



Extending Your Mainframe For More Business Value

Server Consolidation to Linux on IFLs

Typical Scenarios

- Most data centers are not green field projects
- Cost concerns drive typical scenarios:
- Large transactional workloads and database
 - ▶ Scale may compel platform choice
- Adding new workload to an existing System z
 - ▶ The rule of three
- **Server consolidation to Linux on IFLs**
 - ▶ Consolidation Math
- Offloading projects
 - ▶ Proliferation of cores defeats distributed price advantages

Distributed Server Sprawl




**Service Oriented Finance
CIO**

IBM


Distributed Server Sprawl Uses...

- Lots of hardware
 - ▶ Lots of floorspace
 - ▶ Lots of power
 - ▶ Lots of networking
- Lots of software licenses
- Lots of people to manage the systems
- **Consequences**
 - ▶ Low Utilization of Hardware Resources
 - ▶ Complexity
 - ▶ Increased time to respond to business requirements
 - ▶ Difficulty integrating information from various systems



Server Utilization At A Large Financial Institution

Average Server Utilization by Class
Jan-08



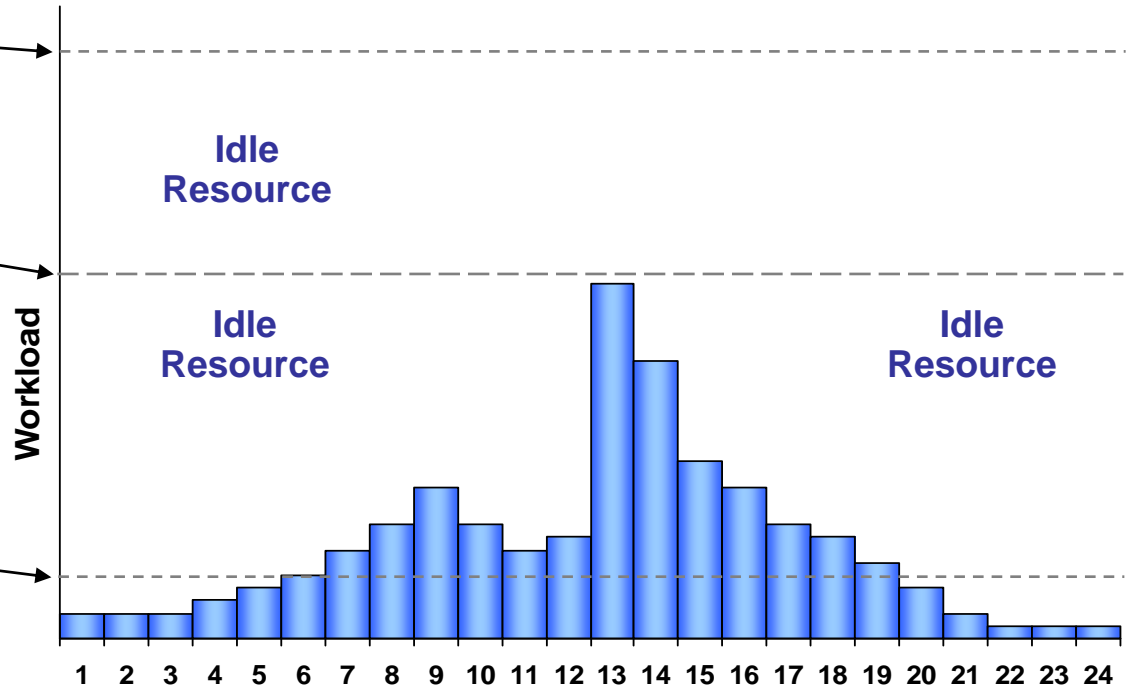
Utilization Of Distributed Servers

Provision for expected growth

Provision capacity for peak workload



Average utilization

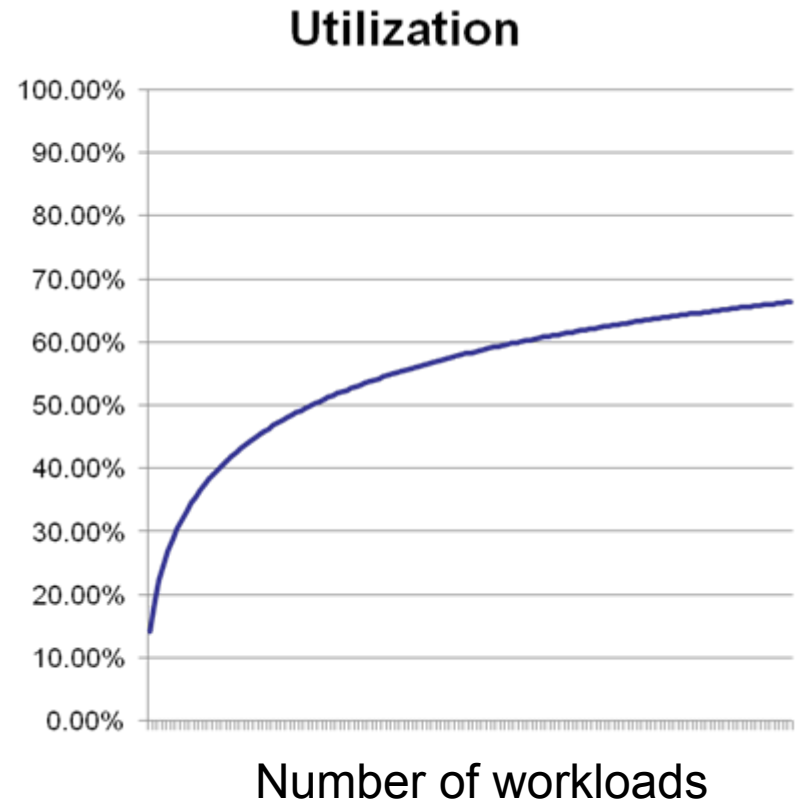


Server dedicated to one application

- ▶ Typical utilization of Windows Servers <5%
- ▶ Typical utilization of UNIX Servers 15 – 20%
- ▶ Typical utilization of System z Servers 70 – 100%

Statistics Can Work In Our Favor

- When the number of workloads is small
 - ▶ To guarantee that work will complete within “specified” time requires more excess capacity
- When the number of workloads is large
 - ▶ Combination of arriving workloads is more statistically predictable
 - ▶ Higher predictability means lower excess capacity required to meet the specified response times
- When many applications are brought from single application servers to a centralized server, wasted utilization can be squeezed out of the datacenter



An Experiment Shows How Combining Workloads On A Shared Server Statistically Improves Utilization

■ Group 1

- ▶ Take 1 die and roll it 10 times. Count the number of times you get a 1, 2, 3, 6
- ▶ Plot your results on a histogram

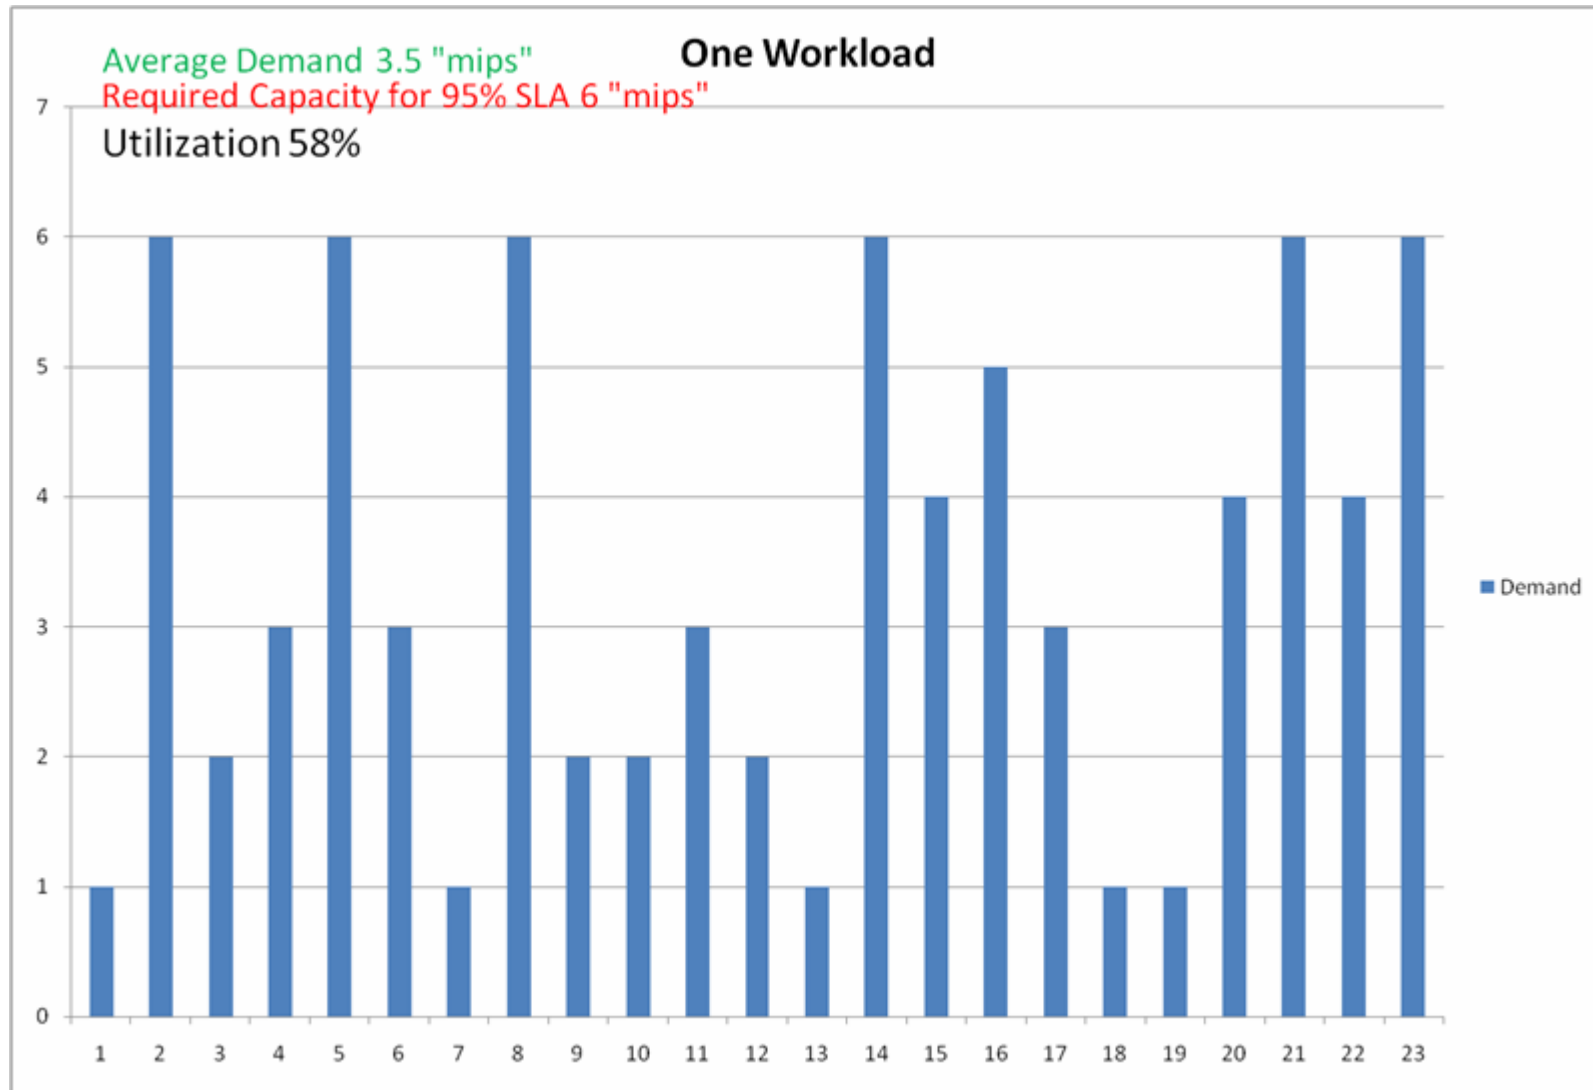


■ Group 2

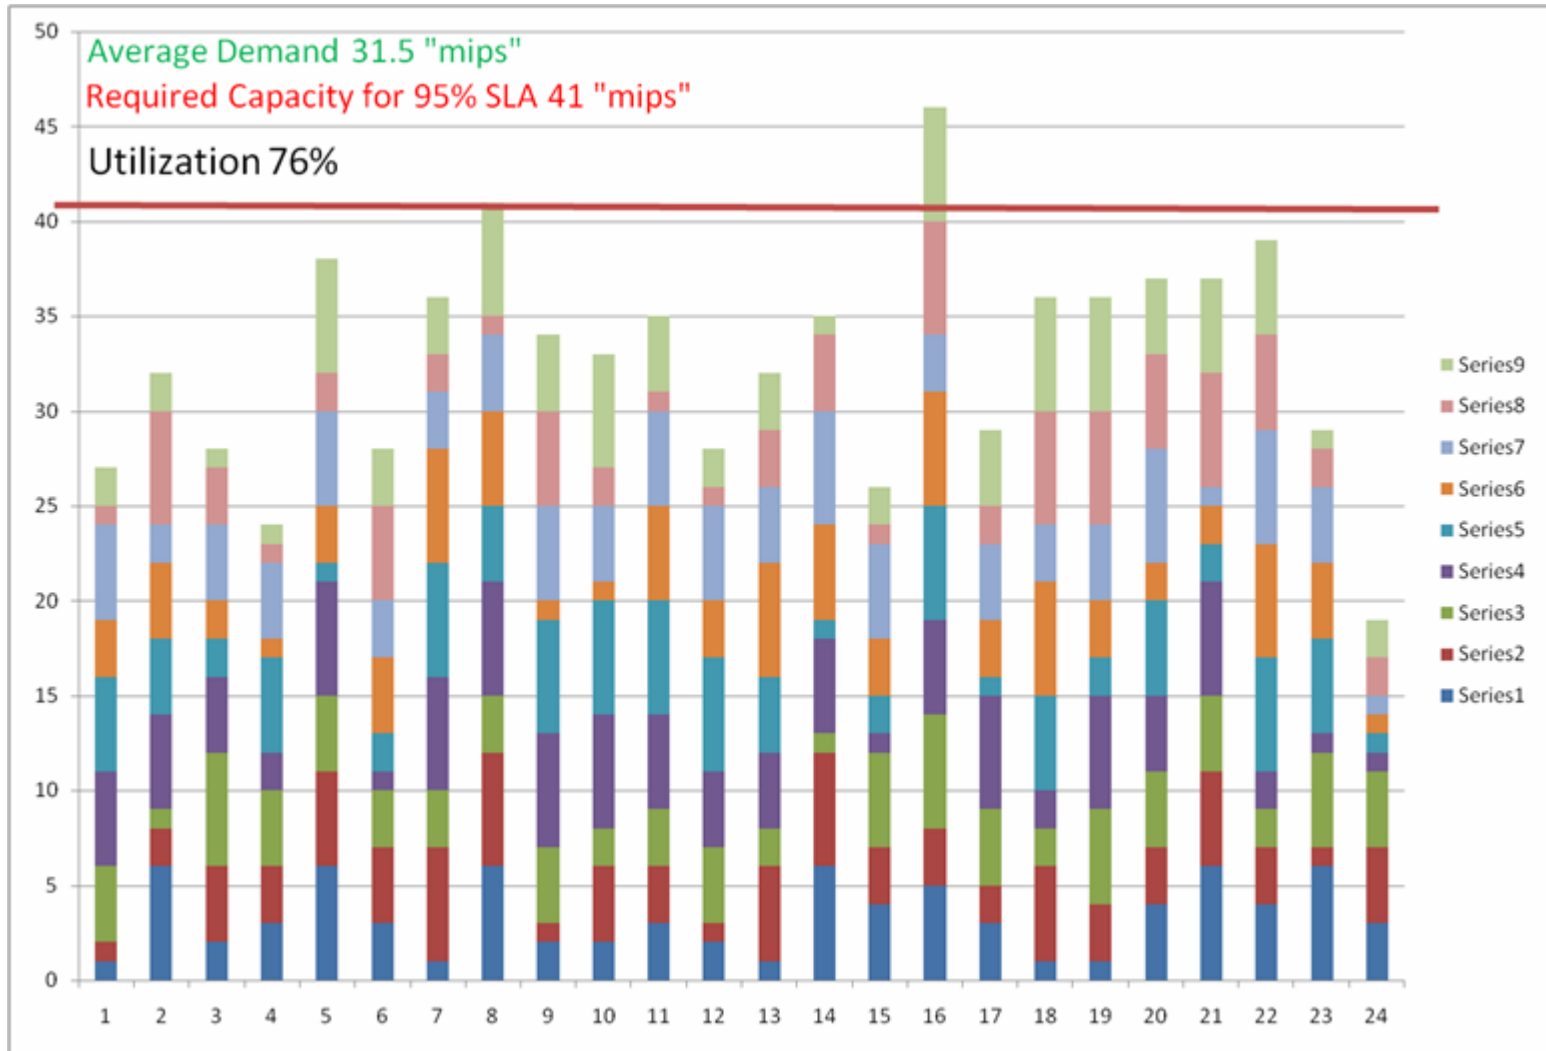
- ▶ Take 9 dice and roll them 10 times. On each roll get the total value on the 10 dice. Count the number of times you get a 9, 10, 11, 12,, 30, 31, ... 54.
- ▶ Plot your results on a histogram.

- What do we see about the “predictability” of the result of a “roll”? (Let’s roll the dice with a computer.)


After Rolling 1 Die (1 Workload) The Distribution Would Look Like This



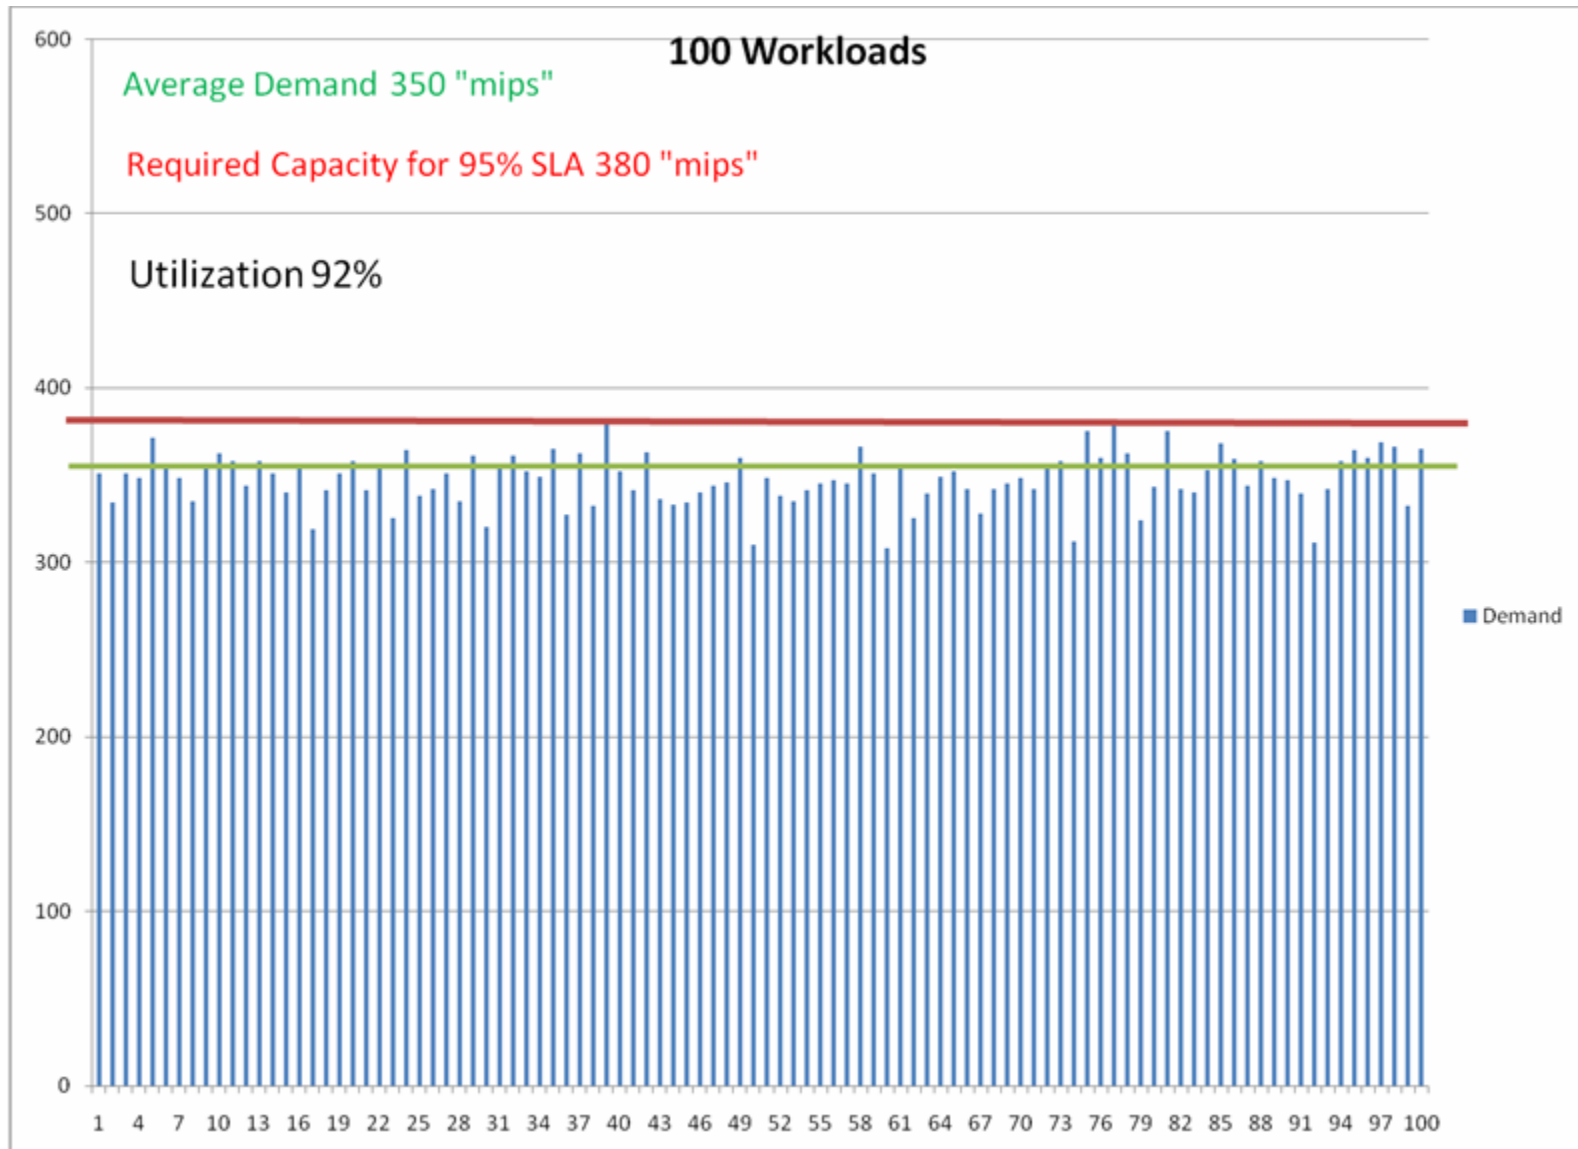
After Rolling 9 Dice (9 Workloads) The Distribution Would Look Like This



After Rolling 30 Dice (30 Workloads) The Distribution Would Look Like This



After Rolling 100 Dice (100 Workloads) The Distribution Would Look Like This



Economics Of Consolidation

- Consolidating workload means running multiple workloads on the mainframe at the same time
- Consolidation achieves greater utilization of assets which minimizes cost per unit of work
- Same principal was applied by Henry Ford at the dawn of the industry era
 - ▶ It still applies today
- Workload consolidation on a mainframe squeezes out cost to achieve maximum efficiency
 - ▶ And return on investment



Copyright © 2006, Toyota Motor Manufacturing Kentucky, Inc.

Consolidation Math

What is the theoretical maximum number of servers that can be consolidated?



N Servers

P_A – Processor Power

U_A – Utilization

C_A – Cores Per Server

One Server

P_B – Processor Power

U_B – Utilization

C_B – Cores Per Server

Ratios

$$P_R = P_B / P_A$$

$$U_R = U_B / U_A$$

$$C_R = C_B / C_A$$

$$N \leq \left(\begin{array}{c} \text{Processor} \\ \text{Performance} \\ \text{Ratio} \end{array} \right) \left(\begin{array}{c} \text{Processor} \\ \text{Utilization} \\ \text{Ratio} \end{array} \right) \left(\begin{array}{c} \text{Cores per} \\ \text{Frame} \\ \text{Ratio} \end{array} \right)$$
$$N \leq (P_R) (U_R) (C_R)$$


Implementation variations from average and practical considerations will constrain this theoretical number

Identify Consolidation Opportunities

The more servers you can consolidate, the more money you will save (Maximize N)


$$N \leq (P_R) (U_R) (C_R)$$

Servers that are candidates to be consolidated



Older servers with slower processor

Servers that are best consolidation platforms



New servers with faster processor

Servers with low utilization

Servers that can achieve sustained high utilization

Servers with a low number of cores

Servers with a high number of cores

Performance Ratio

Utilization Ratio

Core Ratio


Typical Ratios

1.0 - 3.0

10 - 20

1- 64

Observed Consolidation Ratios



The Mainframe Can Consolidate Thousands of Workloads

- Mainframe hardware provides:
 - ▶ Hypervisor assigns processor resources to logical partitions
 - ▶ Intelligent Resource Director supervises this assignment
 - ▶ Virtualized I/O Subsystem

- z/OS provides:
 - ▶ Workload Manager assigns resources within a z/OS image according to service level agreements
 - ▶ Also performs this function across a cluster of z/OS images

- z/VM provides:
 - ▶ Virtual Machine Resource Manager
 - ▶ Complete mainframe virtualization (including memory)

- All of these facilities provide
 - ▶ Business workload oriented goal or velocity definitions
 - ▶ Autonomic and continuous management to those definitions

Other Workload Management Solutions Lack Business Goals and Dynamic Flexibility

- For example, HP-UX workload manager
 - ▶ Marketing description looks like it is referring to z/OS WLM!
 - “...a goal based policy engine...”
 - ▶ In fact, HP’s WLM uses static relationships based on the IT environment, not related to business goals
 - ▶ Can adjust share/number of CPUs available to a workload but only via static relationships:
 - “3 cpu-shares per connected user with a minimum of x, maximum of y”
 - ▶ Can only hard code memory and disk bandwidth shares
 - Dynamic changes not possible so only useful for hard-caps

- No sign of z/OS WLM business goals like:
 - ▶ Application X will achieve <2s response time for 95% of users during US Eastern business hours, <5s at other times

Example Workloads That Can Be Consolidated On A Mainframe

What	Where	Specialty Processor	How
Growth of Existing Mainframe Workload	z/OS	--	Capacity on demand
New CICS or IMS Applications	z/OS	--	Develop
Data Warehouse	z/OS	zIIP	Deploy
SAP Database Server	z/OS	zIIP	Deploy
WebSphere Application Server	z/OS	zAAP	Deploy
WebSphere Portal Server	z/OS	zAAP	Deploy
WebSphere Process Server	z/OS	zAAP	Deploy
.NET Applications	z/OS	zAAP	Mainsoft
Domino	z/OS	--	Deploy

More Example Workloads That Can Be Consolidated On A Mainframe

What	Where	Specialty Processor	How
Linux Applications	Linux on z/VM	IFL	Recompile
Linux Middleware - IBM Brands (DB2, WebSphere, Lotus, Rational, Tivoli) - Oracle Database - etc.	Linux on z/VM	IFL	Rehost
Linux Packaged Applications - SAP - Oracle - etc.	Linux on z/VM	IFL	Rehost
.NET Applications	Linux on z/VM	IFL	Mono, Mainsoft
Open Solaris Applications	Open Solaris on z/VM	IFL	Sine Nomine

Linux On z/VM

We've seen some examples of incremental growth on z/OS

- ▶ Extend new access channels with WebSphere
- ▶ New data workloads with DB2
- ▶ Business insight with DB2 and Information Server
- ▶ Communications backbone with IBM Enterprise Service Bus

Now let's look at some examples of roll-up consolidation to Linux on z/VM



IBM



Nationwide® Saves \$16+ Million With Linux On Your Side™ On System z


■ **Problems:**

- ▶ High TCO including data center power and floor space scarcity
 - New facility would cost \$10M+
- ▶ Long server provisioning process

▶ **Solution:**

- ▶ **350** servers virtualized with **15** z990 IFLs – **23 to 1 consolidation**
 - 12 mission critical applications with 100,000+ users/day
- ▶ 50% reduction in Web hosting monthly costs
- ▶ 80% reduction in floor space and power conservation
- ▶ 50% reduction in hardware and OS support efforts
 - Significant savings on middleware costs
- ▶ Significantly faster provisioning speed (months → days)
- ▶ Mainframe high availability and disaster recovery

Vastly improved TCO, Speed & Simplification



Update (August 2008):

- \$16M savings realized a year earlier than planned
 - ▶ In 2 years not 3

- Up to 18 mission critical applications
 - ▶ Added more WebSphere, Portal, and DB2

- Upgraded from z990 to z9 IFL's

- **517** virtual servers with 1,630 JVM's running on **35** z9 IFL's
 - ▶ So, workloads that would have required **1,630** physical servers are running on **35** z9 IFL's – **47 to 1 consolidation**


Improved TCO, Speed and Simplification

Telemar Roll-up Consolidation Project

Largest provider of fixed-line telecommunications services in South America.



Consolidated 16 geographically dispersed UNIX servers on a centralized System z9 EC
16 to 1 consolidation



Source Robert Frances Group

Benefits:

- Open-standards-based solution
- Maximized manageability, scalability, security and availability of its key business systems.
- Reduced need for server capacity by one-third
- Lowered operating and administration for maintaining email server applications.



Case Study: Québec Government Runs Oracle At IFL Prices

- Running **292** server instances on a z9-EC with **5** IFL's
 - ▶ 200 Oracle, 80 WebSphere, 12 WebSphere messaging
 - ▶ Reduced cost of hardware and software by 30%
 - Saved \$800,000 in licensing cost in the first year
 - ▶ Used RACF for consistent security
 - ▶ Each administrator can manage 100 consolidated Linux images (up from 30)
 - ▶ Easy migration
 - Create new Linux server in 30 min (vs. 1 week – 3 months)
 - Clone Oracle DB instance in 30-45 min (vs. 10 – 14 hours)
 - ▶ Inherited benefits of z platform – workload management, availability, disaster recovery, I/O bandwidth

Case Study: Consolidate On Mainframe vs. Keeping Dedicated Servers

Existing Mainframe



Existing processors:
4 general purpose

Add 1 LPAR for Oracle Server Consolidation



Add three processors:
3 IFLs

Or maintain existing 292 server farm for Oracle data servers



*3 year TCO
\$9.06M*

Annual operating cost \$0.67M

Breakeven in first year

*3 year TCO
\$30.13M*

Annual operating cost \$10.04 M

Case Study: Consolidate On Mainframe vs. Keeping Dedicated Servers

Mainframe Incremental Hardware

OTC		ANNUAL	
3 IFL Processors	\$375,000	Processor ² Maintenance	\$52,524
Conn. + Disk Acquisition	\$639,033	Power/Space ¹	\$47,073
RAM (190GB)	\$1,140,000	Conn. + Disk Maintenance ¹	\$87,480
Migration	\$4,920,492	System Admin ¹	\$386,518
		On-Premise Network Maintenance ¹	\$8,935
TOTAL	\$7,074,525	TOTAL	\$582,530 (year 2, 3)

Mainframe Software

OTC		ANNUAL	
z/VM	\$67,500	z/VM ²	\$16,890
		Oracle S&S ²	\$26,400
		Linux S&S ¹	\$45,000
TOTAL	\$67,500	TOTAL	\$88,290 (year 2, 3)

Dedicated Hardware

OTC		ANNUAL	
Sunk Cost	\$0	Disk Maintenance ¹	\$59,276
		Server maintenance ¹	\$226,884
		Off-Premise Network ¹	\$299,008
		Power/Floorspace ¹	\$501,656
		System Admin ¹	\$5,944,828
		On-Premise Network Maintenance ¹	\$62,196
TOTAL	\$0	TOTAL	\$7,093,848




Dedicated Software

OTC		ANNUAL	
Sunk Costs	\$0	Oracle S&S ¹	\$2,569,600
		Linux S&S ¹	\$379,308
TOTAL	\$0	TOTAL	\$2,948,908

1 – Needs three years maintenance, 2 – Needs two years maintenance

Network Simplification

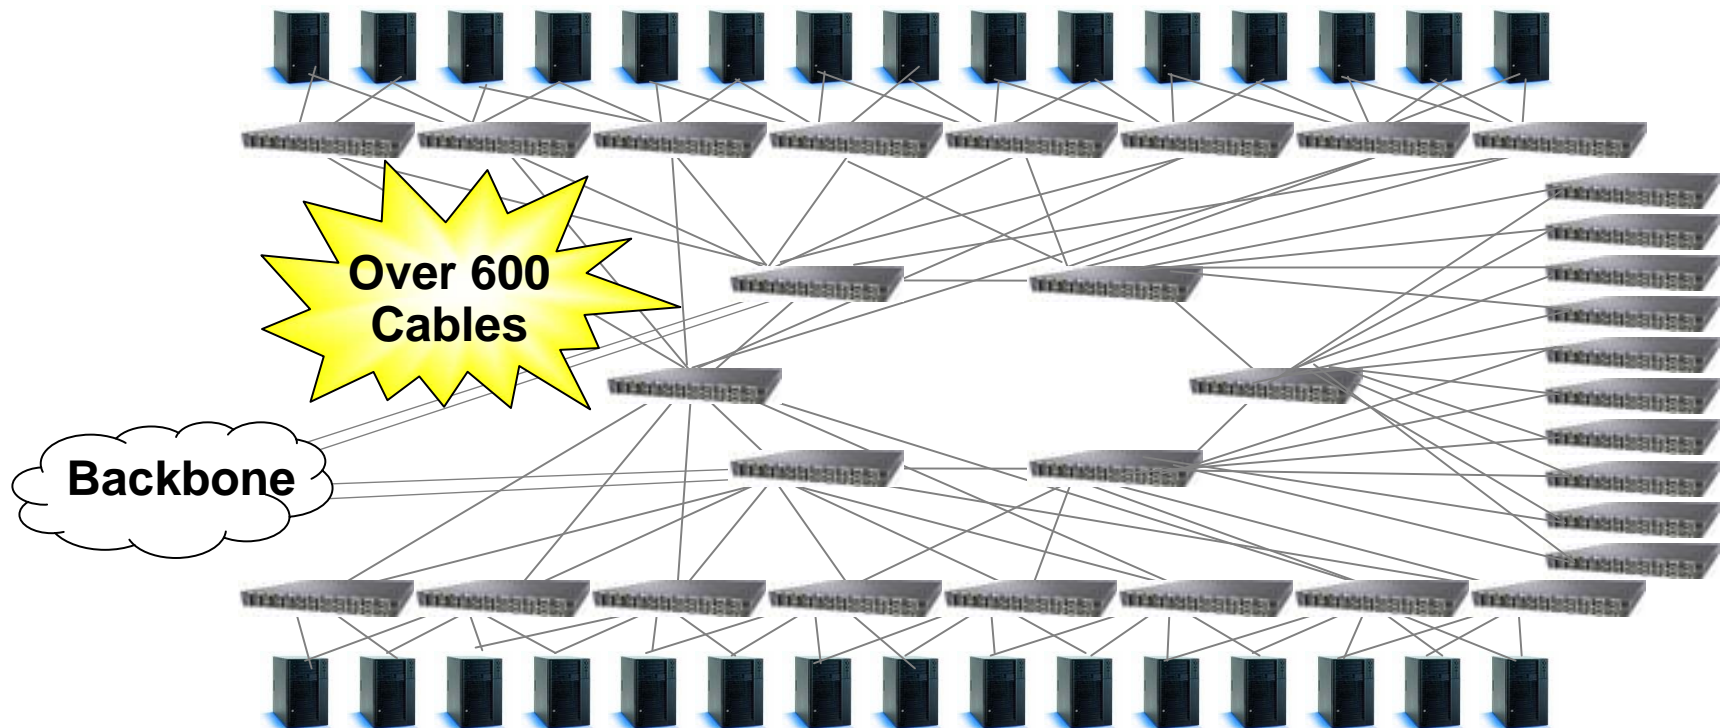
- Consolidation replaces cables and routers with internal connections
- Better performance and security



Case Study: Network Before Consolidation (292 Servers To 1 System z)

Catalyst 3560E-24TD – 24 ports 1gbps	25
Catalyst 3560E-12D – 12 ports 10gbps	6
50 Ft UTP Cable	584
10GB Eth Fiber Cable	60


The diagram shows only **30** of the **292** Servers



Case Study: Network After Consolidation (292 Servers To 1 System z)

Network Simplification!

Better Performance and Security




Other Benefits Of Virtualization

- Fast provisioning of pre-installed and configured images
 - ▶ Minutes instead of days or weeks
 - ▶ No additional space, electric connections or network cables
- Compatible with the data center practice of standardizing on strategic software stacks
 - ▶ Pre-tested stacks
 - ▶ Consistent levels and patch
 - ▶ A management approach to achieve better stability
 - ▶ Jukebox selection of standard enterprise images

Demo: Provisioning using z/VM Center

Need a new machine? Let's see how fast we can get one...



Benefits Of Consolidation On The Mainframe

- Less hardware
- Fewer software licenses
- Less costly to manage
- Consumes less power and floor space
- Responsiveness to the business via faster provisioning
- Inherit the benefits of the mainframe platform
 - ▶ High reliability
 - ▶ I/O bandwidth
 - ▶ Consistent security
 - ▶ Systematic disaster recovery
- Lower annual costs!



Mainframe Labor Costs Per MIP Declining

- IBM Survey five years ago, average MIPS per person
 - ▶ **50** for z/OS
- Typical MIPS per person today
 - ▶ **150 to 700** for z/OS (1,300 to 2,000 for zLinux)
- A major bank went from 128 MIPS/person to 597 MIPS/person in 8 years with no extra people
- Gartner showed the MIPS/person doubling in 3 years at another site
- An outsourcer stated they doubled MIPS with only 20% increase in headcount

IBM Internal Project To Consolidate Over 3,000 Servers

- IBM expects substantial operational annual savings by consolidating 3,917 distributed servers to about 30 mainframes
 - ▶ 86% savings in system administration cost
 - ▶ 85% savings in floor space
 - ▶ 81% savings in power
 - ▶ 57% savings in network management
- \$81M savings per year including
 - ▶ Operational savings above
 - ▶ Hardware and software maintenance

TCO Case Studies Demonstrate Consistent Savings In Annual Operations Costs

Scenarios	Cost of Distributed vs. z	Distributed Cost Ratio	Cores vs. z Processors	Core Ratio
Linux Consolidation				
- Nationwide			1630 → 35	47 → 1
- Quebec Govt			292 → 5	58 → 1
- Hannaford			150 → 1	150 → 1
- Brokerage Firm (Power)	\$12.7M vs \$7.8M	1.6x	112 → 1	112 → 1
- Brokerage Firm (Floor)	\$25.5M vs \$10.7M	2.4x	180 → 2	90 → 1
- Major Bank	\$46.9M vs \$19.9M	2.4x	520 → 14	37 → 1

2.1x

82 : 1

Lower annual costs pay back initial migration investment quickly
 Core consolidation ratio varies with situation


What About Using VMWare On Intel?

- VMWare lacks the consolidation efficiency of z/VM
- Less efficient use of memory and storage
- Less efficient use of processors
- Not fully supported with enterprise software

	z/VM	VMWare
Maximum memory per virtual Linux server	More than 256GB	16GB
Maximum CPU's per virtual Linux server	Up to 64	Up to 4
Maximum "Active virtual memory" supported	Up to 8TB	16,384MB
Maximum virtual CPU's per core	Thousands	Up to 8
Maximum real memory	Up to 256GB	Up to 64GB
Maximum virtual servers per machine	Thousands	128
Call Oracle for support if you have a database problem?	Yes	No - uncertified


Result: Consolidation on z/VM Saves the Most Money

Comparison of consolidation options



Cost of Different Linux Consolidation Solutions (Zoom on 0-300 Servers)

Comparison of Consolidation TCO



■ 292 servers:

70% TCO reduction on mainframe

Only 41% on VMWare

Do YOU Need To Consolidate?

- I/T department whose budget is consumed by operating cost?
- Contemplating new data centers due to power or floor space constraints?
- Need a systematic site failover plan for **all** applications and data?
- Quality of service issues?
- Lots of UNIX or Linux servers?
- Lots of small database servers scattered around (including Oracle)?



Service Oriented Finance Did A Roll-up Consolidation Of Linux Servers

I saved a lot of money by consolidating our Linux servers onto System z!



**Service Oriented Finance
CIO**

