#SHAREorg

SHARE
Technology · Connections · Results

IBM zBladeSizer and IBM zBladeEXTR zBX Capacity Planning Tool and Extractor

Chuck Hackett
IBM Corporation

February 7, 2013
Session 13091

SHARE in San Francisco 2013

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

AlphaBlox*
APPN*
CICS*
CICS/VSE*
Cool Blue
DB2*
DFSMS
DFSMShsm
DFSMSrmm
DirMaint
DRDA*
DS6000
DS8000
ECKD
ESCON*
FICON*
FlashCopy*
GDPS*
HiperSockets
HyperSwap
IBM*
IBM eServer
IBM logo*
IMS
Language Environment*
Lotus*
Large System Performance Reference™ (LSPR™)
Multiprise*
MVS
OMEGAMON*
Parallel Sysplex*
Performance Toolkit for VM
PowerPC*
PR/SM
Processor Resource/Systems Manager
RACF*
Redbooks*
Resource Link
RETAIN*
REXX
RMF
S/390*
Scalable Architecture for Financial Reporting
Sysplex Timer*
Systems Director Active Energy Manager
System/370
System p*
System Storage
System x*
System z
System z9*
System z10
Tivoli*
Tivoli Storage Manager
TotalStorage*
VSE/ESA
VTAM*
WebSphere*
xSeries*
z9*
z10
z10 BC
z10 EC
z/Architecture*
Zenterprise 196*
z/OS*
z/VM*
z/VSE
zSeries*

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:
Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.
IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.
All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.
This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.
All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.
Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.
Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Agenda

- Intro
- IBM zBladeSizer and IBM zBladeEXTR
- What problem is being solved
  - How are we solving it
  - Input data and customer controls
  - Optimization
- Documented results
- Demo
- Q+A

The Problem: Server Migration to the zBX

- Find an optimum solution to a zBX server migration scenario
- Document the solution with a comprehensive set of reports and graphs
- Rapidly investigate alternative scenarios

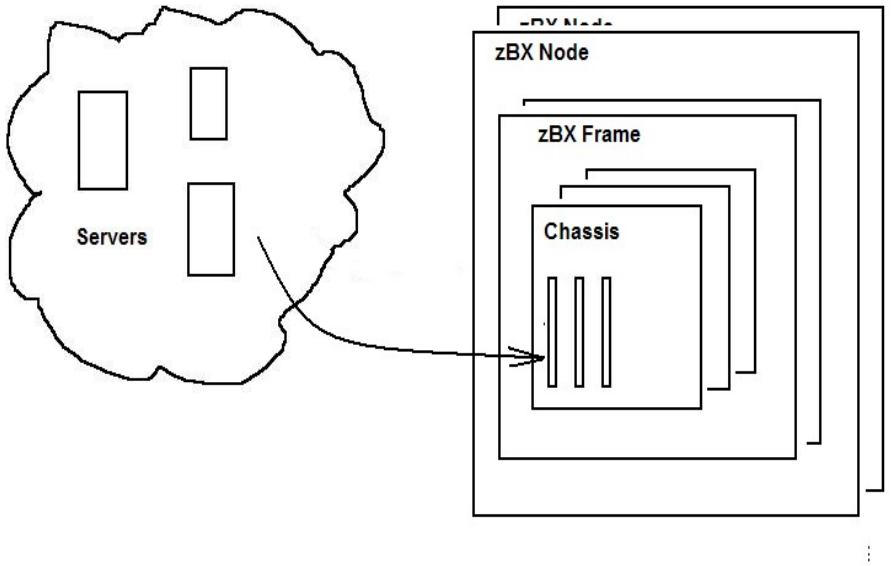

zBX Hardware Overview

1 to 8 nodes per Ensemble

2 TORs per Node

1 to 4 frames per node

1 to 2 BC per frame

1 to 14 blades per BC

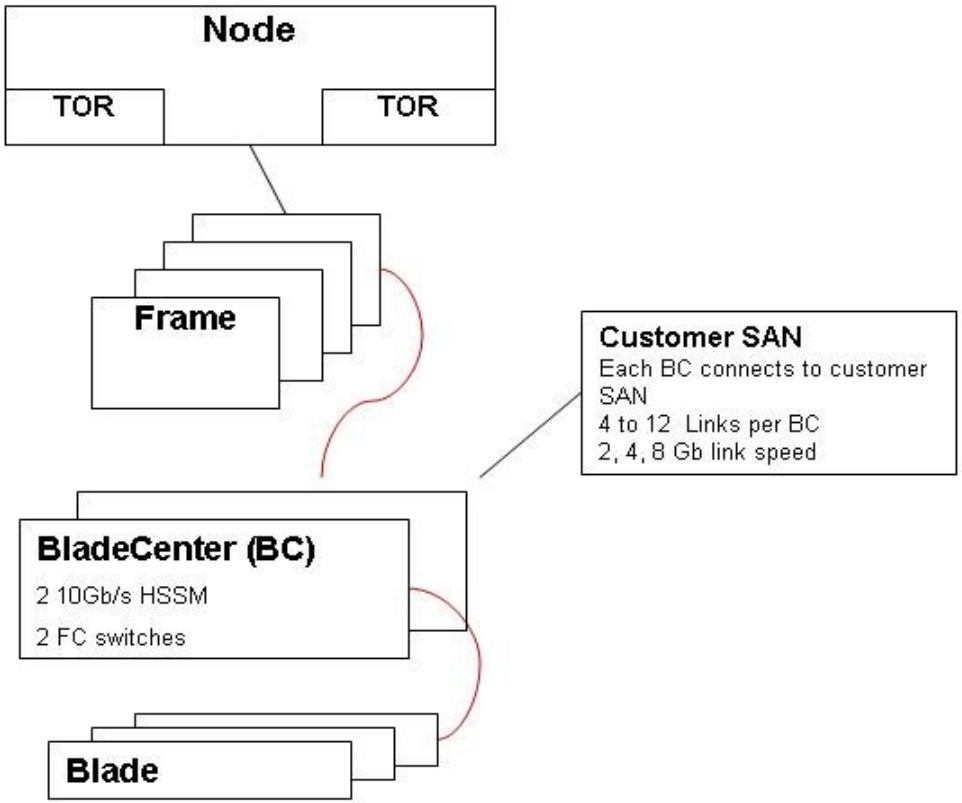

Systems configuration and performance data collected by zBladeEXTR (Power or x86)

Techline receives data and uses zBladeSizer to build a right-sized zBX configuration.

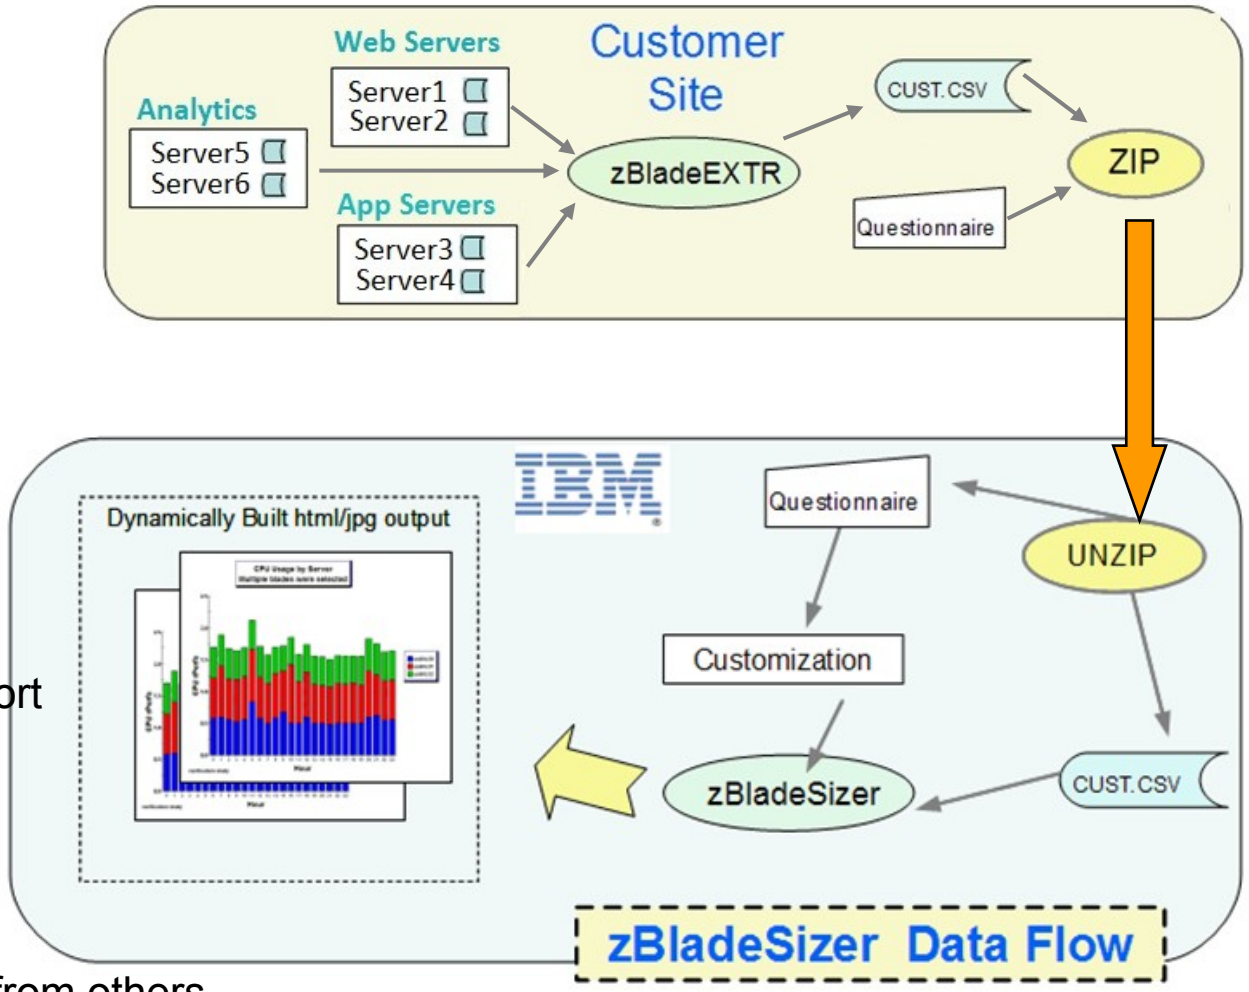

Benefits:

- Rapidly generated solution
- Easy generation of a solution report
- Interactive and easy to use
- Tailorable solution parameters:
  - zBX hardware parameters
  - Segregation of groups of servers by application
  - Isolation of specific servers from others
- Resource demands mapped to a 24 hour profile configure an optimum zBX solution, not just peak hour.
- A solution based on actual performance data including CPU, network and disk I/O, and memory

IBM zBladeEXTR

- Extract intervals from collected performance data
  - Customer can run zBladeEXTR and send file to IBM - OR -
  - Zip and ship data * to IBM to run
- Validity checking
- Output a CSV and error/summary reports
- Prompted to fill out a questionnaire of customer business controls
- Help for data collection:
  - data collection guide
  - zBladeEXTR process wizard
  - Support from the IBM ATS organization

* Performance Data Source:
  - NMON data for System P
  - SAR for System X server LINUX systems

zBladeEXTR Data Gap Analysis

Is this the right set of data for analysis?

- Navigation
- Current status
- Step Detail

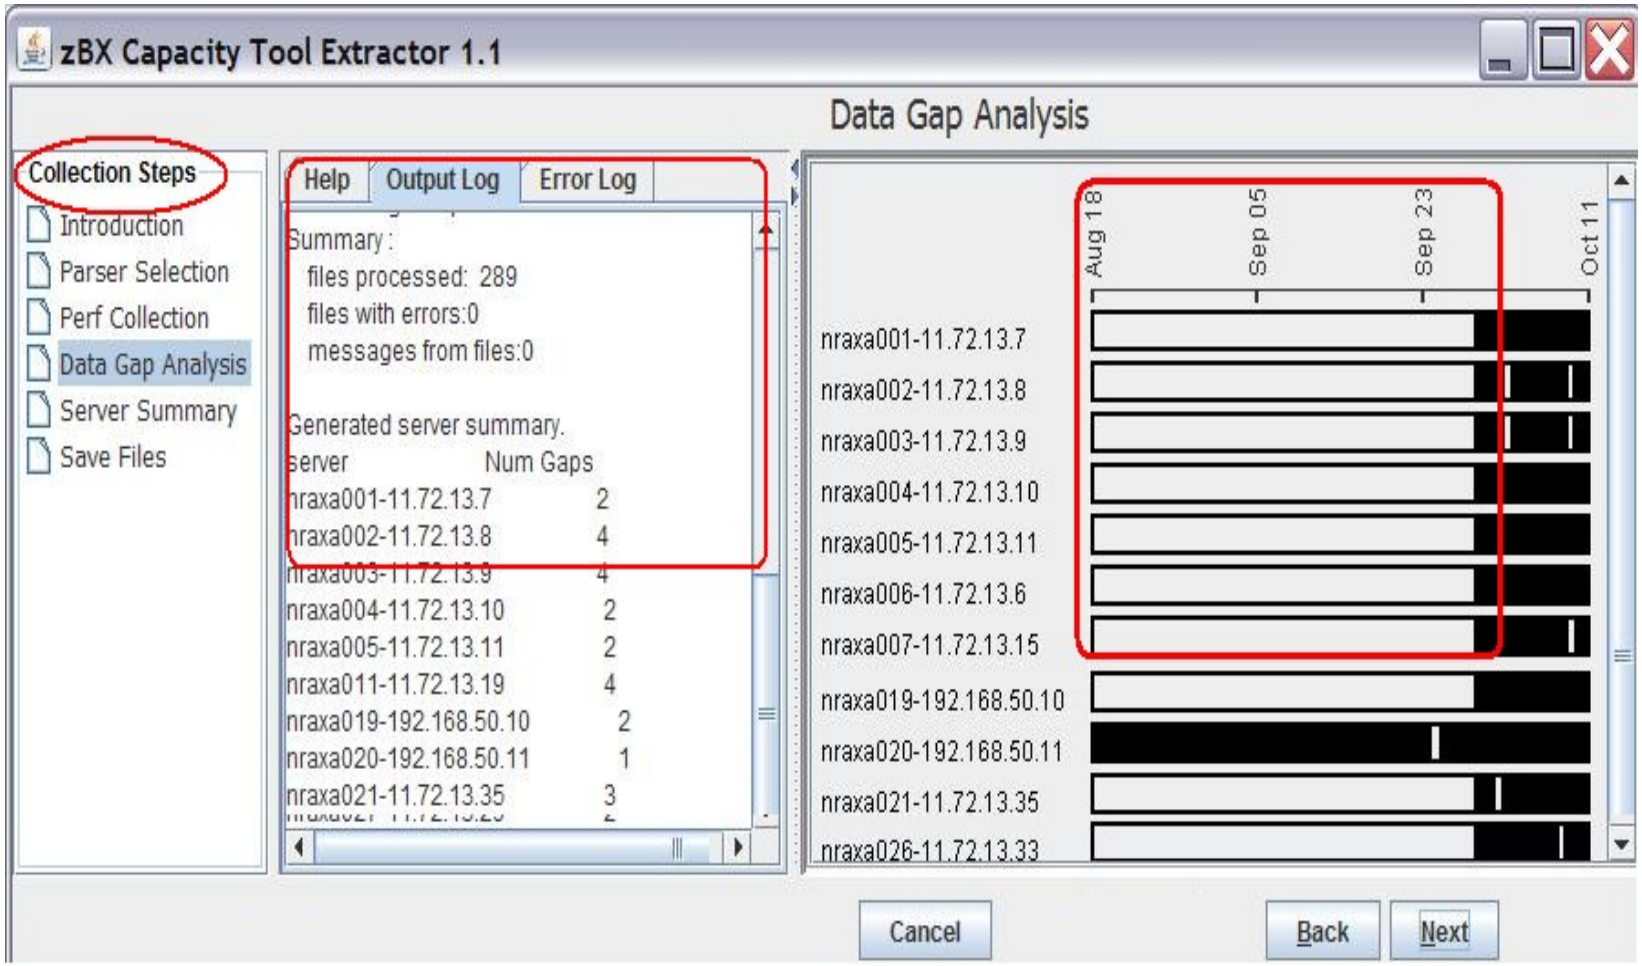

zBladeEXTR Questionnaire
Application Server Groups

Groups are used to demark business functions and priorities

Servers are assigned to a group, and each group of servers can:

- Share blades
- Have different peak hours (or default)
- Have different CPU and memory usage constraint values (for warn and limit).
- Have different growth rates
- Order their placement on zBX blades.

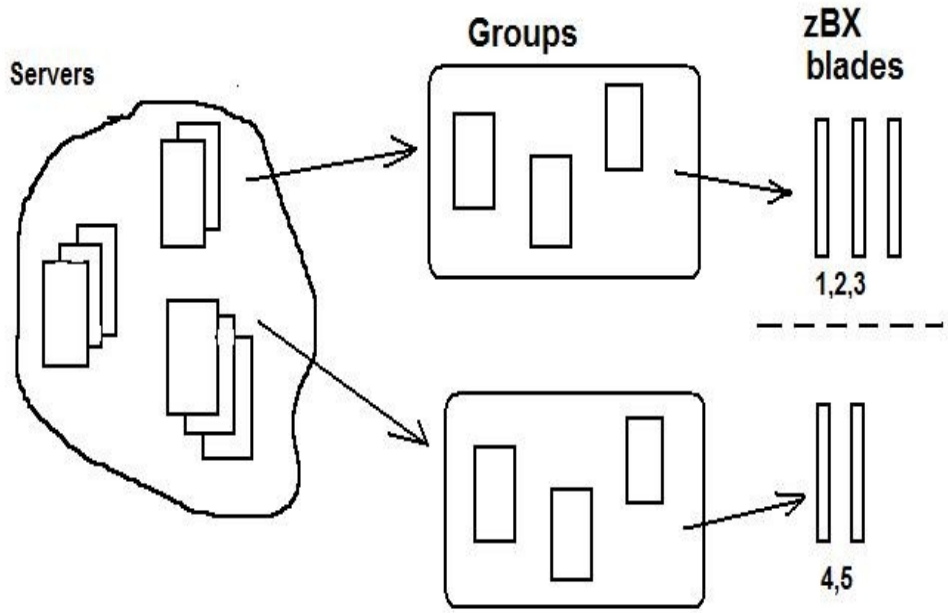

zBladeEXTR Questionnaire
Isolation Subgroups

Isolation subgroups separate servers

- Define isolation subgroups within a group
- Servers in an isolation subgroup will NOT share a blade

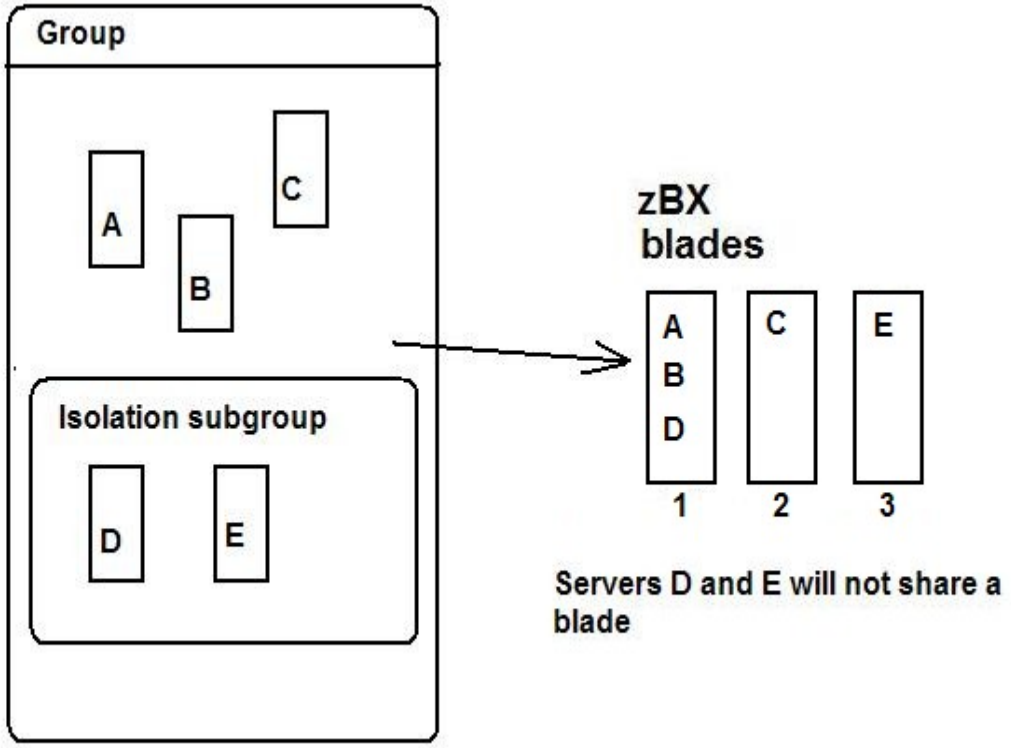

zBladeEXTR **Questionnaire Hardware**

- Blade selection
  - PS701 32,64,128GB
  - E7-2830 64,128,192,256GB
- Number of reserved blades per chassis
- I/O paths (number of links and link speeds)
- I/O path usage constraints for each of the 5x2 connections (10 total because the links are full duplex)

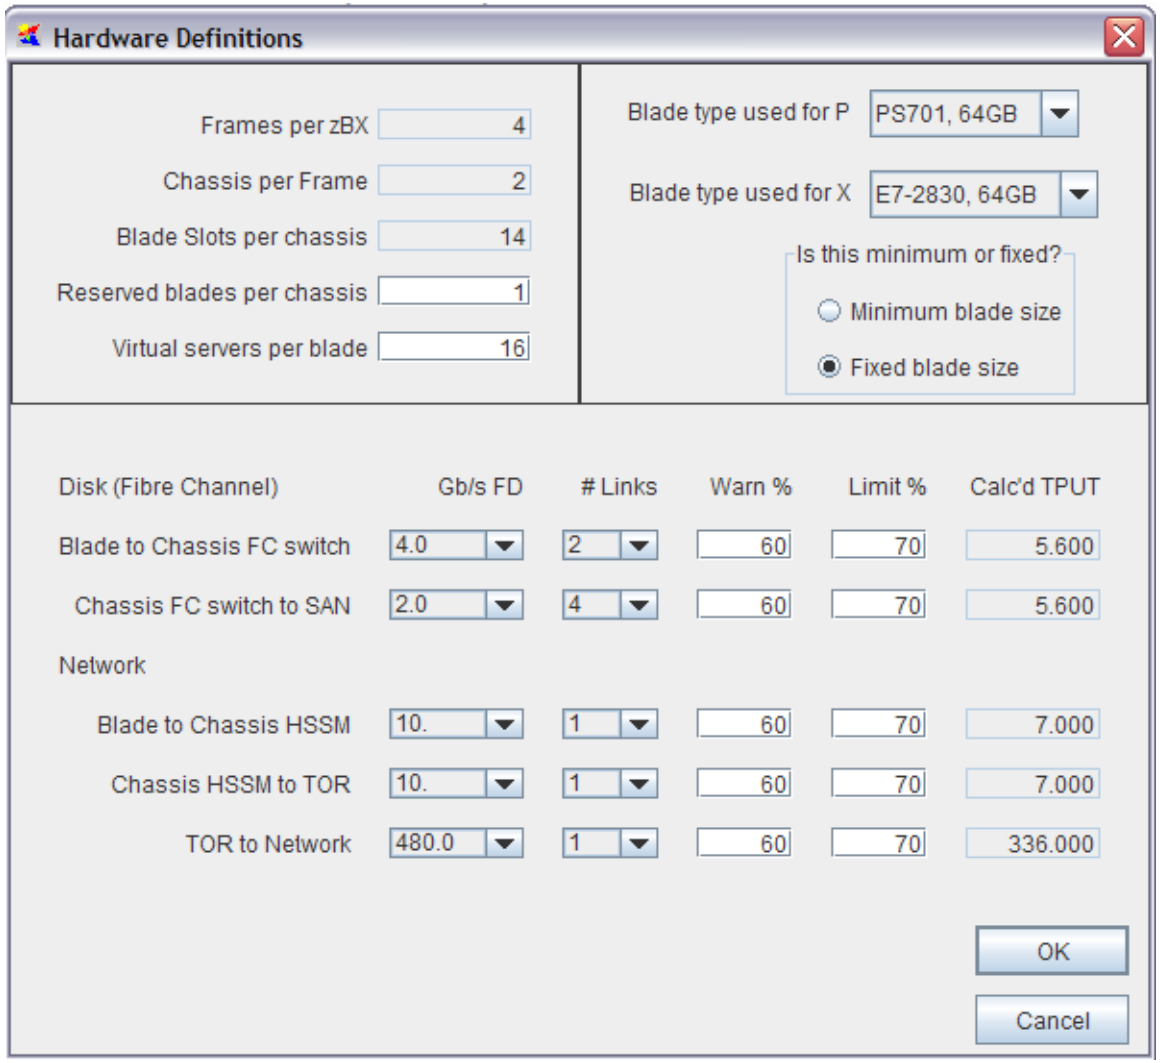

zBladeEXTR Questionnaire
Capacity Planning Considerations

Specification		Customize It
▪ Growth Rate	→	By Group or Server
▪ Peak Hour Selection	→	Selected or Calculated
▪ Solution Constraints	→	Maximum allowable CPU and Memory % by group, and disk and network IO bandwidth % by configuration

zBladeSizer: How it Solves the Server Migration Problem

- Data from server collection agents
- 24 hour statistical profiles each server, each resource
- Customer business controls
- Search for optimum solution
- Output best configuration
  - graphs and analysis
  - deliver report to customer

Optimization in zBladeSizer

- Resource supplies from hardware definitions and usage constraints
- Resource demands from server usage data (and growth rates)
- Supply and demand are matched to find the configuration with the least number of blades
- The hybrid algorithm runs in parallel by solving each application server group independently, then combines the group results into a final configuration
  - For small sized groups, uses an exhaustive search
  - For larger groups, uses a genetic algorithm
  - Heuristics employed to reduce turnaround time

Outputs from zBladeSizer

- Graphs and reports showing usage for 6 resources:
  - CPU
  - Memory
  - disk read and write IO
  - network read and write IO
- Shown by relevant hardware component,
  - Blade, chassis (blade center), frame, node
- Shown by group and server
- Reports show the layout of the zBX with server to blade assignments
- Reports can be saved to disk as HTML for import into a word processor

Server to Blade Assignments

- Sample main view of the application showing server to group assignments and blade assignments

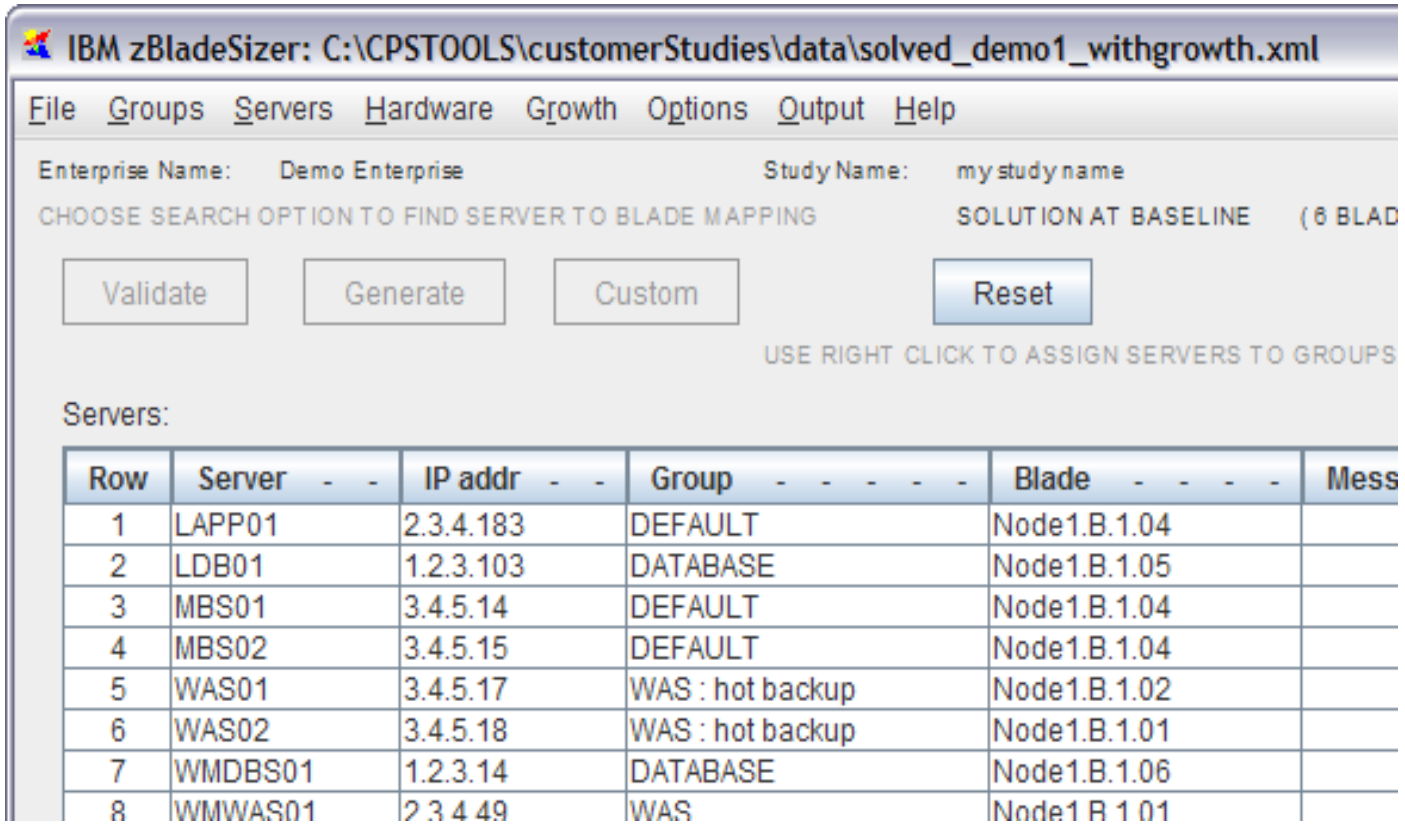

Row	Server	IP addr	Group	Blade	Mess
1	LAPP01	2.3.4.183	DEFAULT	Node1.B.1.04	
2	LDB01	1.2.3.103	DATABASE	Node1.B.1.05	
3	MBS01	3.4.5.14	DEFAULT	Node1.B.1.04	
4	MBS02	3.4.5.15	DEFAULT	Node1.B.1.04	
5	WAS01	3.4.5.17	WAS : hot backup	Node1.B.1.02	
6	WAS02	3.4.5.18	WAS : hot backup	Node1.B.1.01	
7	WMDBS01	1.2.3.14	DATABASE	Node1.B.1.06	
8	WMWAS01	2.3.4.49	WAS	Node1.B.1.01	

Servers on Blades

- Sample report output showing server assignment to blades, by blade. Also has isolation subgroup “hot backup”

Node	Chassis	Blade Number	Blade Model	Sizing Group	Isolation SubGroup	Virtual Server	Server IP address	Peak Hour rPerfs	Max Memory (GB)	Ne… Ban… (M…
Node1	B.1	01	PS701	WAS		WMWAS01	2.3.4.49	1.04	8	
			PS701	WAS	hot backup	WAS02	3.4.5.18	0.17	4	
		02	PS701	WAS	hot backup	WAS01	3.4.5.17	0.46	4	
		03	PS701	DEFAULT		app1	5.6.7.63	9.20	4	
			PS701	DEFAULT		asprd	2.3.4.91	0.14	16	
			PS701	DEFAULT		rdb1prod	6.7.8.15	2.60	36	
		04	PS701	DEFAULT		app2	5.6.7.64	0.91	4	
			PS701	DEFAULT		LAPP01	2.3.4.183	0.93	8	

2 Sample Charts Following

- CPU Usage by blade
  - The highest average CPU usage is at 35%
  - the CPU is relatively lightly used and does not stress the capacity of the PS701
  - one blade is very low usage, and that was because of the way the groups were setup

- Memory usage by blade
  - Can see how the memory has significant resource demands and contributes to the sizing of the zBX
  - All blades are sized identically

Example: CPU Each Blade

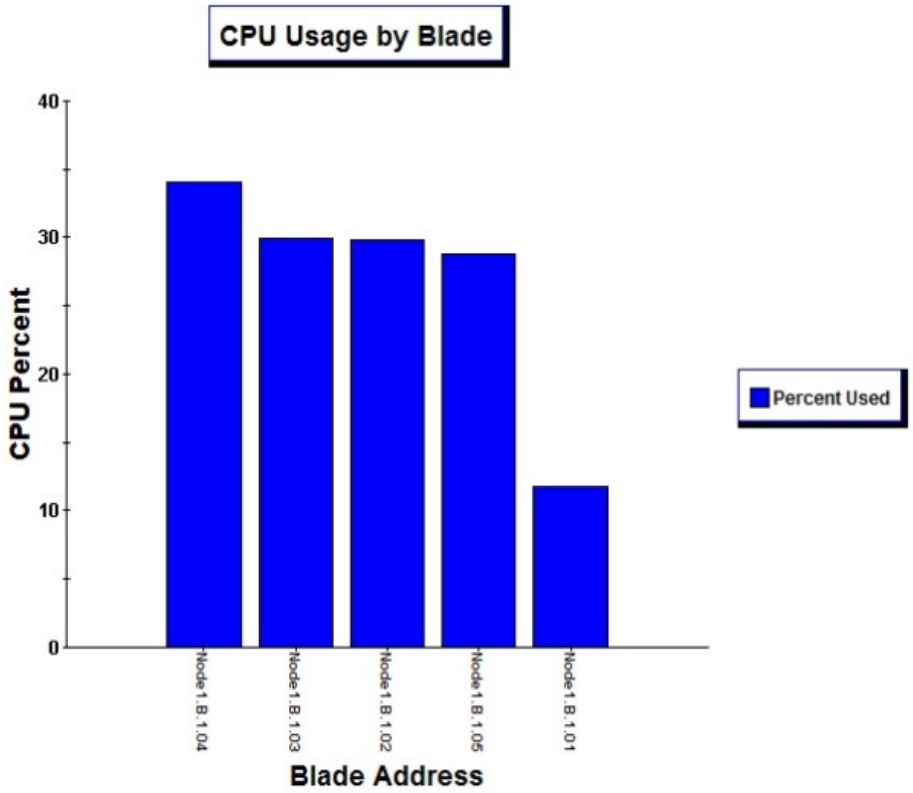

Blade	CPU Percent Usage	Sizing Group	Peak Hour	Was peak selected or calculated?
Node1.B.1.04	34.09%	DATABASE	1	Calculated
Node1.B.1.03	29.90%	DEFAULT	1	Calculated
Node1.B.1.02	29.85%	DEFAULT	16	Calculated
Node1.B.1.05	28.77%	DATABASE	1	Calculated
Node1.B.1.01	11.74%	WAS	22	Calculated

Example: Memory Each Blade

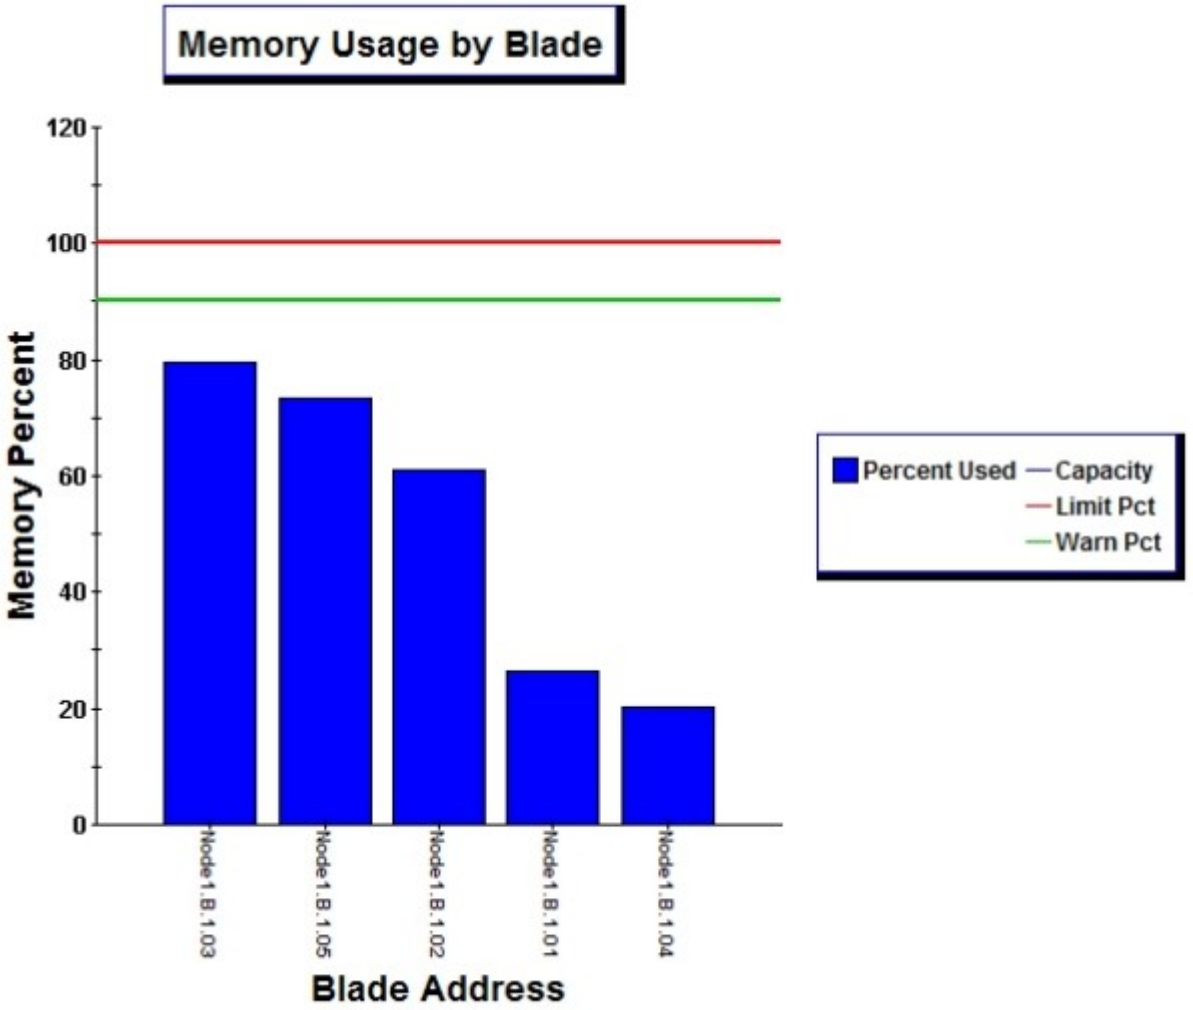

X86 Server Platform Identification

- SAR data isn't as comprehensive as NMON, so we need to identify the source server platform
- Using vendor data which may change to a different vendor
- Previous decisions are retained for a rapid restart
- Filtering dialog to provide assist

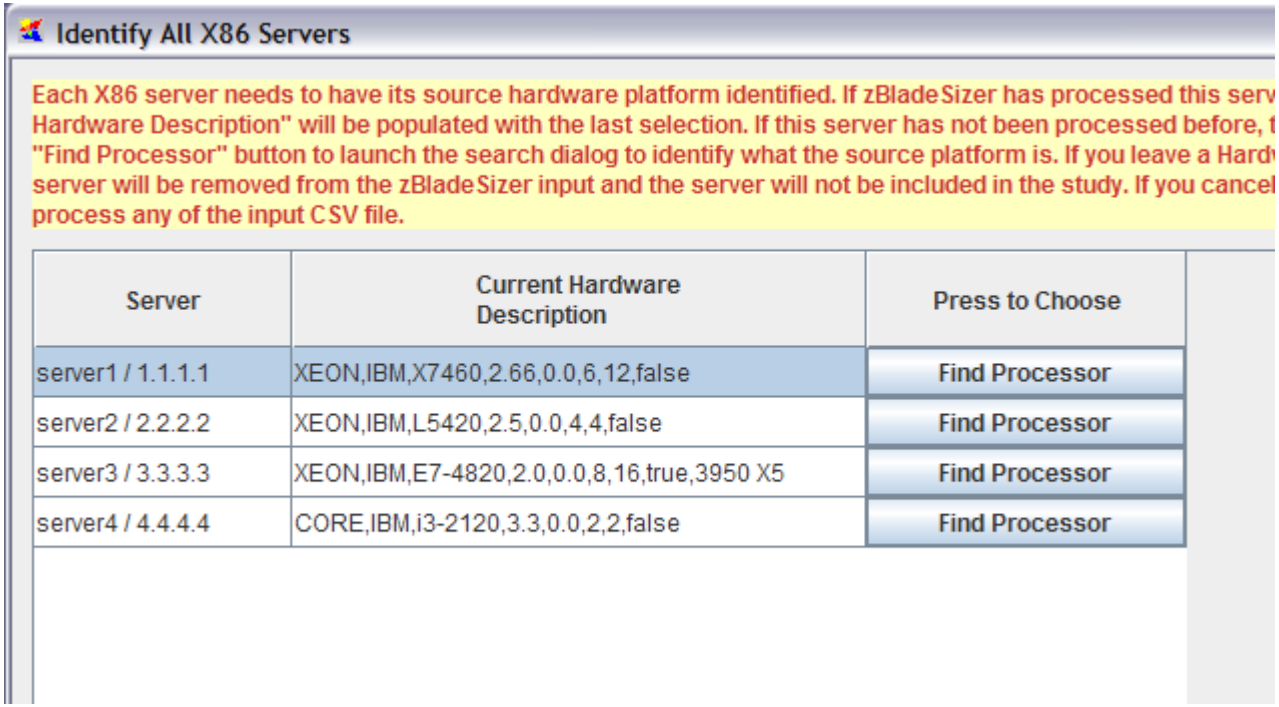

Selecting the X86 platform

- Display gathered info
- Each column has an automatic filter

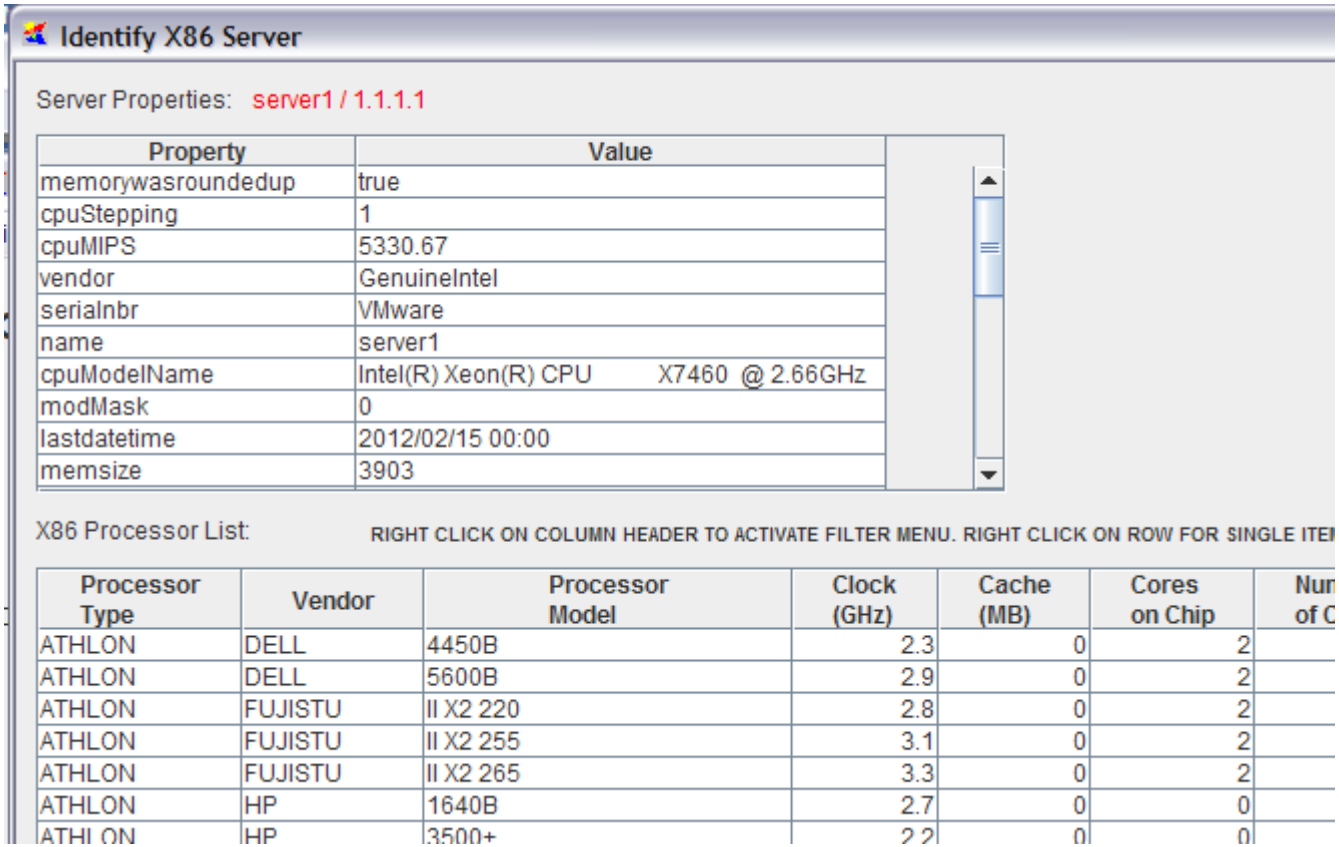

Identify X86 Server

Server Properties: server1 / 1.1.1.1

Property	Value
memorywasroundedup	true
cpuStepping	1
cpuMIPS	5330.67
vendor	GenuineIntel
serialnbr	VMware
name	server1
cpuModelName	Intel(R) Xeon(R) CPU X7460 @ 2.66GHz
modMask	0
lastdatetime	2012/02/15 00:00
memsize	3903

X86 Processor List: RIGHT CLICK ON COLUMN HEADER TO ACTIVATE FILTER MENU. RIGHT CLICK ON ROW FOR SINGLE ITE

Processor Type	Vendor	Processor Model	Clock (GHz)	Cache (MB)	Cores on Chip	Nur of C
ATHLON	DELL	4450B	2.3	0	2	
ATHLON	DELL	5600B	2.9	0	2	
ATHLON	FUJISTU	II X2 220	2.8	0	2	
ATHLON	FUJISTU	II X2 255	3.1	0	2	
ATHLON	FUJISTU	II X2 265	3.3	0	2	
ATHLON	HP	1640B	2.7	0	0	
ATHLON	HP	3500+	2.2	0	0	

Column filter, example using Processor

- Cascading menu for choosing from a big collection

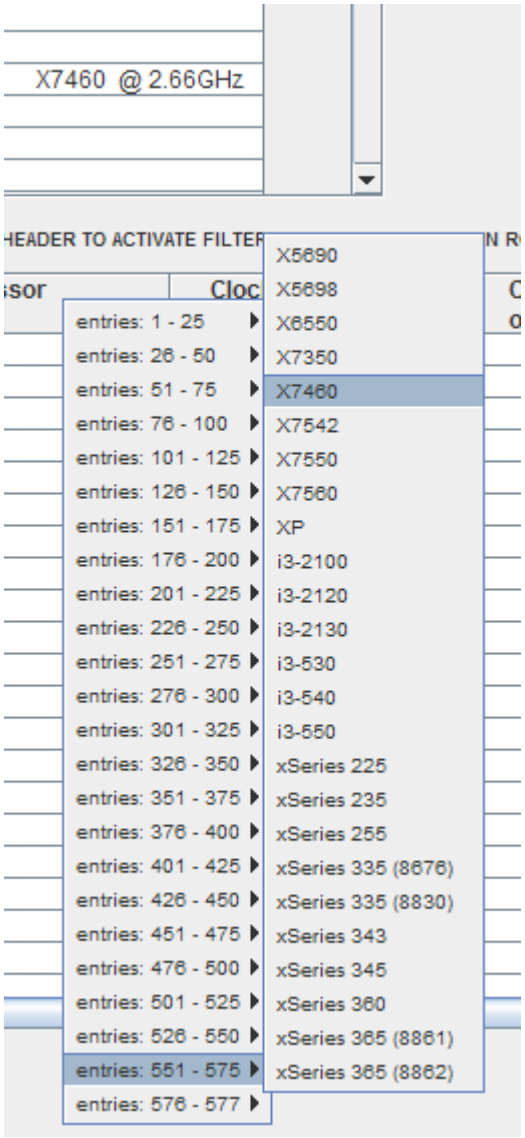

Plans for CY2013

- Data Preview (example next slide), with interactive excludes
- Change P servers to use percent of blade (zBPCI) like X servers
- Combine P and X servers in the same study
- Secondary studies after a zBX already installed
  - Phase 1: hardened, start counting beyond
  - Phase 2: zManager for CPU, I/O
- Summary by day

Preview Data (v1.5 under development)

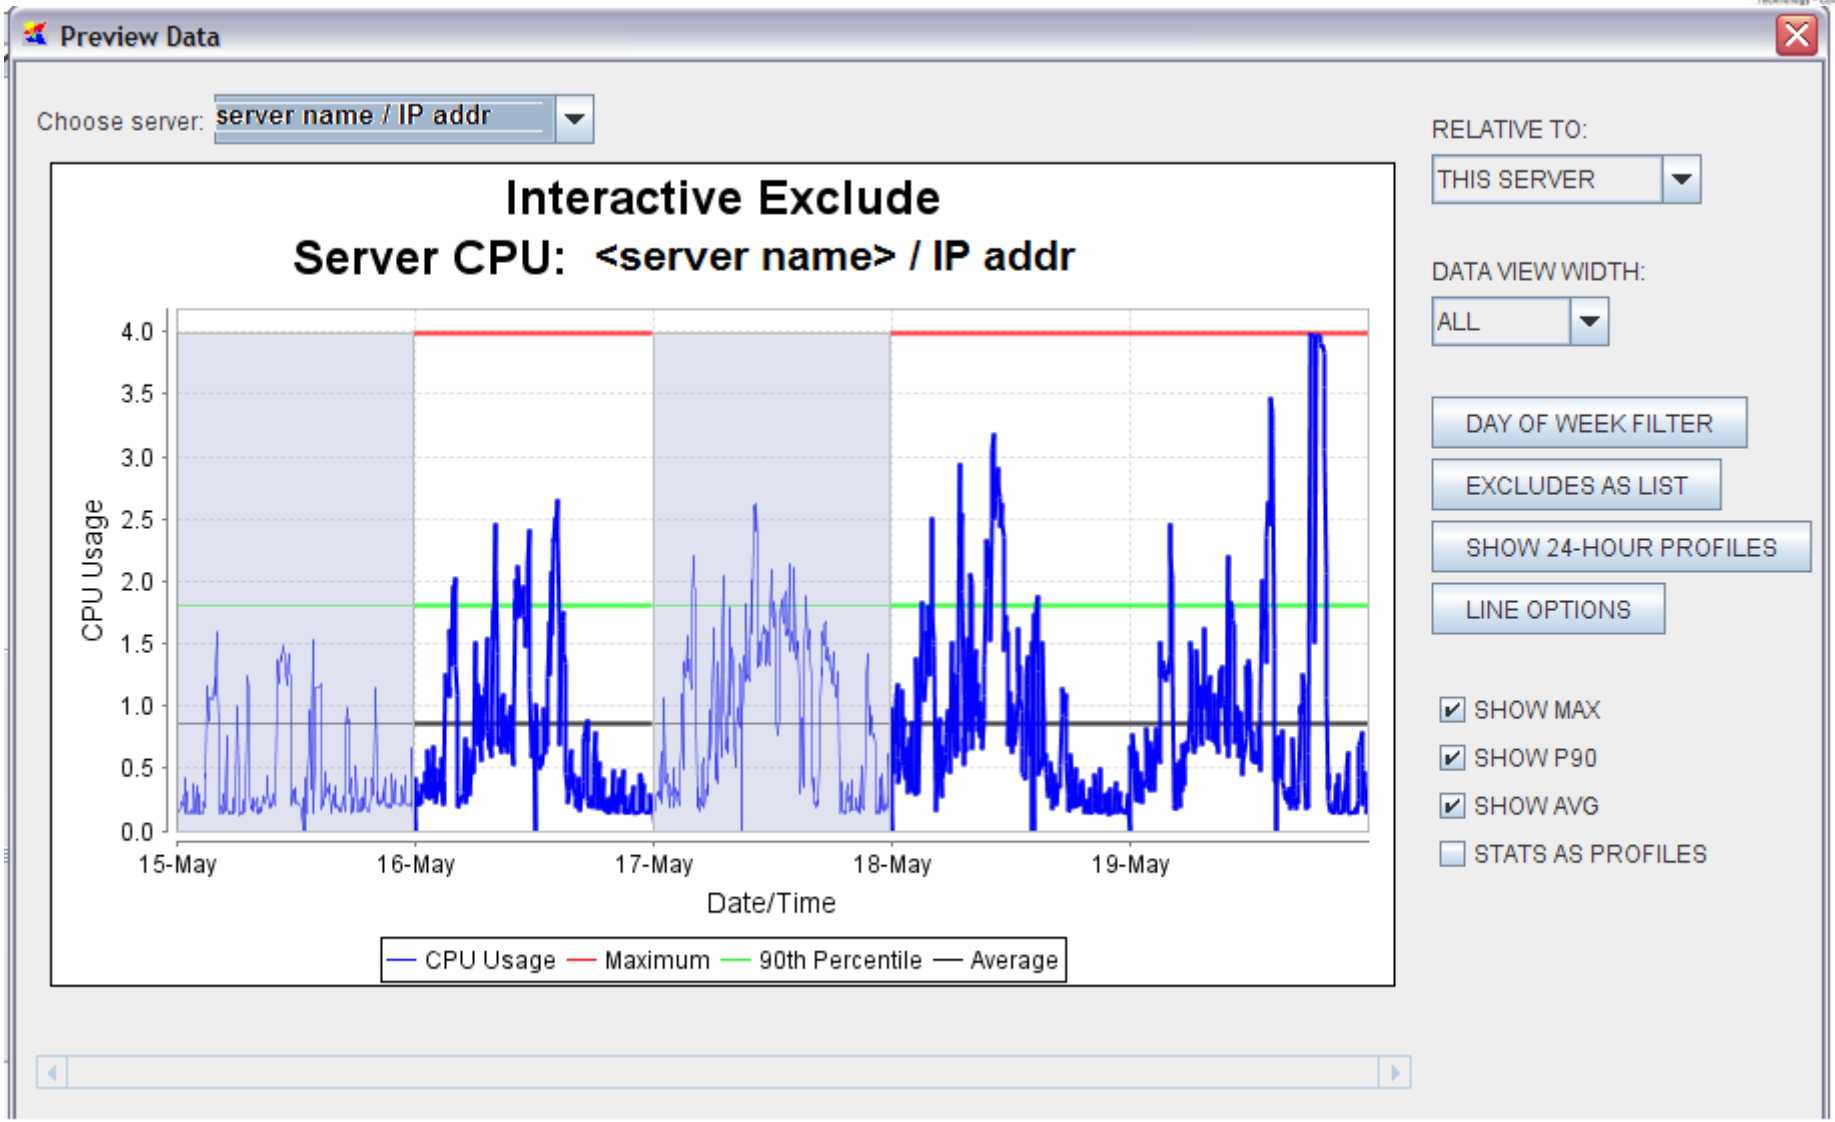

How to get started

EMAIL

zbladesz@us.ibm.com

inside IBM URLs

CPSTools : http://w3.ibm.com/support/americas/wsc/cpsproducts.html

zBladeSizer :
http://w3.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS4966

Techline Request Generator:
http://w3-03.ibm.com/support/techline/na/trg/form.html

- Design Assessment (***not*** System Design and Configuration)
- System z capacity (***not*** HW: System z)
- Solution area is n/a

zBladeSizer Addendum on Hybrid Optimization Algorithm

- Constrained optimization problem, a search through solution space
- Possible solution set and feasible solution set, unconstrained search
- Group segregation to make parallel
- Group solution is a server to blade assignment
- Fitness score is the number of blades for a group, tie breaker std dev of CPU
- Genetic algorithm (GA) vs. exhaustive search for small groups
- The GA population as a parallel solver
  - Ranking by fitness, cull by score
  - Breeding and selection, cross over, mutation
- 24 hour profile and the six resources, supply [* constraints] and demand [* growth rates]
- Configuration build is a combination of groups and the application of additional constraints
- Multiple independent runs to improve solution space coverage