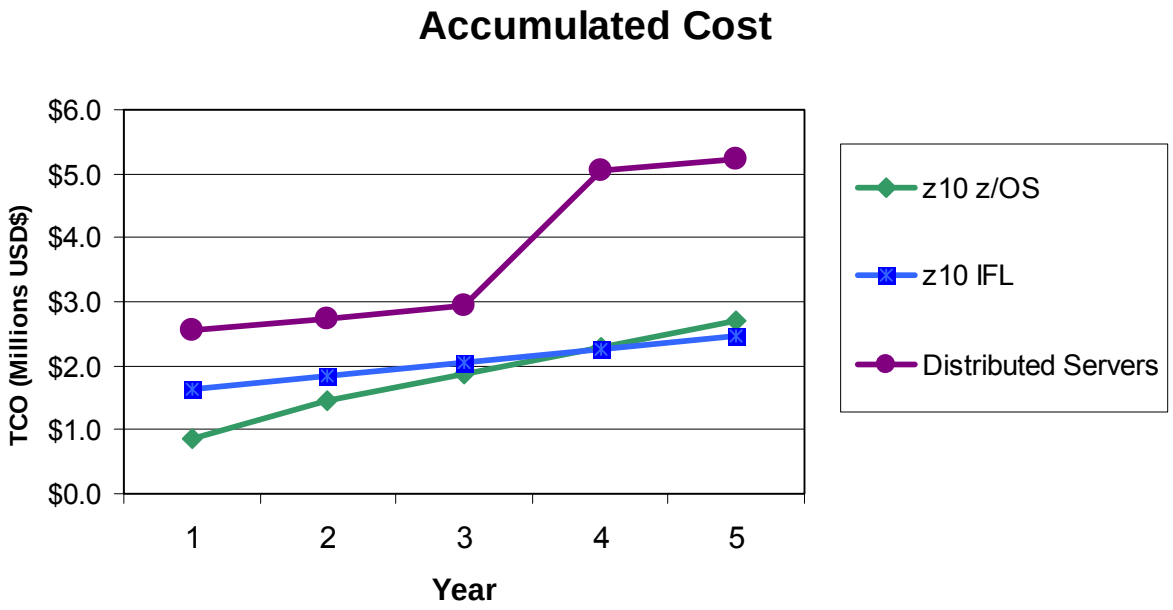

System z TCO Checklist – Incremental

New Workload

- ☑ Have you considered only the incremental cost if using an existing mainframe?
- ☑ Have you used LPARs and sub-capacity pricing to limit incremental cost?
- ☑ Have you used zIIPs and zAAPs for new workload?
- ☑ Are you co-locating your database and transaction monitor?
- ☑ Have you upgraded to the latest hardware to get improved price/performance of specialty engines?
- ☑ Have you extended your existing applications to get decreased costs/transaction?
- ☑ Do you have an ELA or OIO contract with IBM?
- ☑ Is your IBM seller aware of the latest pricing plays?
- ☑ Are you aware of the various Capacity on Demand capabilities, and are you using them?

Consolidation

- ☑ Do you understand the savings in software licensing?
- ☑ Have you considered System z's ability to over-commit memory by 3x?
- ☑ Have you examined the savings in
  - network complexity
  - storage required
  - power and cooling
  - labor productivity ?
- ☑ Have you considered how to avoid server hardware refresh?
- ☑ Are you using sub-capacity pricing where appropriate?
- ☑ Have you consolidated as much workload as possible on your System z?

- ☑ **Have you engaged with the zCPO TCO Studies team?**

Craig Bender +1-518-929-2290

Some Large Core Processing Workloads Can Only Run Efficiently On The Mainframe

Banks
Financial Services
Reservations
Transaction Accounts
Batch Workloads…

No effective alternative on distributed

An Existing Mainframe Can Be Incrementally Extended To Run New Workloads At A Lower Cost Than Distributed

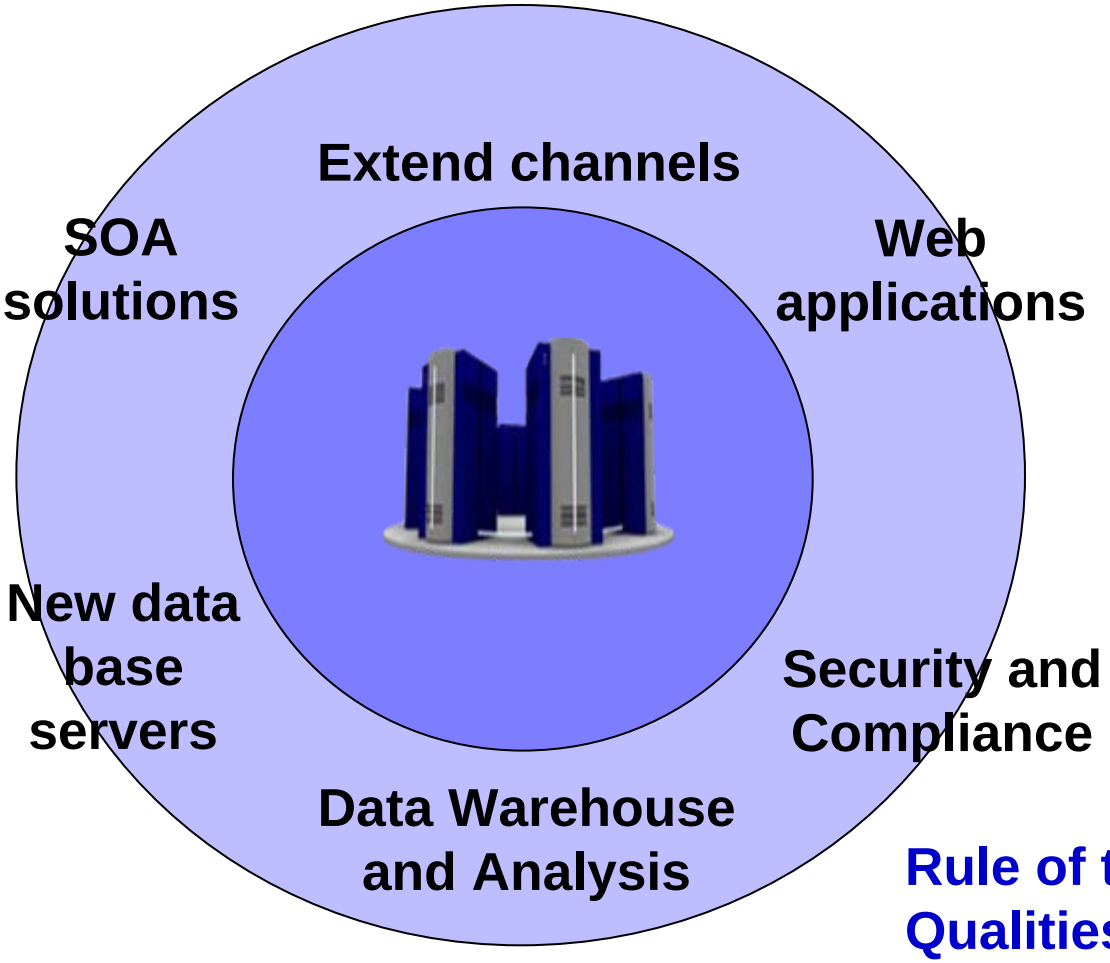

Rule of three!
Qualities of service
Co-location
Lower cost

Distributed Linux Workloads Can Be Consolidated To Cut Costs

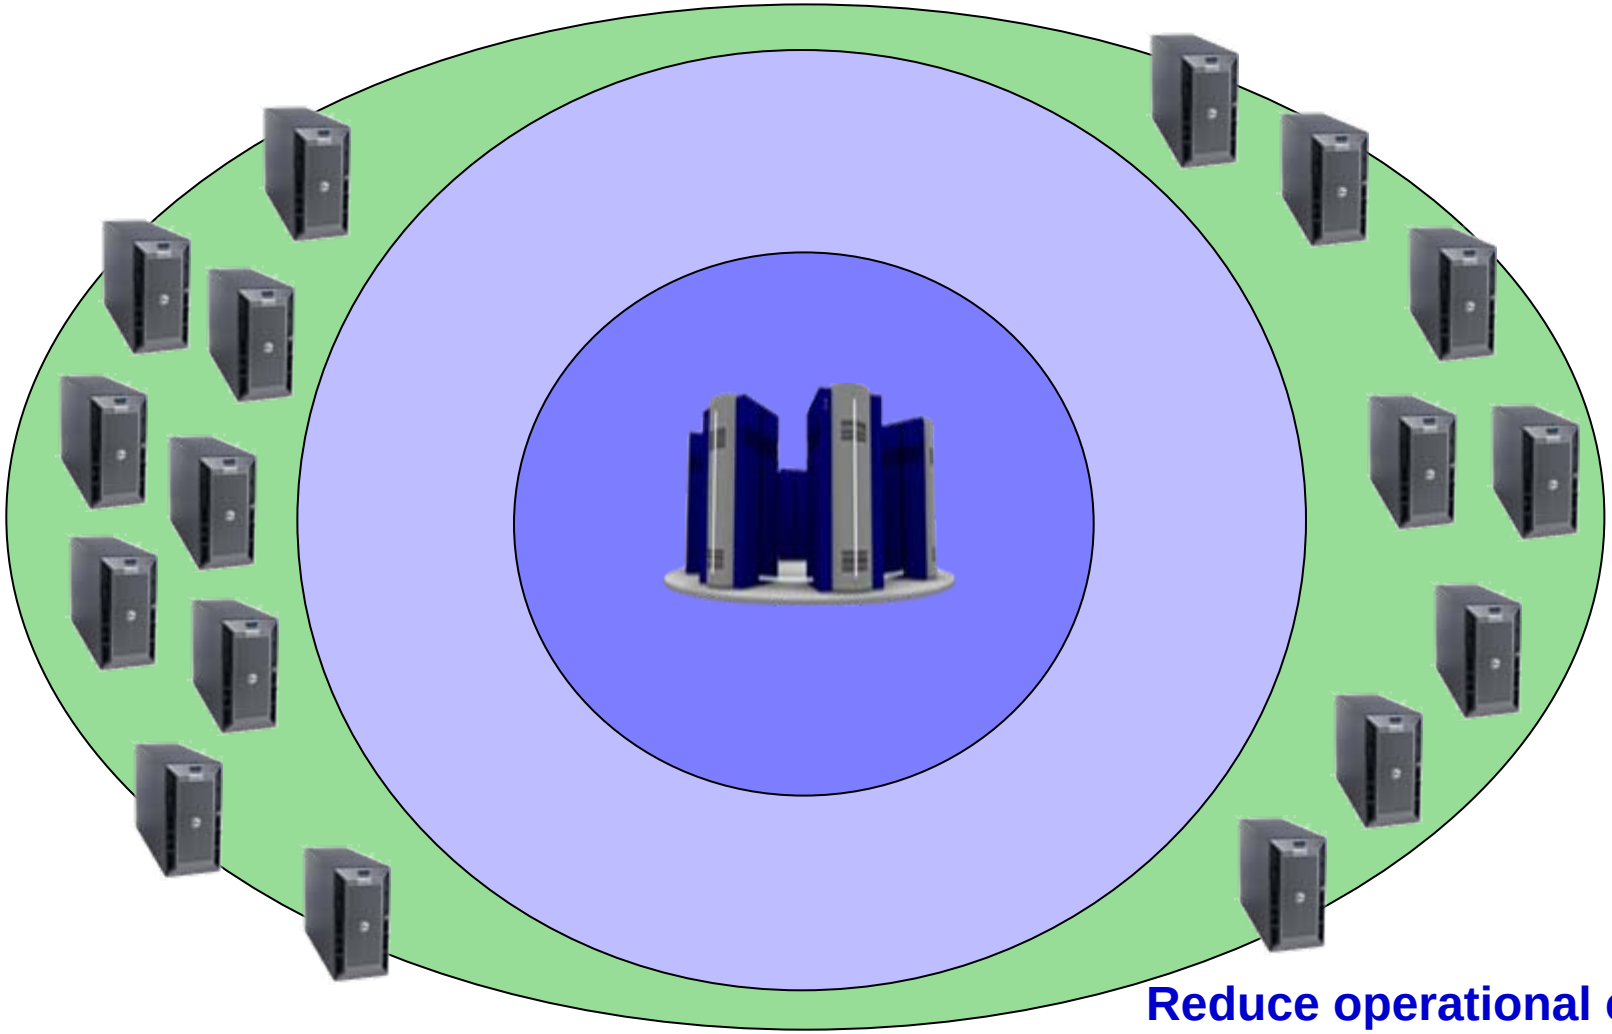

Reduce operational costs
Faster provisioning
Environmentals

A Fully Leveraged System z = Lowest Cost Per Unit Of Work