Advanced Technical Skills (ATS) North America

Getting Started in (z/OS) Capacity Planning (Topics in Capacity Planning) Part II

SHARE
Technology · Connections · Results

March 2011

Ray Wicks
561-236-5846
RayWicks@us.ibm.com
CPS Tools Group, Gaithersburg. MD
CPSTools@us.ibm.com

Advanced Technical Skills
TECHNICAL SALES NORTH AMERICA

Advanced Technical Skills (ATS) North America

Bibliography

Ray has spent most of his career at IBM in the performance analysis and capacity planning end of the business in Poughkeepsie, London, and now at the Washington Systems Center. He is the major contributor to IBM's internal PA & CP tool zCP3000. This tool is used extensively by the IBM services and technical support staff world wide to analyze existing zSeries configurations (Processor, storage, and I/O) and make projections for capacity expectations.

Ray has given classes and lectures worldwide. He was a visiting scholar at the University of Maryland where he taught part time at the Honors College.

He won the prestigious Computer Measurement Group's A.A. Michelson award in 2000. His recent virtual sessions "Getting Started in Performance Analysis & Capacity Planning" workshop held for attendees in China and India was well accepted.

© 2010 IBM Corporation

Trade Marks, Copyrights & Stuff

Many terms are trademarks of different companies and are owned by them.

- **On foils that appear in this presentation are not in the handout. This is to prevent you from looking ahead and spoiling my jokes and surprises.**

Abstract

This tutorial is a **two** part introductory level session designed to introduce the student to the concepts required for Performance Analysis and Capacity Planning.

Emphasis is placed on large processor systems and examples will be largely drawn from z/OS but the concepts apply to all operating systems and hardware. The tutorial is organized to review the architecture where appropriate (albeit briefly). Topics:

- Conceptual and Perceptual structures for performance analysis and capacity planning,
- Using the Forced Flow law in PA & CP
- Performance Analysis queries for capacity planning,
- Processor performance data (ITRRs & MIPS),
- Resource Metrics for use in the Balance System model,
- Sample selection,
- Data preparation in z/OS,
- Using the utilization growth process in capacity planning,

Advanced Technical Skills (ATS) North America

zPCR

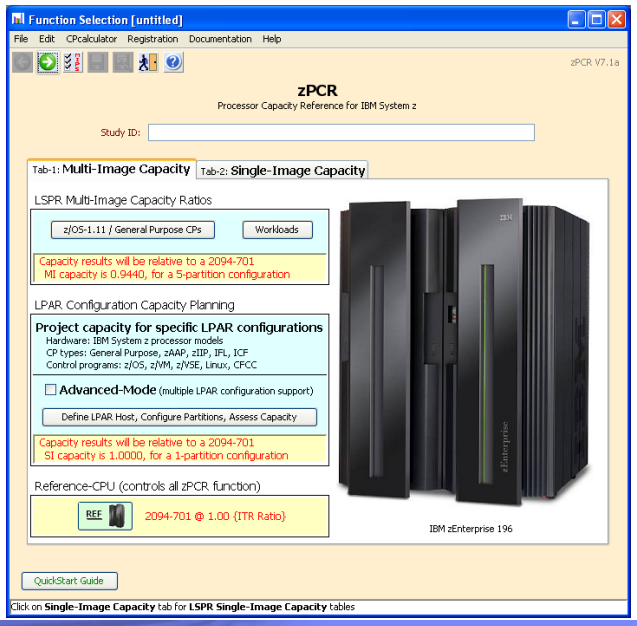

Advanced Technical Skills (ATS) North America

zPCR

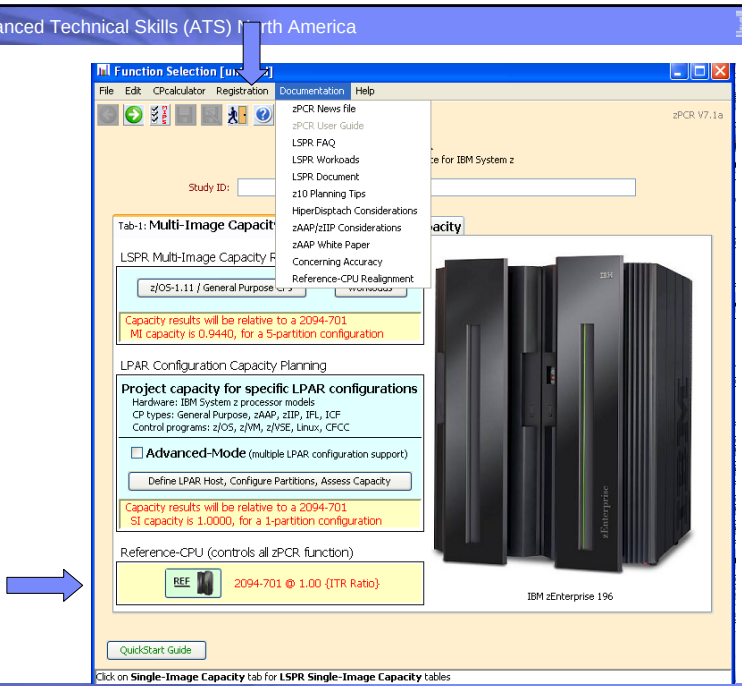

© 2010 IBM Corporation

Establish Power value in MIPS?

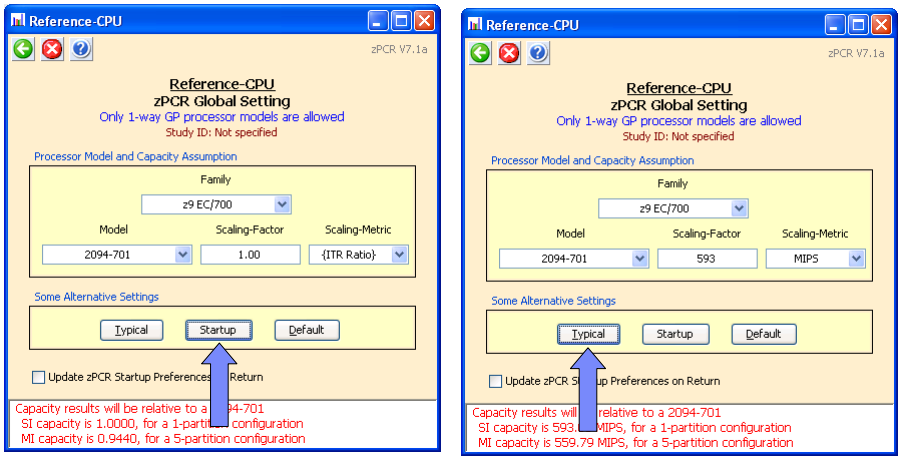

Processor Power: zPCR

Configuration Input

- **Manual**
- **RMF Listing**
- **EDF file (CP3KEXTR)**

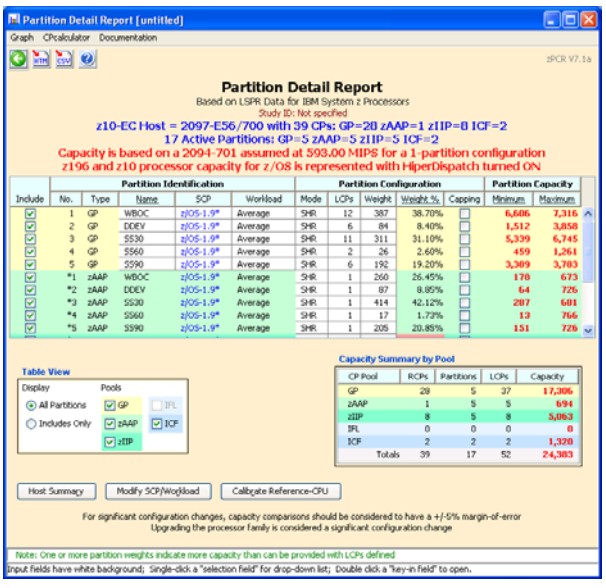

Advanced Technical Skills (ATS) North America

Processor Power: zPCR

Partition Configuration				Partition Capacity	
LCPs	Weight	Weight %	Capping	Minimum	Maximum
12	387	38.70%		6,606	7,316
6	84	8.40%		1,512	3,858
11	311	31.10%		5,339	6,745
2	26	2.60%		459	1,261
6	192	19.20%		3,389	3,783
1	260	26.45%		178	673
1	87	8.85%		64	726
1	414	42.12%		287	681
1	17	1.73%		13	766
1	205	20.85%		151	726

Capacity Summary by Pool

CP Pool	RCPs	Partitions	LCPs	Capacity
GP	28	5	37	17,306
zAAP	1	5	5	694
zIIP	8	5	8	5,063
IFL	0	0	0	0
ICF	2	2	2	1,320
Totals	39	17	52	24,383

Maximum: MIPS available if other partitions are idle given logical configuration.

Minimum: MIPS entitled to if other partitions are demanding their fair share (Weight) given the logical configuration.

2097-E56 summary with this logical configuration

Advanced Technical Skills (ATS) North America

ITRRs Reality? A Fuzzy Patch?

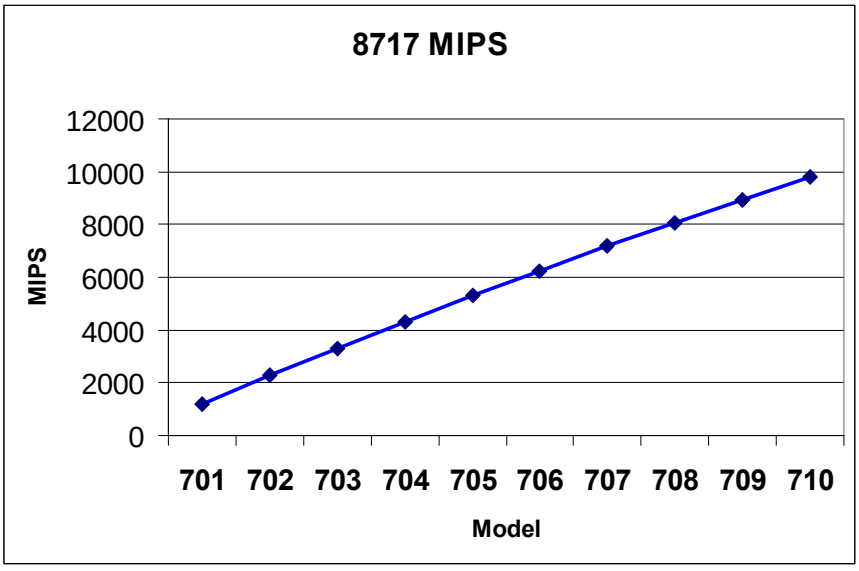

ITRRs Reality? A Fuzzy Patch?

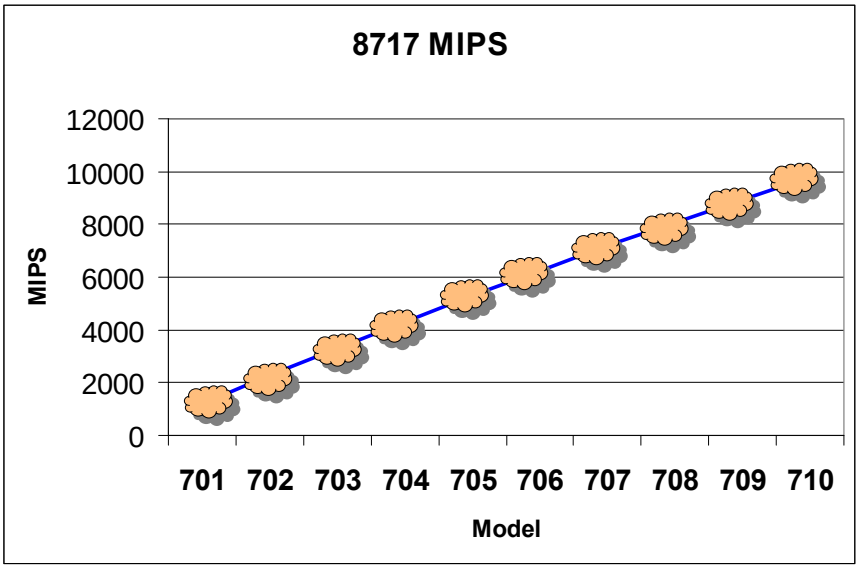

Comparing Fuzzy Numbers

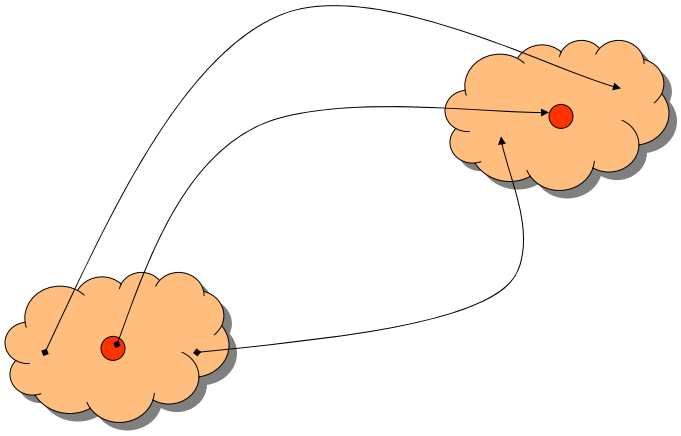

Summary: In a CPC Migration

Some Fiction, Some Truth

- New_CPU% = Old_CPU% x Old_Power/New_Power
  - Migrating to a CPC twice as fast: would 50% old = 25% new?
- The Workload stays the same.
- The Workload changes
  - The number of transactions changes
  - Response time changes impact user behavior
  - Latent demand can rear its ugly head
  - Different workloads are impacted differently
- The Software Changes
- The Memory Changes
- The I/O Changes

I/O Intensity

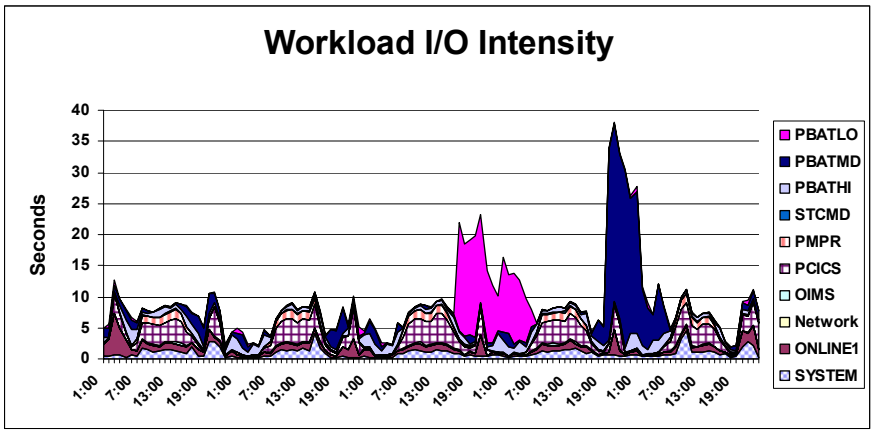

Intensity = Rate * Response_Time (Note Little's Law N = λ*T)

Traffic = Rate * Service_Time

Who's using the I/O? And When?

Advanced Technical Skills (ATS) North America

Workload Contention Analysis
(Latent Demand?)

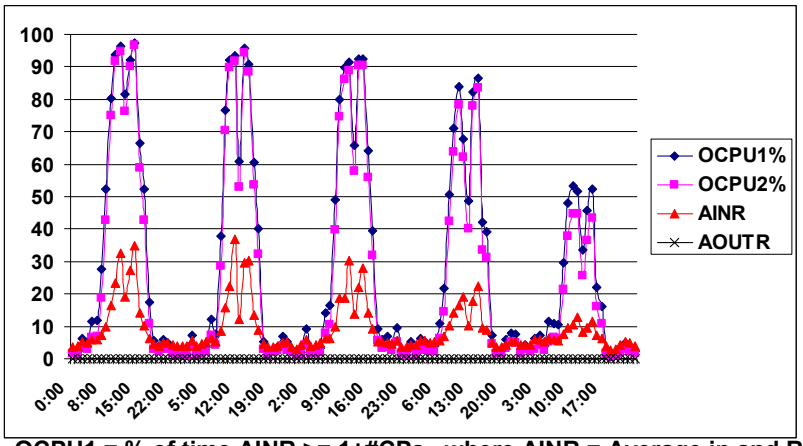

OCPU1 = % of time AINR >= 1+#CPs where AINR = Average in and Ready

Is someone not getting it? Who?

Advanced Technical Skills (ATS) North America

CPU Delay By Workload

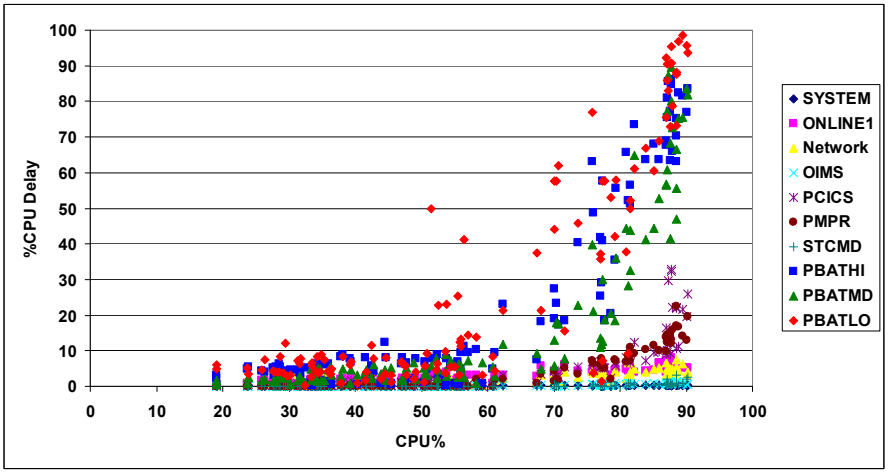

Ah the low priority is not getting it. So?
How about Single Task Multi thread?

CPU Delay = F(CPU%)

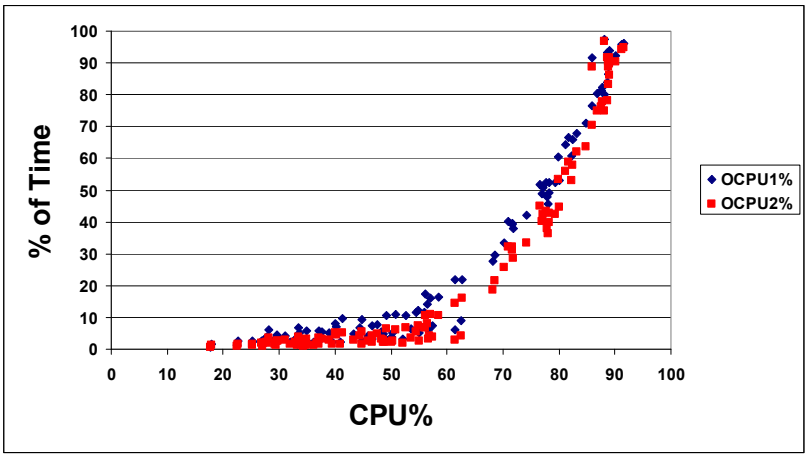

Does this look like a Queuing theory graph? But who is being delayed? This is a nice graph because of the stable environment.

I/O Delay by Workload

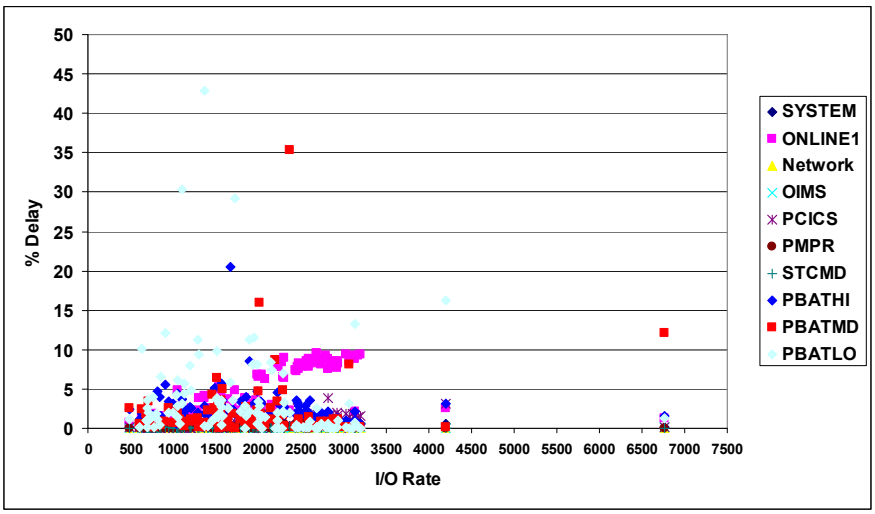

4 Typical Tasks

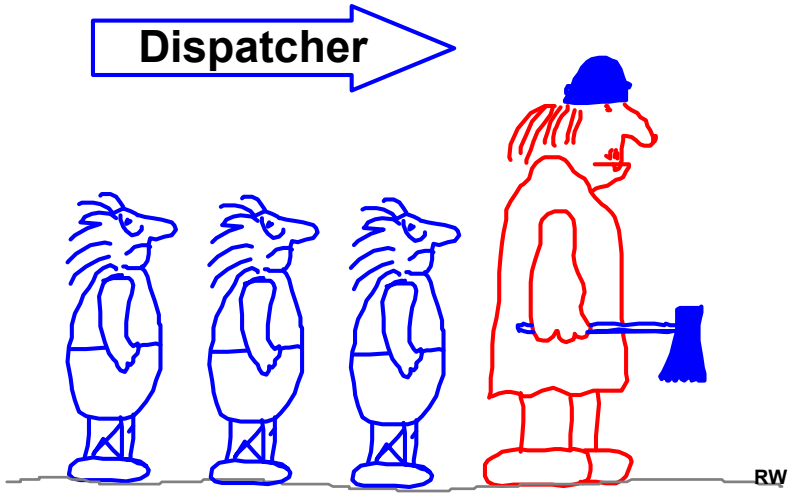

Guess who gets all they want? How much do we, the little people, get if there's one server? Two servers? Four servers? What's best?

4 Typical Tasks on the Ready Queue
(Latent Demand)

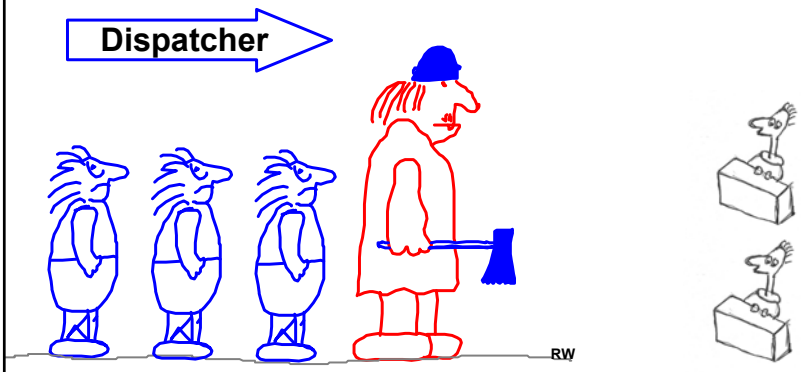

- **What happens when you migrate from 2 slow servers to 1 fast server?**
- **Who complains?**
- **How can you anticipate the result?**

Exceptions

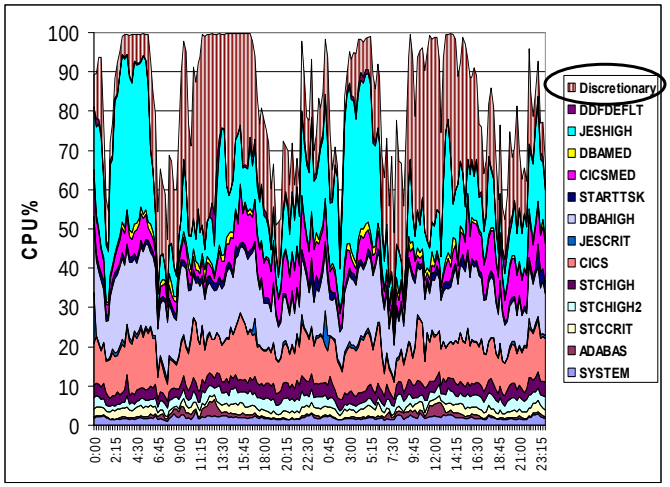

How busy is this system? Pick an interval to model.

Remove Discretionary?

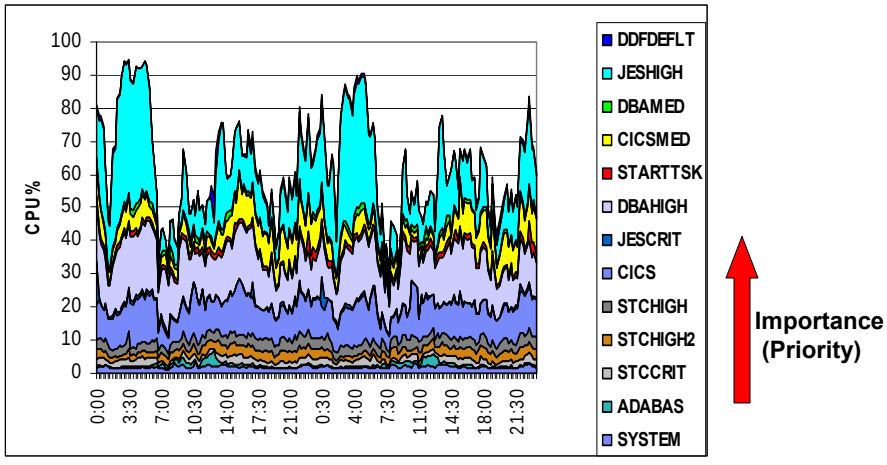

How busy is this system? If discretionary work is excluded, it peaks around 92%. Otherwise it peaks at 100%.

Workload Profile (NFU, BU)

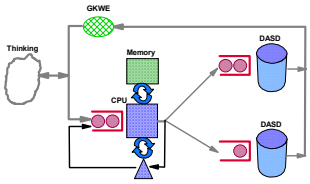

Resource Demands at Modeled Service Centers
- **Homogeneous?**
- **Heterogeneous?**
- **Grouped?**

SMF Data Groupings RMF (70s)
- **Performance Groups**
- **Service Classes**

SMF Data (30s, 42s, etc)
- **Job Name**
- **Program Name**
- **Dept. Name**
- **RACF ID**
- **....**

Resource Description
- **CPU Time**
- **I/O Rate**
- **I/O Devices, Data Sets, Rate, Resp**
- **Storage (MB)**
- **Transaction Rate**
- **Transaction Resp**
- **Sysplex**

Growth Laws

Moore's Law describes an important trend in the history of computer hardware: that the number of transistors that can be inexpensively placed on an Integrated circuit is increasing exponentially, doubling approximately every two years.

Gates' Law says that the speed of software halves every 18 months.
- **Maintenance adds path length → MIPS per Transaction goes up.**
- **User transactions increase in complexity → MIPS per Transaction go up.**

Growth Laws

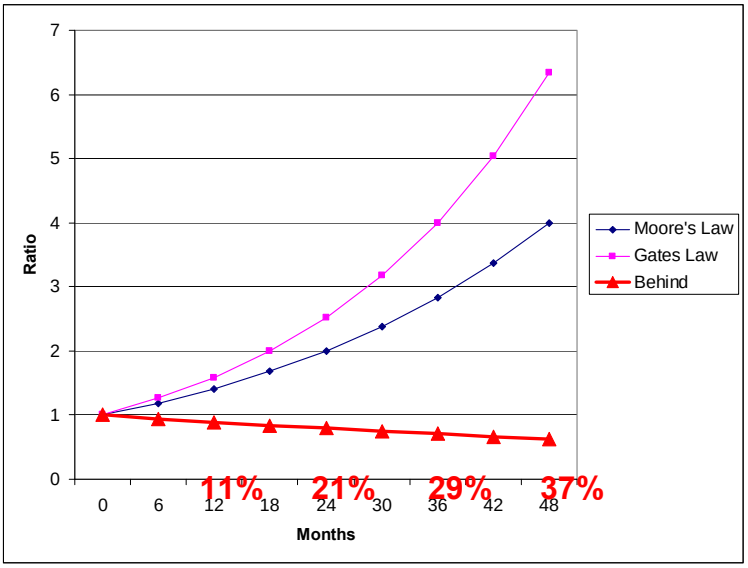

Capacity Planning

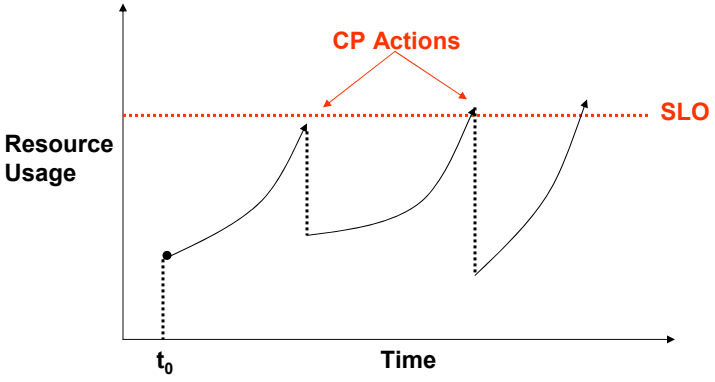

Starting from a single point at t_0, we project growth until some threshold is reached – a Service level Objective or Agreement (SLO, SLA).

Then we take action.

SLO or SLA for a Workload Group

- **Interactive**
  - **For some number of users**
  - **At some threshold transaction rate**
  - **At some threshold Response Time**
  - **At some amount of power & I/O**
- **Batch**
  - **For some number of Jobs**
  - **At some rate**
  - **At some about of power & I/O**
  - **At some turn-around threshold**

Capacity Planning Actions

- **Upgrade Hardware**
  - **Add CPs (PUs)**
  - **New Model**
  - **Add Another CPC**
- **Move Workload to another image**
- **Split Workload and move a piece**
- **Tune it?**
- **Continue to Suffer**

How Accurate Is It?

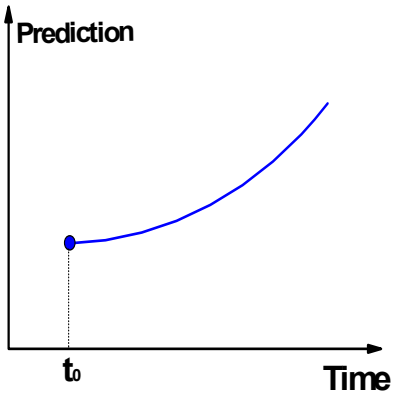

Starting from an initial point of maybe dubious accuracy, we apply a growth rate (also dubious) and then recommend actions costing lots of money.

Accuracy

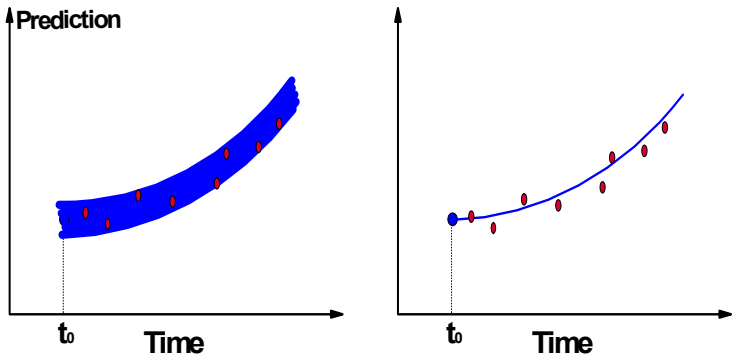

Accuracy is found in values that are close to the expected curve. This closeness implies an expected bound or variation in reality. So a thicker line makes sense.

Fuzzy Patches

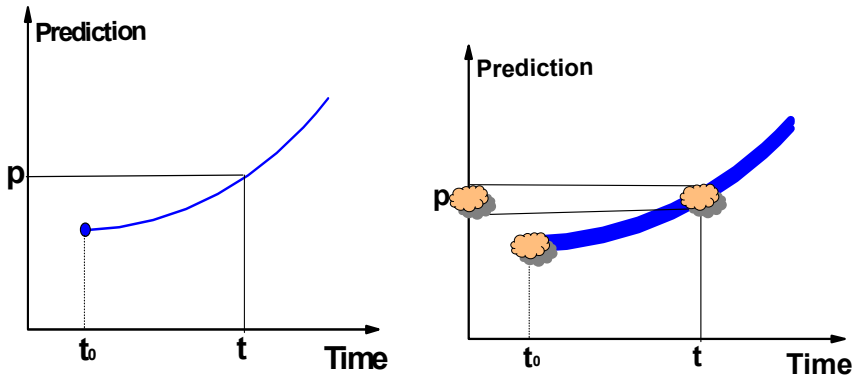

At time t, is the prediction a precise point p or a fuzzy patch?

Fuzzy Factors

- **Basis for prediction is a single sample taken from a set of samples with some distribution.**
- **Growth Factor applied may be just better than fiction.**
- **Prediction compounds the fuzz and is itself fuzzy.**
- **Niels Bohr: "Prediction is very hard to do. Especially about the future."**

Modeling Interval Selection

- Overall Usage
  - By pool: GCPs, zIIPs, zAAPs, IFLs, ICFs
  - Percentile selection: 90th, 95th, or Peak? What does your SLA say?
- Filter by
  - Date (avoid weekends and holidays)
  - Time (shift or business period)
  - Target CPC(s)
    - Target partitions
    - Target workloads (by importance)

Interval Selection Step 1

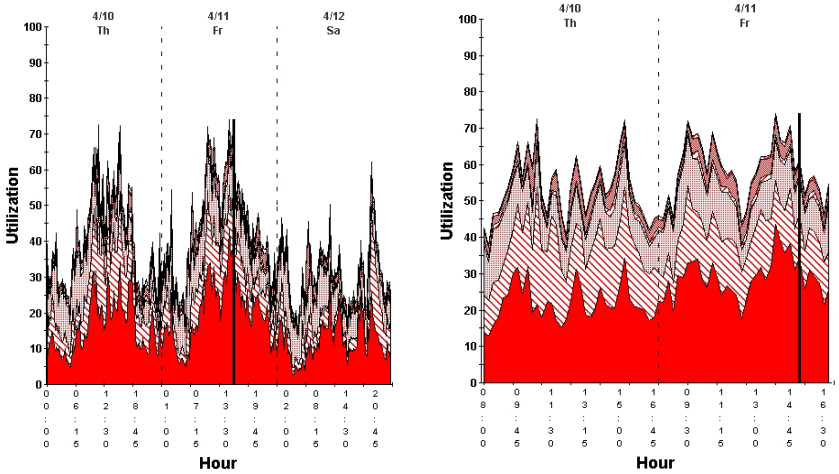

Eliminate unwanted days (weekend) and hours (just prime shift)

Advanced Technical Skills (ATS) North America

Interval Selection Step 2

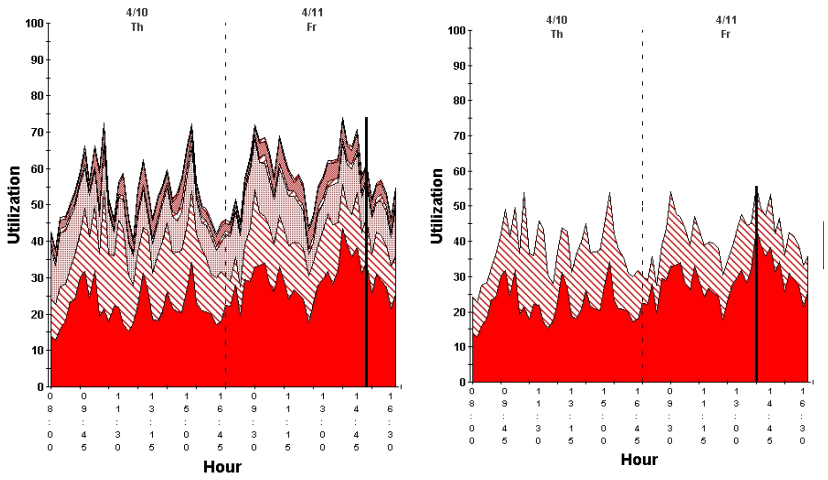

Select interval for business partitions. Peak? 90th percentile?

Advanced Technical Skills (ATS) North America

GCPs by Sample Interval

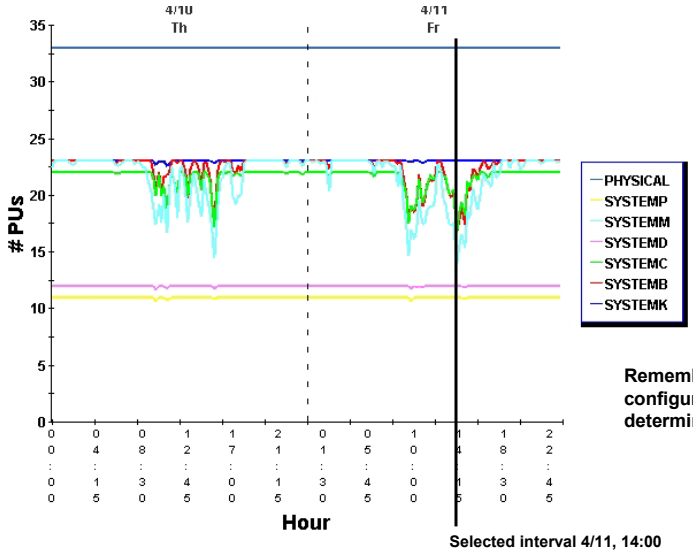

Remember: Logical configuration determines MIPS.

Selected interval 4/11, 14:00

Advanced Technical Skills (ATS) North America

CPC – GCP Pool

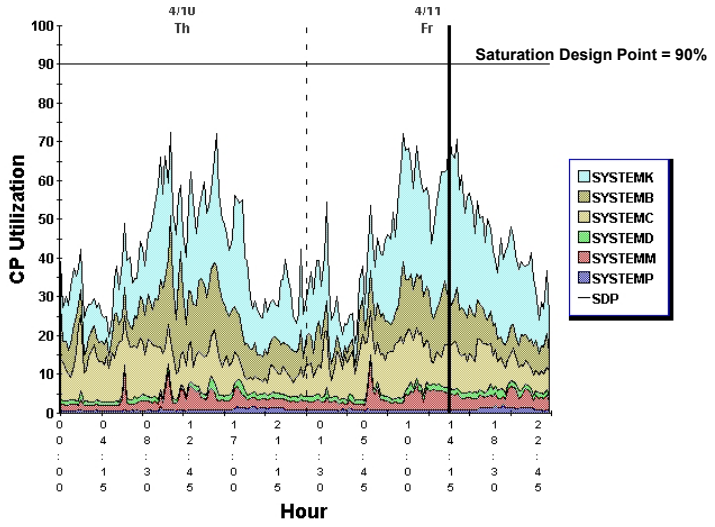

Advanced Technical Skills (ATS) North America

zAAP Pool

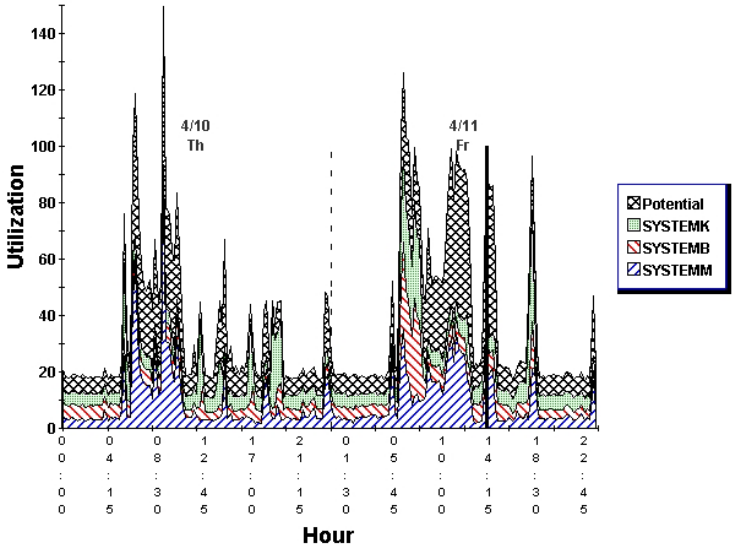

Just Growth Utilization Analysis at 40% GCP Pool

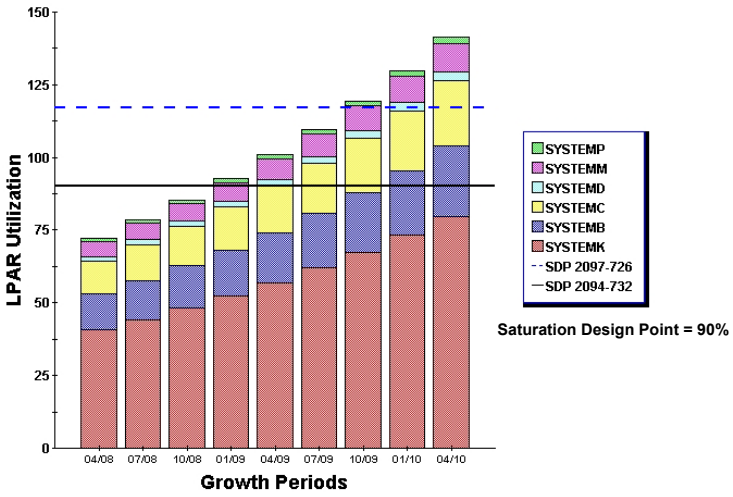

Saturation Design Point = 90%

zAAP Pool at 40%

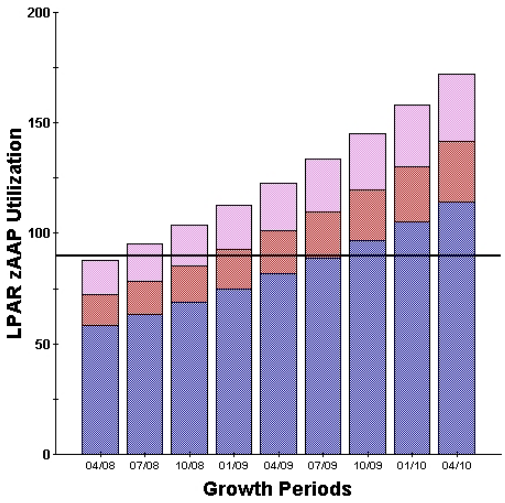

Potential zAAP MIPS (GCP MIPS) Added to zAAP MIPS.

Logical GCP Workload View

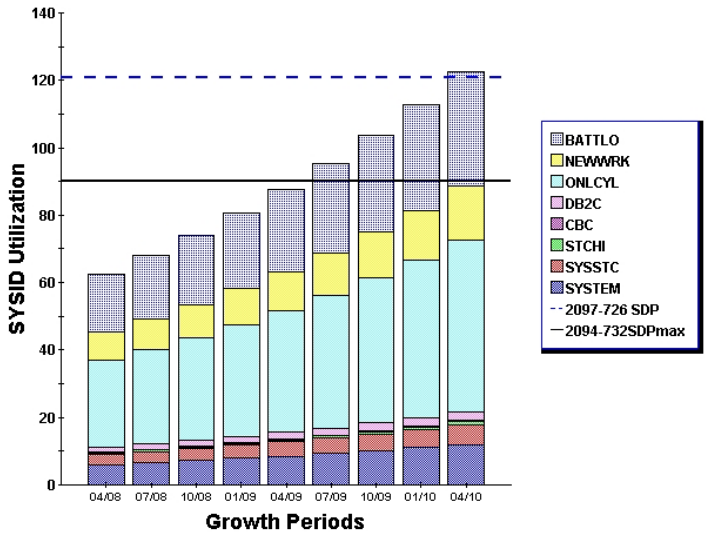

Adjusted Workload Data

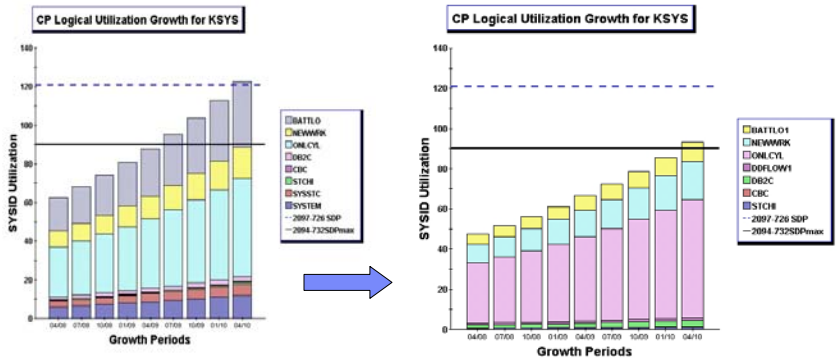

Remove SYSTEM workloads, adjust capture ratios.
Delete BATLO and add back 5% for BATLOW1.

Logical Projection

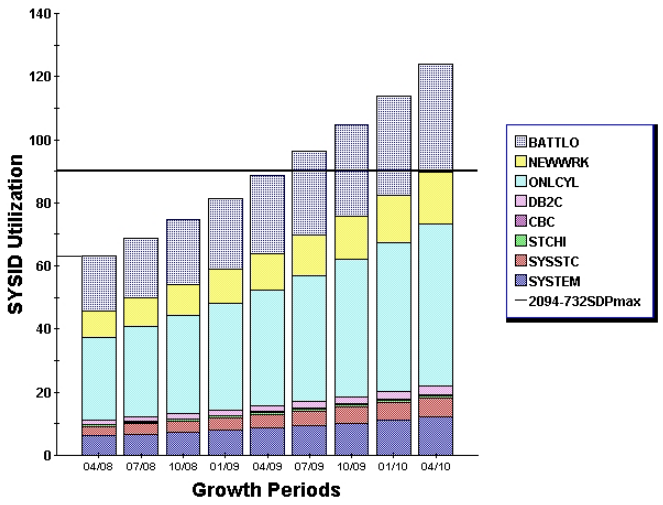

Physical Projection

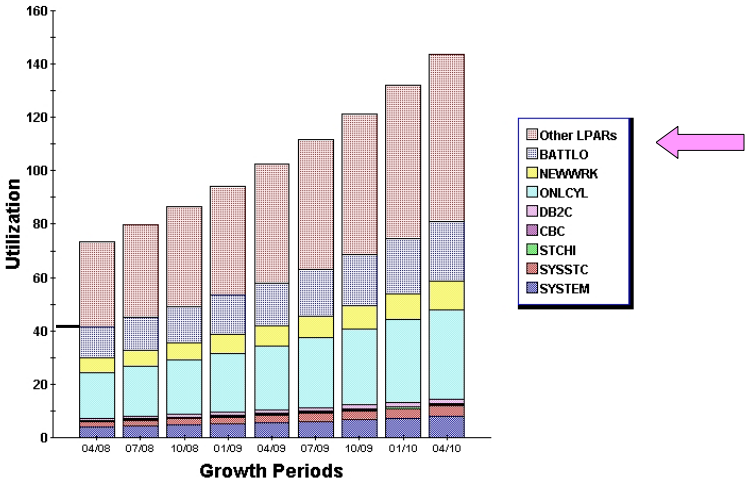

Physical CPU% * Growth Factor. Include other partitions.

What Really Happens

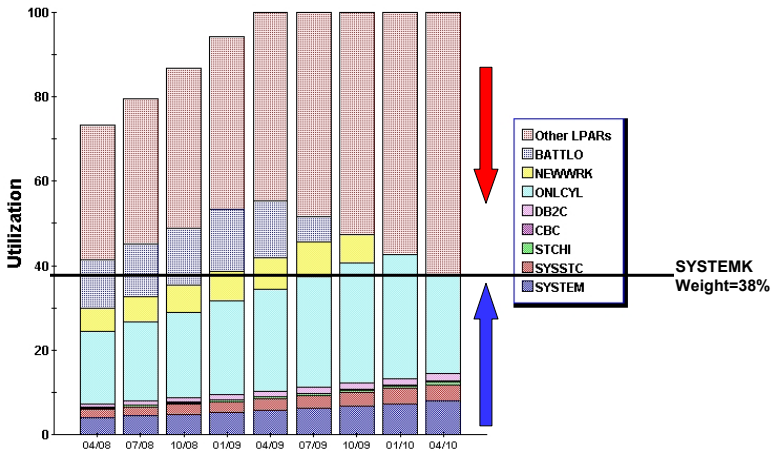

After CPC reaches 100%, CPU% is distributed by weight Within SYSTEMK, the CPU% is distributed by priority (importance). Lower priority work gets squeezed.

Storage Projection

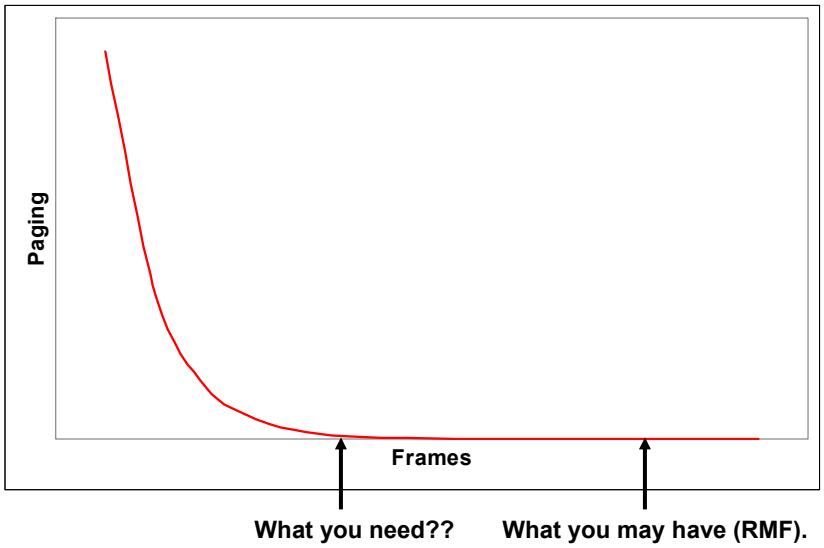

Storage = F(CPU%)?

Who is Using Storage?

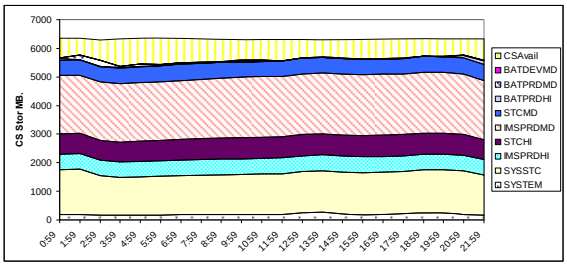

Who is Using CPU?

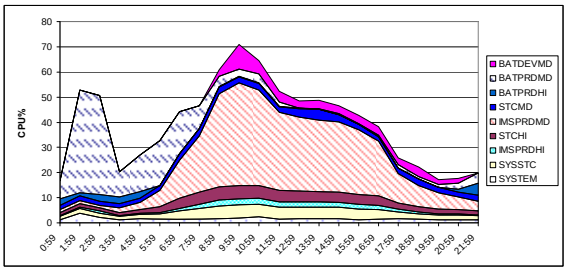

Expectations for Batch

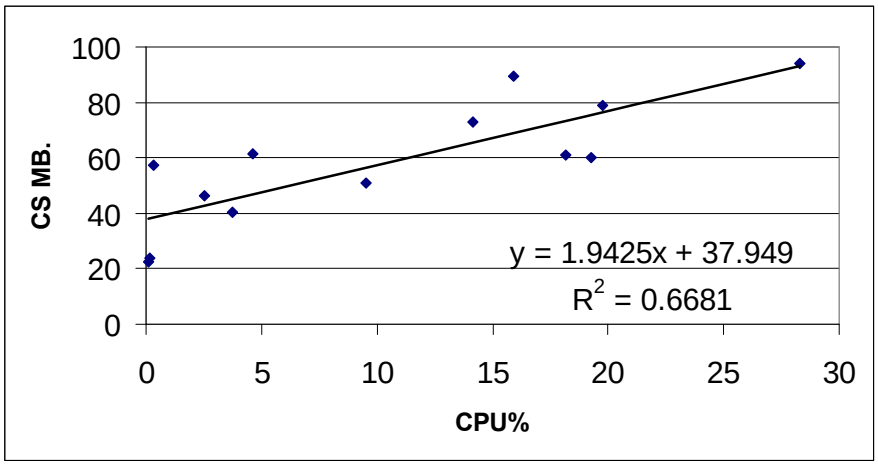

Expectations for DB/DC

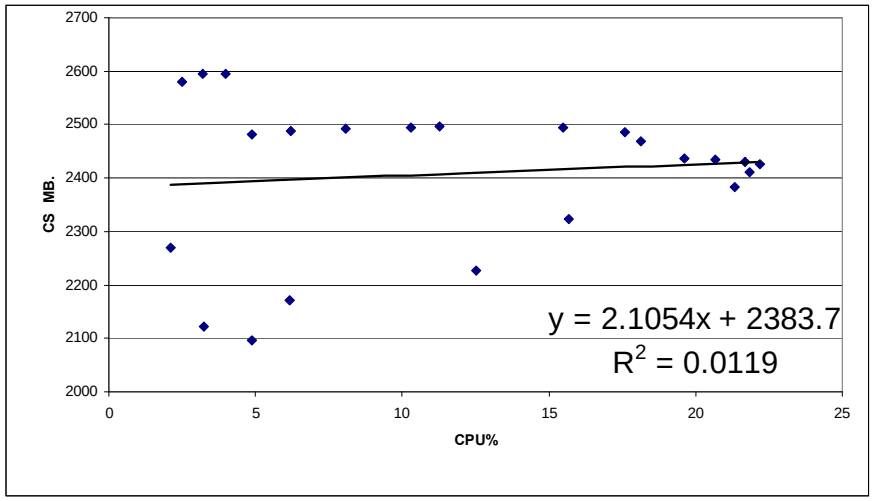

DB/DC Scaled to 0

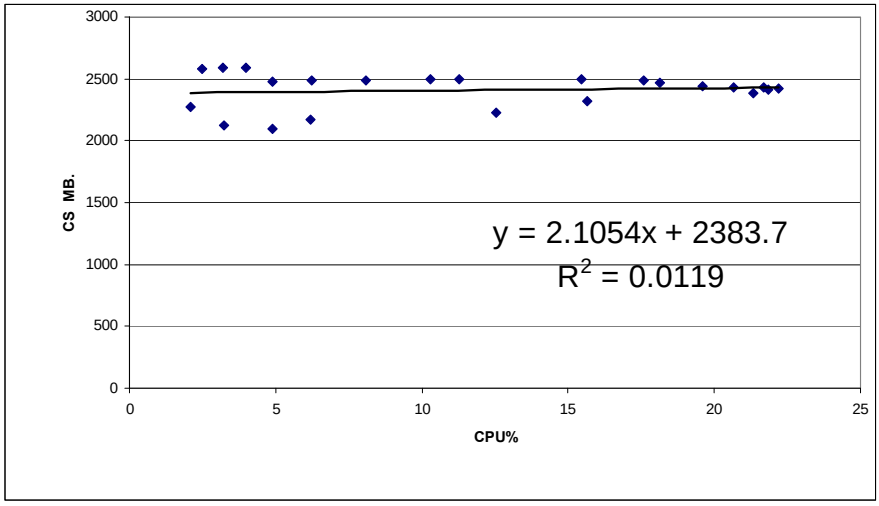

Note: R^2 doesn't change whereas perception does.

Storage Growth

- Size depends upon the number of address spaces
- If growth ⇗ More Address Spaces ⇗ more storage

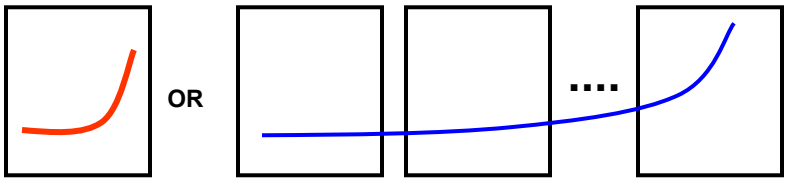

Some applications storage requirements grow as the load grows. Typically it's those applications where growth means more address spaces (TSO, batch). DB/DC often has workload growth without an increase in storage. Only after the number of address spaces increase does the storage change.

Local Storage Hierarchy

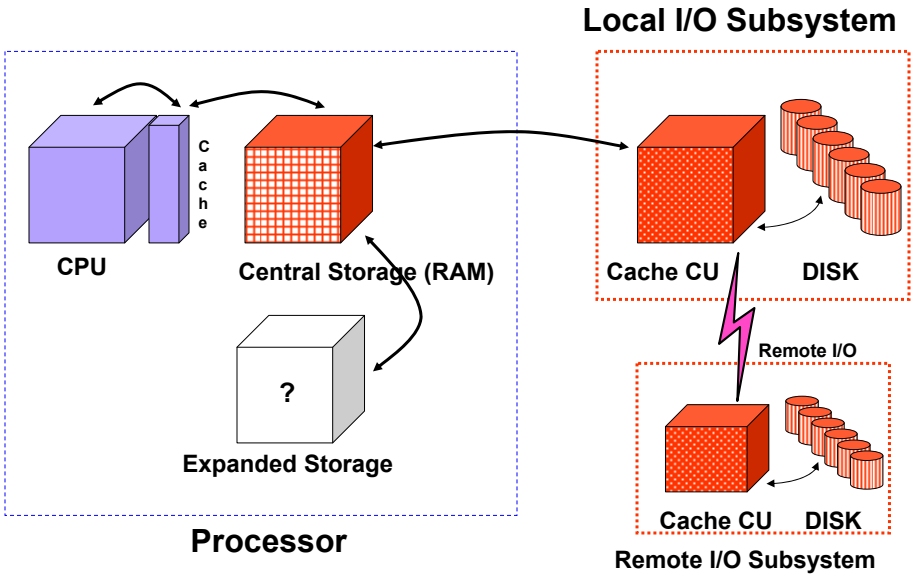

RMF Partition Report

CENTRAL STORAGE

	MIN	MAX	AVG
AVAILABLE	1450630	1917651	1638857
SQA	17,478	18,443	17,840
LPA	9,842	9,842	9,842
CSA	43,884	43,946	43,913
LSQA	178,761	180,850	179,946
REGIONS+SWA	4379708	4845902	4658493
TOTAL FRAMES	6553600	6553600	6553600

FIXED FRAMES

NUCLEUS	2,629	2,629	2,629
SQA	15,528	16,493	15,890
LPA	90	90	90
CSA	12,187	12,187	12,187
LSQA	65,380	67,270	66,451
REGIONS+SWA	60,380	73,810	61,791
BELOW 16 MEG	76	98	78
BETWEEN 16M-2G	38,193	39,699	39,097
TOTAL FRAMES	157,643	170,842	159,039

Sample Partitions

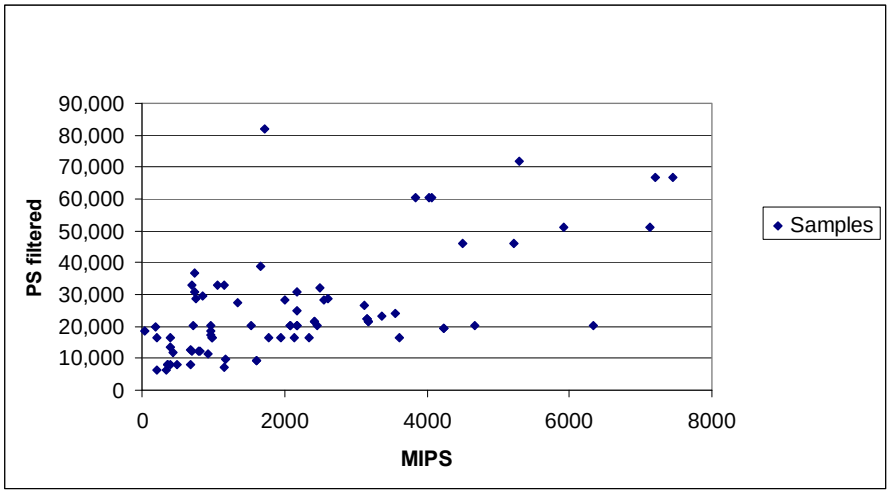

Advanced Technical Skills (ATS) North America

Projection: Approach 1

Use and exponential function CS=4000 + 0.04*(MIPS_Used)^D

Known **Known**

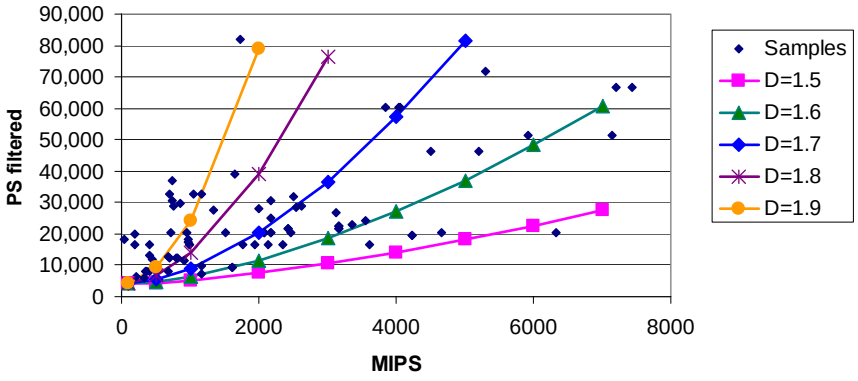

Advanced Technical Skills (ATS) North America

Projection Approach 1

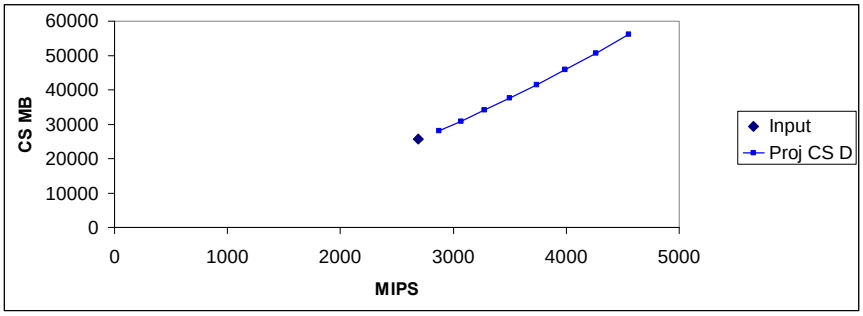

CS=25600, MIPS=2689
IF 25600 = 4000 + 0.04*(2689)^D
Then D=LOG(((25600-4000)/0.04),10)/LOG(2689,10)
D= 1.671451
Apply 30% per annum growth to MIPS
Use formula CS=4000+0.04*(MIPS)^1.671451 for future MIPS to project CS.

Projection Approach 2

Use ratio CS/MIPS
CS=25600, MIPS=2689; CS/MIPS = 9.52
CS/MIPS metric is [10%, 50%, 90%] = [6.77, 14.06, 37.31]
CS = 9.52 * MIPS

Result

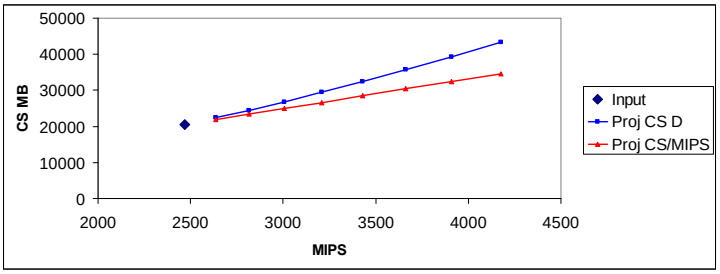

	Quarter	CS	MIPS	Ratio	Proj CS D	Ratio	Proj CS/MIPS	Ratio	D	CS/MIPS
Input	0	25600	2689		25600		25600		1.671451	9.520268
	1		2872	1.068	28111	1.098	27341	1.068		
	2		3067	1.068	30913	1.100	29200	1.068		
	3		3276	1.068	34041	1.101	31186	1.068		
	4		3498	1.068	37533	1.103	33306	1.068		
	5		3736	1.068	41431	1.104	35571	1.068		
	6		3990	1.068	45782	1.105	37990	1.068		
	7		4262	1.068	50638	1.106	40573	1.068		
	8		4552	1.068	56059	1.107	43332	1.068		

Note 6.8% quarterly growth is 30% per annum.

CP Alternatives

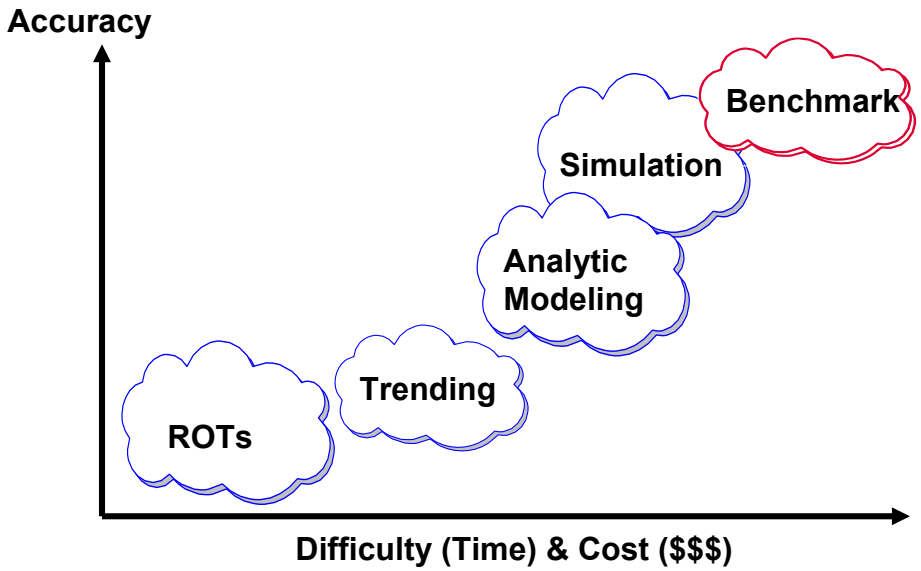

Rules of Thumb

- ROTs (thresholds) are often quite adequate for CP

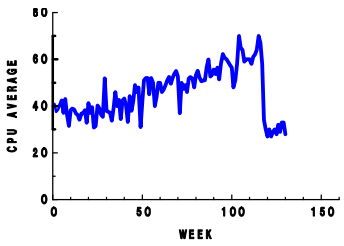

- Useful for Health Check

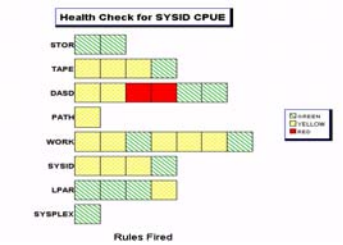

Rules of Thumb

- **Honor your Father & Mother**
- **Do unto others as you would have them do unto you.**
- **Do unto others before they do unto you.**
- **Keep your CPU%<90%**
- **Don't swim soon after eating.**
- **It is better to give than receive.**

Philosophical Remark

$$\forall x \Phi x \quad .\equiv. \quad \sim \exists y \sim \Phi y$$

All swans are white ≡ There does not exist a swan which is not white

We understand a Rule by trying to break it.
Or
Learn the rules so you know how to break them correctly.

Balanced System ROTs

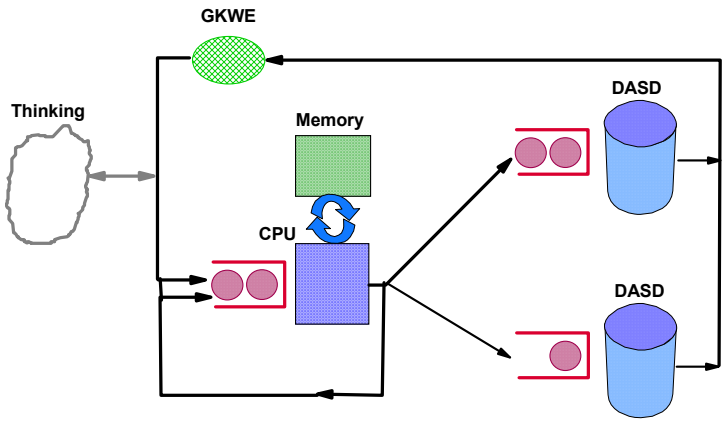

MIPS used, memory used, I/O used should be in some proportion.

Expected Bounds

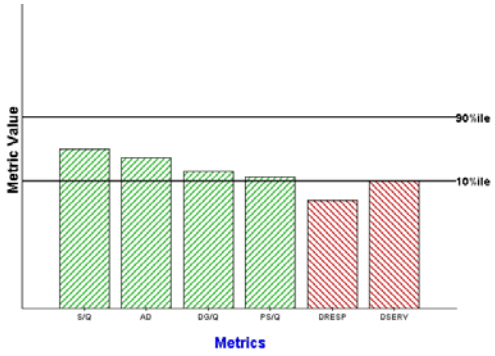

The resource ratio is shown as a bar. If the bar is above the 90%ile line, it means that the value was in the top 10% of the samples reviewed. Similarly, if the bar is below the 10%ile line, the value is in the bottom 10%. Neither is good or bad., it's an flag to examine the amount of resource available.

Trending

- ❑Trending predicts the future if the future looks like the past.
- ❑Time Series Trending is complicated.
- ❑Trending can answers overall CP questions.

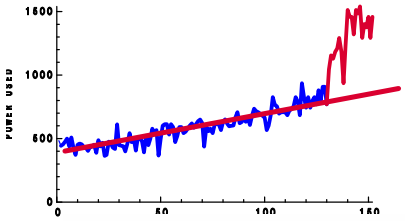

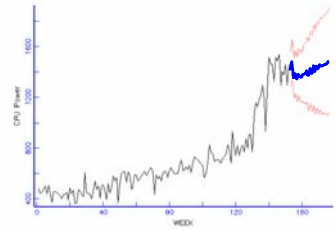

Analytic Modeling

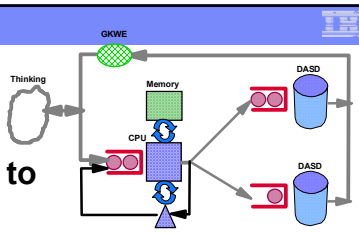

- ❑ Pre-built packages can be fast to solve and relatively easy to use.
- ❑ Flow is statistically driven and usually predefined.
- ❑ Accuracy?
  - ❑ Utilization within 5%
  - ❑ Response times within 30%
- ❑ Data acquisition is key.
- ❑ Calibration can be tough.
- ❑ Custom analytic models are really tough.
- ❑ Requires technical staff.
- ❑ Services are Available.

ROTs & Analytic Modeling

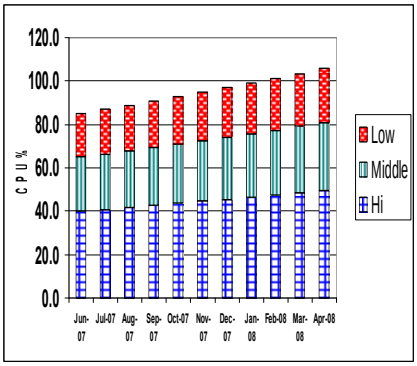

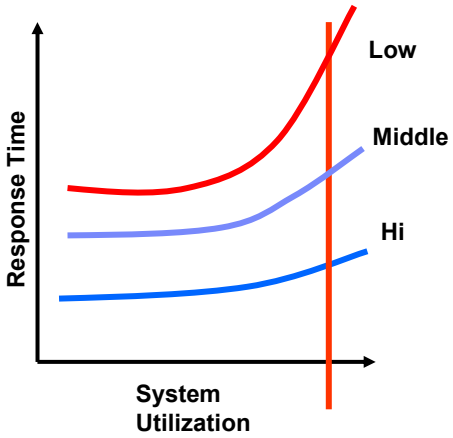

The relationship between Utilization and Server Response is sensitive to the priority of the workload. Utilization in Response time is "perceived utilization". Watch out for: Logical vs physical utilization and single task workloads.

Simulation

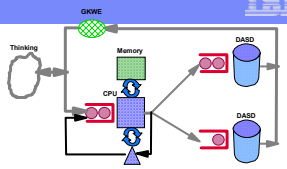

- ❑ **Pre-built packages are slower to solve and can be relatively easy to use.**
- ❑ **Flow is statistically driven and usually predefined but can be customized. (Application modeling.)**
- ❑ **Accuracy?**
  - ❑ **Utilization within 5%**
  - ❑ **Response times within 30%**
- ❑ **Data acquisition is key.**
- ❑ **Calibration can be tough.**
- ❑ **Custom models are build from service center building blocks.**
- ❑ **Simulation languages do exist.**
- ❑ **Specialized staff.**
- ❑ **Services exist.**

Benchmark

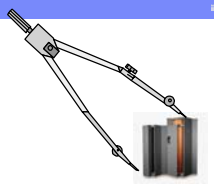

- A lot of work in preparation
  - Hardware/Software
  - Workload
  - Lot's of time.
- It does mimic the running environment the best.
- Software flow & queuing
- Software usage
- It's expensive.
- Variations limited by resources.
- Given the resources the benchmark can be complicated.
- **Tests the environment** - does it work?

CP Alternatives

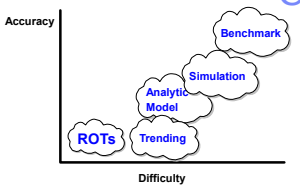

- What questions do you have?
- What questions must you answer?
- Cost of the answer?
- Cost of getting it wrong?
- Time line?
- What happens if you get it wrong?

Things to Remember

- **Be aware of the exceptions to the rule.**
- **A framework helps but it can make you see things that just aren't there.**
- **Impeccable mathematics does not replace knowledge of the facts.**
- **Protect yourself.**
- **Business decisions can override technical issues.**
- **Sometimes being understood is more important than being very accurate.**
- **Being "very" accurate may be a luxury of the idle.**
- **Other than the technicalities, there may be a hidden agenda.**

Bibliography - I

The Art of Computer Systems Performance Analysis, by Raj Jain, Wiley. I like this one. It is thorough and complete. A very good reference.
Capacity Planning for Web Performance, by Daniel A. Menasce and Virgilio A.F. Almeida, Prentice Hall. A good book on network structure and terminology and introduction to the topic.
Probability, Statistics, and Queuing Theory, by Arnold O. Allen, Academic Press Inc. This is the classic in queuing theory.
Performance by Design: computer capacity planning by example. By Daniel A. Menascé, Virgilio A. F. Almeida, and L. W. Dowdy. The web site http://cs.gmu.edu/~menasce/perfbyd/ has a lot of .xls modeling worksheets.
MVS I/O Subsystems, by Gilbert E. Houtekamer and H. Pat Artis, Performance Associates. More than you want to know about the I/O subsystem. A definitive source but is a little out of date. Is available from Intellimagic or perfassoc.com.
Exploring IBM S/390 Computers, by Jim Hoskins and George Coleman, Maximum Press. A general introduction to S/390 hardware and architecture. (with IBM G326-3006-06)
Statistical Concepts and Methods, by Gouri Bhattacharyya and Richard A. Johnson, John Wiley & Sons.
The Practical Performance Analyst, by Neil J. Gunther, Authors Choice Press. A very good book.
Almost any volume of the Computer Measurement Group (CMG) Proceedings is worth looking at for performance and capacity planning articles. Web Site: http://www.cmg.org/measureit/

Bibliography - II

GC28-1761 MVS™ Planning: Workload Management. A guide to WLM.
SC28-1950 Resource Measurement Facility Report Analysis. A guide to report reading.
SC28-1951 Resource Measurement Facility Performance Management Guide. A good tutorial to get started.
SG24-5975 IBM zSeries 900 Technical Guide. A good hardware architecture and implementation Red Book.
LY28-1042 RMF™ Support for LPAR Management Time. Want to know how LPAR works?
SC28-1187 Large Systems Performance Reference by John Fitch. John goes into detail about the LSPR data.
SG24-4356 System/390® MVS Parallel Sysplex Performance. A good Red Book on Parallel Sysplex RMF reports and data.
SG24-4680 System/390 MVS Parallel Sysplex Capacity Planning . A good Red Book on the function and capacity of Parallel Sysplex.

A great URL for z/Series documents in general:
http://www-1.ibm.com/servers/eserver/zseries/zos/bkserv/
RMF in particular:
http://www-1.ibm.com/servers/eserver/zseries/zos/bkserv/r4pdf/rmf.html
For zPCR,
search www.ibm.com for “zPCR” & “SoftCap”

Bibliography - III

EXCEL:

Applied Statistics For Engineers and Scientists Using Excel and MINITAB, by David Levine, Patricia Ramsey, Robert Smidt, Prentice Hall. This comes with a CD containing handy Excel Add-Ins.
Excel Data Analysis by Jinjier Simon, Wiley. Nice basic reference concentrating on data presentation.

Other Good Stuff:

The Black Swan: The Impact of the Highly Improbable, by Nassim Nicholas Taleb, Random House. This is an informative and entertaining approach to statistical analysis among other things.

Statistics as Principled Argument, Robert Abelson, Erlbaum Assoc. Publishers, 1995. Go good discussion of the use of statistics without the ugly formulae.

Judgment under uncertainty: Heuristics and biases, Kahneman, Slovic, & Tversky, Cambridge University Press. The first chapter alone is worth reading. It’s a summary of the pitfalls with intuitive thinking.

Ray Wicks' Monographs

CPS document for this presentation and other interesting monographs can be obtained from your favorite IBMer. Goto:
ftp://cpstools.washington.ibm.com/zcp3000/win
Look for: Getting Started In CP.
These have been published in cmg.org/measureit.