



Doing a (Dis)Honest Linux TCO Analysis or better Cost and Value

Share Orlando 2008 Session 9261

***Erich Amrehn
Romney White***



Agenda

- TCO Factors
- IT Cost Trends
- TCO Case Studies
- Mainframe Advantages
- Univar Example
- References

Attention / Achtung!


- Not all charts are in the handout due to customer confidentiality requirements
 - We felt it would be better to show you the information than not
-
- As of February 20, 2008, one € ≈ \$1.38385

Major TCO Factors

- Hardware
- Software
- Other (e.g., Environmental)
- People

TCO

Production



Hidden costs

Other TCO Factors

- Availability
 - ▶ High availability
 - ▶ Hours of operation
- Backup/Restore/Site Recovery
 - ▶ Backup
 - ▶ Disaster Scenario
 - ▶ Restore
 - ▶ Effort for Complete Site Recovery
 - ▶ SAN effort
- Infrastructure Cost
 - ▶ Space
 - ▶ Power
 - ▶ Network Infrastructure
 - ▶ Storage Infrastructure
- Additional development and implementation
 - ▶ Investment for one platform – reproduction for others
- Controlling and Accounting
 - ▶ Analyzing the systems
 - ▶ Cost
- Operations Effort
 - ▶ Monitoring, Operating
 - ▶ Problem Determination
 - ▶ Server Management Tools
 - ▶ Integrated Server Management – Enterprise Wide

Other TCO Factors ...

- Security
 - ▶ Authentication / Authorization
 - ▶ User Administration
 - ▶ Data Security
 - ▶ Server and OS Security
 - ▶ RACF vs. other solutions
- Deployment and Support
 - ▶ System Programming
 - Keeping consistent OS and SW Level
 - ▶ Middleware
 - SW Maintenance
 - SW Distribution (across firewall)
 - ▶ Application
 - Database Effort
 - Technology Upgrade
 - Non-disruptive System Release change
- Operating Concept
 - ▶ Development of an operating procedure
 - ▶ Feasibility of the developed procedure
 - ▶ Automation
- Resource Utilization and Performance
 - ▶ Mixed Workload / Batch
 - ▶ Resource Sharing
 - Shared nothing vs. shared everything
 - ▶ Parallel Sysplex vs. Other Concepts
 - ▶ Response Time
 - ▶ Performance Management
 - ▶ Peak handling / scalability

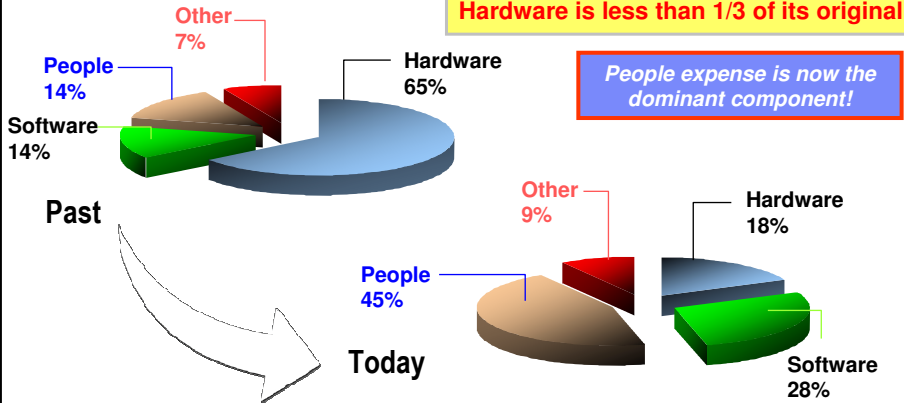
Other TCO Factors ...

- Integration
 - ▶ Integrated Functionality vs. Functionality to be implemented (possibly with third-party tools)
 - ▶ Balanced System
 - ▶ Integration of / into Standards
- Skills and Resources
 - ▶ Personnel Education
 - ▶ Availability of Resources
- Further Availability Aspects
 - ▶ Planned outages
 - ▶ Unplanned outages
 - ▶ Automated Take Over
 - ▶ Uninterrupted Take Over (especially for DB)
 - ▶ Workload Management across physical borders
 - ▶ Business continuity
 - ▶ Availability effects for other applications / projects
 - ▶ End User Service
 - ▶ End User Productivity
 - ▶ Virtualization

Cost dynamics of supporting corporate IT infrastructures has changed significantly


We typically see . . .

**People expense has tripled as a %
Software expense has doubled as a %
Hardware is less than 1/3 of its original %**




Source: IBM Scorpion studies – analysis of typical CIO budget for Mainframe plus UNIX/Windows servers


The IT Infrastructure as a Complexity Item



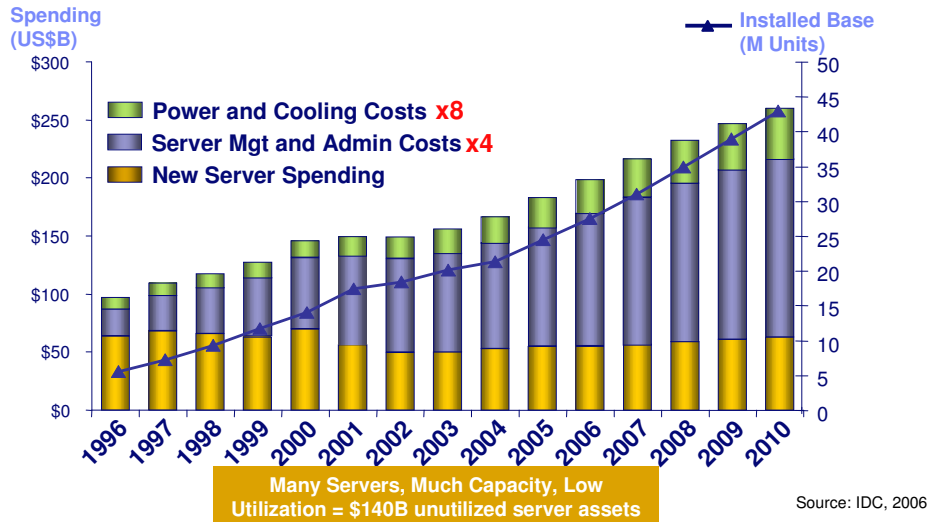
Cost of People vs. Spending on New Systems



- Size & Volumes
- Control
- Management costs

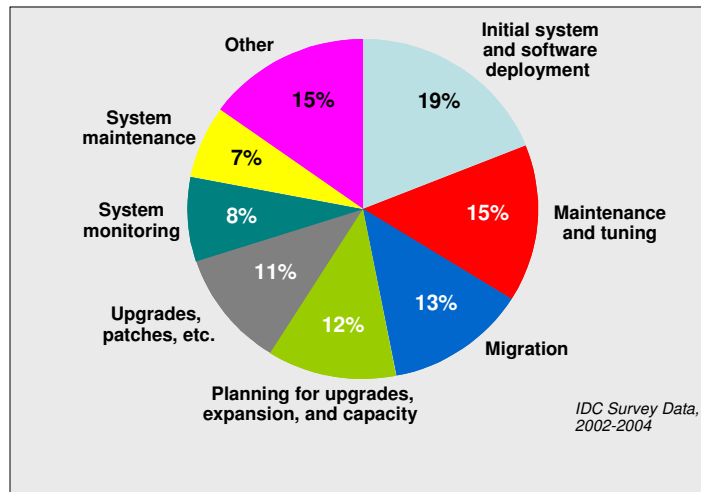


Worldwide IT Spending on Servers, Power and Cooling, and Management/Administration



Categories Of Costly System Management Tasks

Server Management and Administration Costs (2004 - 95B\$)




Many common management tasks must be simplified or eliminated to achieve significant improvements in IT productivity

Business Issue of Availability

Financial Impact of Downtime Per Hour
(by various Industries)

Brokerage Retail	\$6.5 Million
Credit Card Sales Authorization	\$2.6 Million
Airline Reservation Centers	\$90,000
Package Shipping Services	\$28,250
Manufacturing Industry	\$26,761
Banking Industry	\$17,093
Transportation Industry	\$9,435

Source: ©Eagle Rock Alliance, LTD. All Rights Reserved 2003



IDC 2005

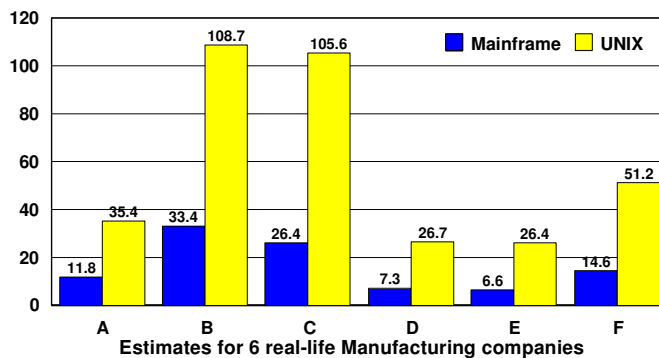
- **Loss of business**
- **Loss of customers – the competition is just a mouse click away**
- **Loss of credibility, brand image and stock value**



- On demand challenges
 - ▶ Downtime unaffordable
 - ▶ Heterogeneous by nature
 - ▶ Complex to manage
- Customer pressures
 - ▶ Application availability
 - ▶ Operations complexity and costs
 - ▶ Automation implementation and maintenance costs
 - ▶ Rapid change of I/T infrastructure
 - ▶ Adding new workloads

Downtime Costs


5 year cost of downtime - millions of dollars



Estimates for 6 real-life Manufacturing companies

Source: Business Value of Availability, Bottom Line Impact of SAP R/3 Platform Choices, ITG, November 2003

TCO / TCA / TVO / RCO / ROI ? MORE ??



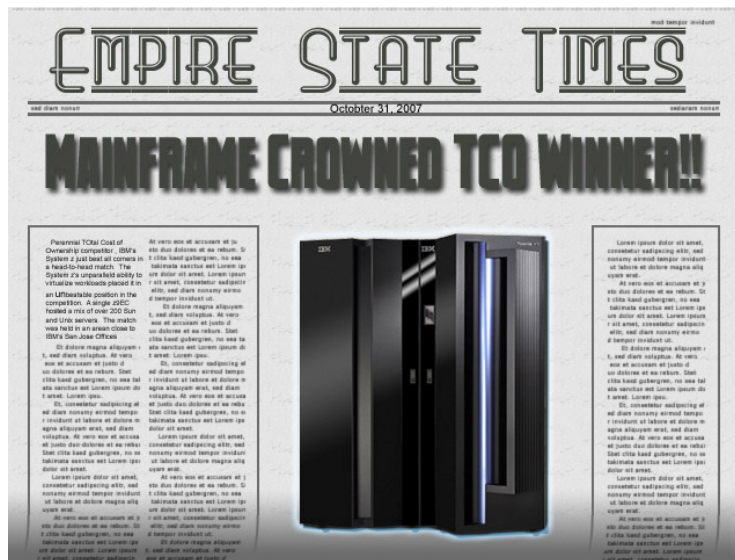
Is the complexity of your infrastructure costing you more than you can afford?

A power / cooling crisis is upon us...

“Sometimes we run out of power, sometimes we run out of cooling, usually we run out of both”



Anonymous

Much of the crisis is due to unrestricted server sprawl without regard to power/cooling and space



Because IT Complexity Drives Many Hidden Costs TCO

- Do you recognize this description?
 - ▶ Thousands of lightly loaded servers
 - ▶ Hundreds of application instances
 - ▶ IT everywhere across the business
 - Physically
 - Logically
 - ▶ Thousands of distributed control points
- The Result:
 - ▶ Massive complexity
 - ▶ Spiraling people, power, cooling and server costs
- Compounded by the inability to allocate IT costs to the business
 - ▶ Lack of internal costing methodologies



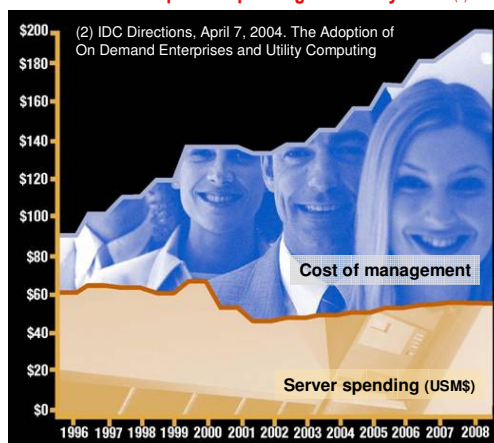
Virtualization and infrastructure mgt standards are the only hope to intercept these trends !

Cost of Complexity

- Operational efficiency (1)
- Pinpointing and managing issues
- Flexibility

Predicted average cost per end user by 2010 (1)	
Mainframes	\$6,250
UNIX® Minis	\$19,000
PC Servers	\$24,000

Cost of People vs. Spending on New Systems (2)



(1) 2010 prediction: Dinosaur myth, 2004 Update. Arcati Research Limited, Strategic Analysis for Enterprise IT. Based on current and projected end user costs for hardware, software, maintenance, and labor in mainframe environments

Energy Efficiency

Case study: Retailer consolidated 600 x86 servers onto a single mainframe

eliminated an **entire floor** of servers

reduced the administrative staff by 10

cut power and cooling costs by 80%



Source: RFG analyst report, "The Mainframe: It's Baaack!" July 2006

"Mainframe is the coolest place in the datacenter"




© 2006 IBM Corporation

THE WALL STREET JOURNAL

November 14, 2005

Power-Hungry Computers Put Data Centers in Bind

- *Today's distributed servers draw too much electricity and generate too much heat*
- If planners miscalculate, servers overheat, damaging circuitry or causing shutdowns
 - ▶ The Uptime Institute, an organization that represents data-center managers, predicts that power-related problems this year will cause four of the 20 major failures typically experienced by members annually, up from two of 20 last year. "The people who buy computers often aren't the people who have to manage them"
- Outcomes:
 - ▶ Rackspace Ltd., a San Antonio service that manages servers for clients, has seen its power needs swell to eight megawatts from three megawatts in the past three years -- sending its monthly utility bill up roughly fivefold to nearly \$300,000.
 - ▶ The University of Buffalo was surprised that their new system used 50% more power than planned, causing \$20K in electrical upgrades and \$150K in cooling upgrades.



Alternative Techniques for Removing Paint





Customer example

Consolidation z/VM & Linux on system z


- Customer is a distribution company
- Some core applications run on **two System z9** (model 705) but ...
- Most of new applications run on **hundreds of x86 Linux or Windows** servers distributed in 3 locations
- Main issues:
 - ▶ **Disaster recovery for distributed environment** is not efficient at all
 - ▶ Data centers may become full if the number of physical servers continues to grow
- **Server consolidation** using virtualization is key to support new business growth
- Initial scope of analysis: focus on **103 Linux x86 servers** (171 cores)
 - ▶ AMD Opteron and Intel Xeon processors (2.6 & 2.8 GHz) – mainly dual cores and some quad cores
- Scope reduced to a set of **75 servers** excluding:
 - ▶ Servers already consolidated using VMware,
 - ▶ Sysbase and PeopleSoft AS applications not available on Linux for System z platform

Application analysis

Linux x86 Software	Type	Linux z Software
Apache HTTP	Web application	Apache HTTP
WAS	Web application	WAS
VPSX	Pdf convertor	VPSX
Oracle DB	DB	Oracle DB
PeopleSoft DB (Oracle)	CRM application	PeopleSoft DB (Oracle)
RYO mail appl.	Mail servers	RYO mail appl.
RYO applications	Core applications	RYO applications
RYO SW distribution	Software distribution	RYO SW distribution
RYO XXX application	Core application	RYO XXX application
RYO EDI package	EDI application	RYO EDI package
RYO Network scripts	Network management	RYO Network scripts
CA Access Control	Security	CA Access Control
Veritas Netbackup	Backup management	Tivoli Storage Manager (TSM)

Check application availability on the target environment

Number of CPU cores per application group




**75 servers / 149 processors / 157 cores / 334 GB of memory installed
 459,524 RPEs installed**

Server activity during period of analysis

2 days of Server activity


Server Utilization during Consolidated peak time (14h10)



Average % Utilization = 5.82%
Maximum % Utilization during consolidated peak time = 7.70%

Consolidation effect on System z9

Total MIPS Consumed for All Servers for 48 hours on 07/01/08 in 5 minute intervals




157 x86 cores



11 System z9 IFLs (14 : 1 ratio)

Total MIPS

5-year Comparison and Cost structures



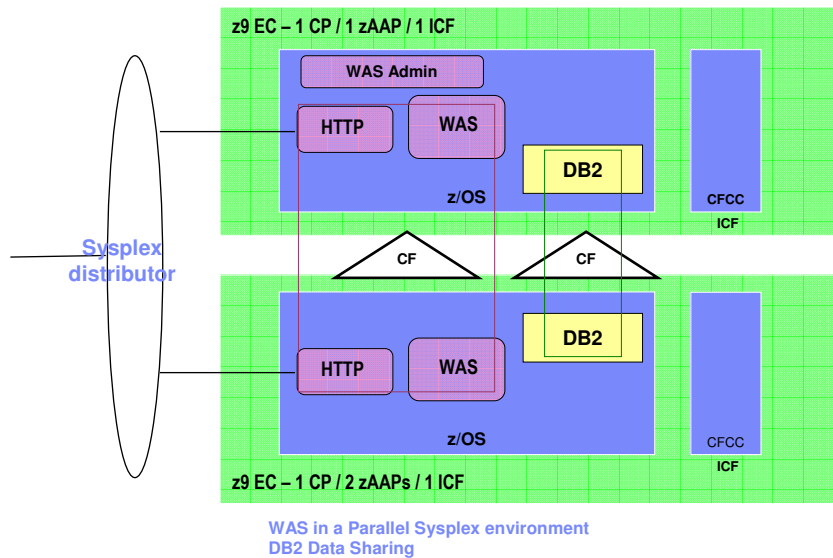
Customer benefits not included in TCO : Quality of service improvement, setup of an efficient Disaster Recovery plan, fast ROI, no need for a new Data Center

→ **Customer example** 2

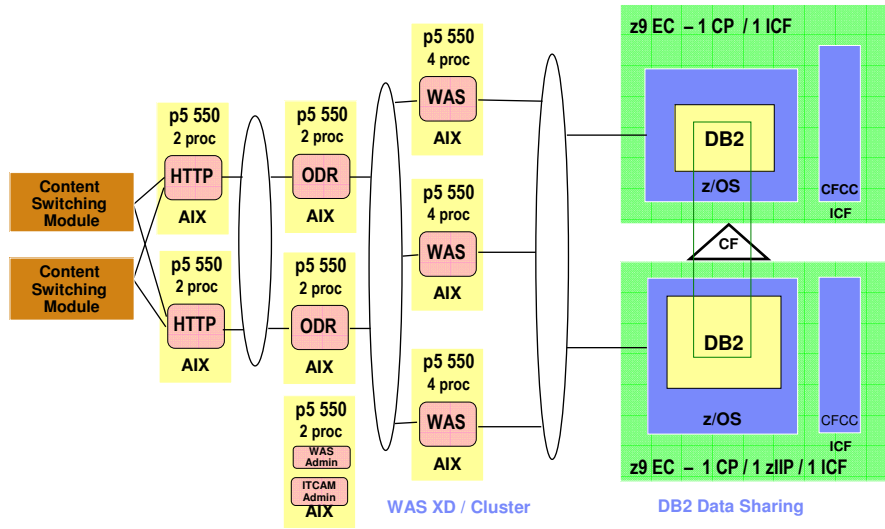
Deployment of new Java applications

- Large European banking customer
 - ▶ The bank wants to perform a platform technology selection for the deployment of **new Web applications**
 - ▶ Large System z9 environment running core banking applications, based on IMS and DB2
 - ▶ Part of customer's centralization strategy, **new databases will be implemented on the current System z9** platforms based on DB2 for z/OS
 - ▶ Technical and cost assessment study to compare WebSphere applications running on **AIX vs. z/OS**, accessing a DB2 database located in a z/OS System z9 LPARs in a Parallel Sysplex environment.
 - ▶ Application utilization estimated to be at **100 transactions/second** during peak time.

System z solution architecture (100 TPS)

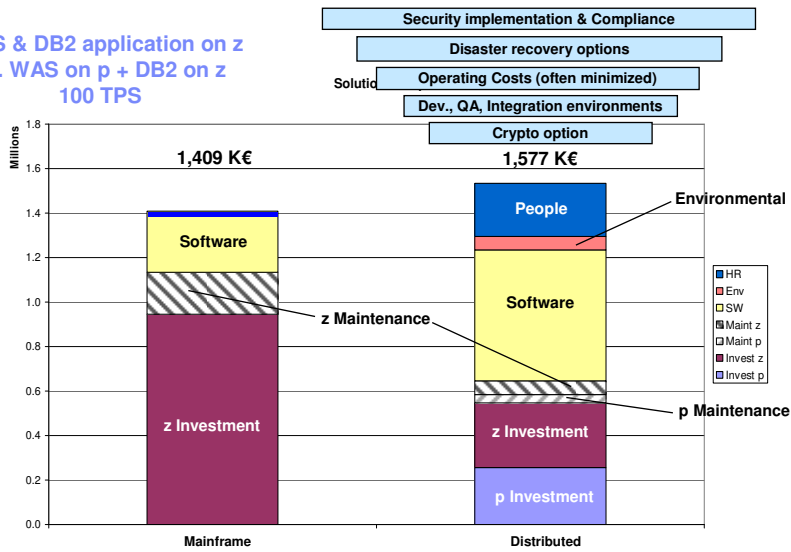


System p+z solution architecture (100 TPS)



3-year Cost Comparison @ 100 TPS

WAS & DB2 application on z vs. WAS on p + DB2 on z
100 TPS



→ **Customer example**

4

Telco Industry Benchmarks Allow Direct Comparison of HP Superdome to IBM z9


HP Superdome – 64 x 1.5GHz Itanium2



64 Unix processors


Call initiation rate:
1,000/second

z9 24-way Benchmark



24 processors


Call initiation rate:
6,568/second



Sources: CommuniGate-Superdome-VoIP-Benchmark.pdf & IBM-CommuniGate-z9.pdf from <http://www.communigate.com/Papers>


Telco Industry Benchmarks Allow Direct Comparison of HP Superdome to IBM z9

HP Superdome – 64 x 1.5GHz Itanium2



(64 x 6,568/1,000)

z9 24-way Benchmark

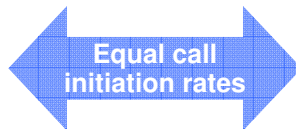


420 Unix processors
(719,590 RPE's)

96 RPE's/MIP

24 processors
(7,509 MIPS)


Call initiation rate:
6,568/second



Call initiation rate:
6,568/second


Sources: CommuniGate-Superdome-VoIP-Benchmark.pdf & IBM-CommuniGate-z9.pdf from <http://www.communigate.com/Papers>

Telco Industry Benchmarks Allow Direct Comparison of HP Superdome to IBM z9



- Energy consumption comparison (without the energy for cooling)
 - Superdomes: 1401600 kw/y
 - Z9 EC: 95484 kw/y

Telco Industry Benchmarks Allow Direct Comparison of HP Superdome to IBM z9



- Heat production/waste of the solutions
 - Superdomes: 541000 BTU/h ca. 180 PKW Aircooling unit
 - Z9 EC: 32000 BTU/h ca. 11 PKW Aircooling unit

Telco Industry Benchmarks Allow Direct Comparison of HP Superdome to IBM z9

HP Superdome – 64 x 1.5GHz Itanium2




\$23.8M
(HW, SW for 3 years)

k€ 1220 *
(Energy for 3 years)

Total: about US\$ 25,5 M

z9 24-way Benchmark



\$10.3M
(HW, SW for 3 years)

k€ 83 *
(Energy for 3 years)

Total: about US\$ 10,5 M

· *) Energy costs in Germany 2007 – 14.5 Cent / kw/h.
· UPS and availability/recovery features are not evaluated

Customer Studies

Representative customer - Websphere

Hardware

- 5000+ MIPS
- 1000+ servers (25% Unix)

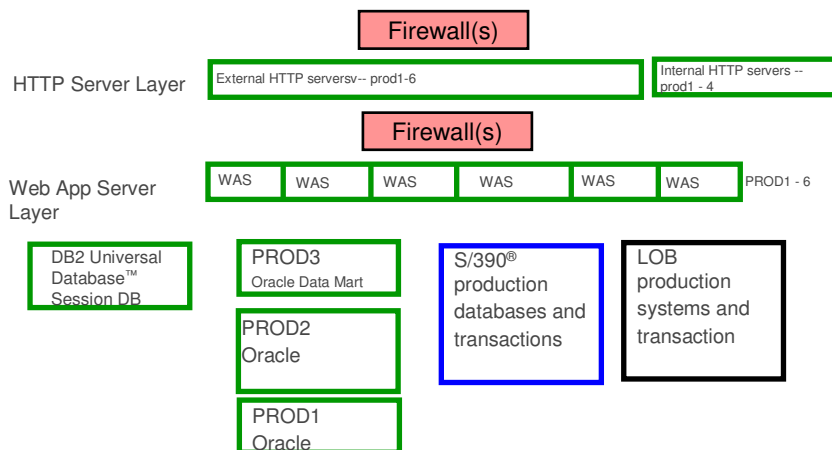
Software

- Websphere currently on Solaris
- Oracle and DB2

Customer perception:
Solaris environment is 1/5 the cost of the mainframe



A Real Client Example of finding "All the Costs" Current Production WebSphere Environment



Source: Scorpion Study 1999 - 2004

The Total Cost of Acquisition may be a Misleading Indicator for Large Enterprise Deployments


Source: Gartner; ID Number: IGG-03102004-01 - CIO Update: The March of Linux in the Enterprise

George J. Weiss

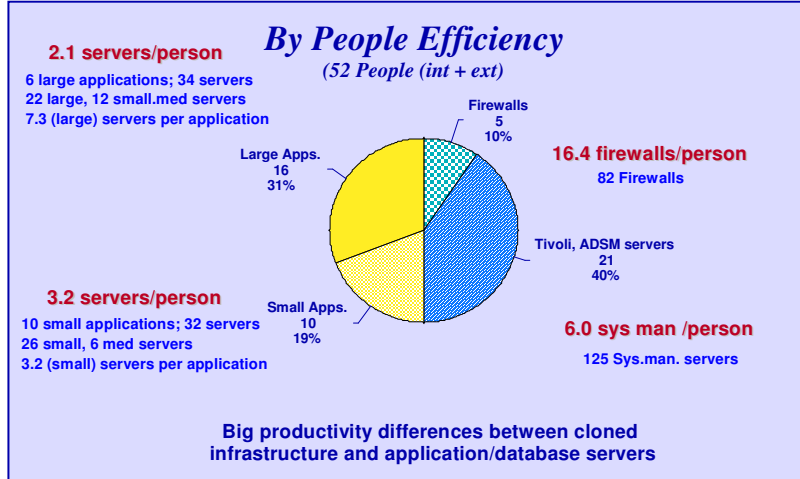
Vice President and Distinguished Analyst in Gartner Inc.

Only the Holistic View will give you the Entire Picture!

Server Analysis




Server Analysis



Server Demographics – Avg Unix and Intel CPU Utilization

This graph represents the portfolio of customer studies performed from 2000 thru 2004 categorized by average CPU utilization



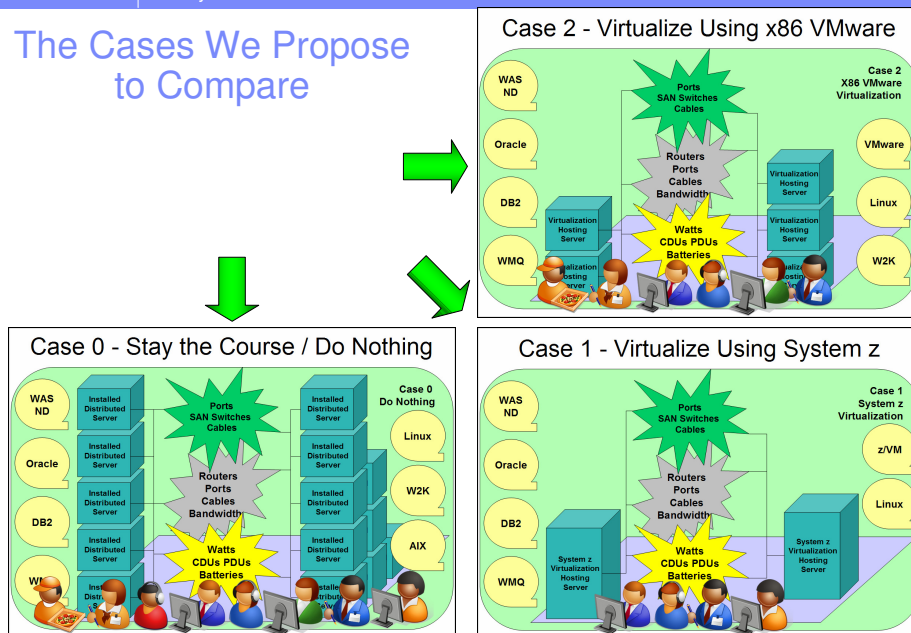
What is a zRACE Rapid Assessment?



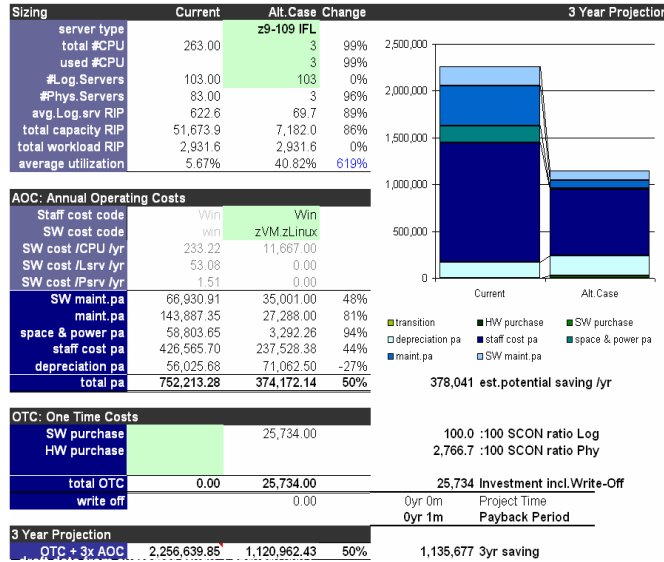
▪ zRACE:

- ▶ Is an assessment methodology used to develop a Total Cost of Ownership scenario comparison for our clients ...
- ▶ Is used to evaluate potential benefits of consolidating workloads from Intel and/or UNIX platforms onto System z IFLs (specialty engines running Linux) or WebSphere Applications on zAAP engines
- ▶ A 5 year TCO analysis comparing current case versus and a distributed alternative ...
- ▶ Takes advantage of new or existing System z footprints) as the target environment ...
- ▶ Uses a combination of assumptions, estimates, industry numbers and actual client data to develop cost models

The Cases We Propose to Compare



Zodiac helps build solution-by-solution business cases

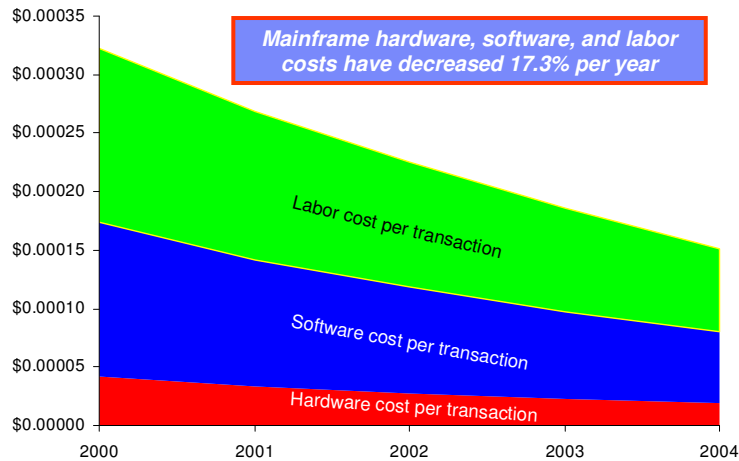


... and business cases that show a Greener Environment!

Environmentals	Current	Alt. Case	Change	Difference
avg RackU / Server	6.4	3.0	53%	3.4
Total RackU	533	9	98%	523.5
30U Racks	18	0	98%	17.5
Total kW	109	1	99%	109
Adjusted kWh/yr	961,428	5,282	99%	956,146
Heat BTU/hr	249,715	1,372	99%	248,343
CO2 tonnes /yr	411	2	99%	409
Carbon tonnes /yr	112	1	99%	112
RIPs /kW	473	11,970	2430%	11,497
RIPS / tonne CO2	126	3,162	2417%	3,036
W /m2	12,305	4,000	67%	8,305

3yr CO2 Needs 4,059 Trees

Mainframe Cost Trends



Kevin Campbell [Video](#) UNIVAR

“For Univar, the ROI model that we have used for this consolidation takes into account all of the hard costs that we’re comfortable with predicting;

- the elimination of hardware leases,
 - the elimination of software maintenance
- suggests that we should see a return on our migration investment within three to four years.

What that doesn’t try to quantify are the soft benefits such as simplifying the process we currently have to move data from platform to platform. It doesn’t attempt to quantify the costs inherent in maintaining all of that. Also - eliminating racks of equipment - we have drastically reduced the cooling and power supply demands on our datacenter. All of these are what we regard as soft benefits that are well worth having, but which we didn’t attempt to quantify”.

Driving Down Costs Requires Knowing What You Have

If there aren't accurate measures . . . how do you know it's effective?

. . . on average, approximately 15% of Global IT budgets are attributable to mainframe-related purchases, contracts, and activities, but, at the same time, 25-30% of the IT budget is recovered via billing for mainframe-resident services . . .

. . . in the past, chargeback systems focused on isolating IT system events that could be relatively easily tracked and could be shown to generate sufficient 'revenues' to cover the IT budget. Now, however . . . , the focus is on building systems that reflect the real underlying relationship between IT resource consumption and cost accrual. Despite this change in focus, mainframe platforms remain the keystone for chargeback architectures, particularly in the financial services industry . . .

Will Cappelli, Vice President, META Group

IBM Mainframes provide technology and tools to accurately track and report resources consumption in mixed workload environments to help better manage cost and improve investment decisions

Meta Consulting has developed a chargeback methodology to enable enterprises to allocate costs more equitably and accurately across all platforms in the entire IT infrastructure.

Solutions from newly acquired Isogon provide customers the tools they need to manage their software costs.

Logical Steps

1. 'Full' Inventory of IT infrastructure servers
2. Group the Servers into 'homogeneous' ISLANDs
 - A manageable ISLAND should contain not more than 500 servers (300 is the best number)
 - Group the servers by Location/Application/Function, not by platform
 - A typical Island should contain less than TEN different applications
 - The application grouping is very important for understanding the consolidation potential
3. Rank the servers by consolidation potential (A,B,C,D....)
 - Easy to consolidate on any platform (Infrastructural Servers)
 - Easy to consolidate on a similar platform (Data Base or Middleware)
 - Not so easy to consolidate (Need for Porting process)
 - Cannot consolidate
 - Out of Scope
4. Run the Zodiac or zRace Business Cases by using:
 - The Application/The Consolidation Potential
5. Run Zodiac or zRace
 - Obtain a cost picture
 - Obtain a savings picture
 - Obtain a target model (no more than one)

Mainframe Advantages


- Least expensive except for smallest multi-user systems
- Biggest factors in wide Total-Cost-per-User differences
 - ▶ Much lower operating/support staff level/costs on mainframe compared to UNIX and Windows
 - ▶ Mainframe software costs very visible; distributed costs hidden and duplicated
- UNIX or Windows distributed platforms
 - ▶ Require 2.5X to 3X more staff than today's mainframe to support similar workloads
 - ▶ Profusion of server/storage hardware and software needed in enterprise configurations, often 10-fold more "iron" than zSeries
- Dramatic reductions in mainframe staffing levels
 - ▶ 10-fold reduction in mainframe staffing (operators and systems programmers) per MIPS over last 7 years
 - From IBM's major advances in self-healing, self-managing, self-protecting, autonomic technologies for the mainframe
 - ▶ Scalability – doubling users adds 90% for mainframe but 125% for distributed
 - ▶ Expect continuing mainframe advances will halve this again over next five years

Source: Arcati

How Could This TCO Scenario Be True?






IBM System z 

Bibliography

- IBM System z: Platform Star for Linux Open Source Software. Ptak, Noel & Associates, November 2006.
- The Greening of IT: Cutting Power Consumption is Just the Beginning. Enterprise Systems, November 2006.
- IBM System z: Platform Star for Linux Open Source Software. Ptak, Noel & Associates, November 2006.
- The new TCO and the value of the mainframe. The Mainstream, January 11, 2007.
- IBM System z TCO: Man Bites Dog? Illuminata Inc., November 29, 2006
- SearchDataCenter: Mainframe server consolidation should be on your hit list
- Research Brief - Does System z Offer Better TCO than Scale-up, Distributed Unix and Linux Servers? Clabby Analytics, Nov 2006.
- Research Brief - IBM's System z: Technology Update. Clabby Analytics, November 2006.
- SearchDataCenter: Mainframe TCO getting more bang for the buck, analyst report says. By Mark Fontecchio December 20, 2006.
- Mainframe As A Green Machine -- And More: ROI -- 2007 to 2013, Susan Eustis, Wintergreen Research, October 2005



SHARE Orlando 2008 Session 9261 88