

Capacity Planning for z/OS

Marty Deitch mdeitch@us.ibm.com

- Overview & Methodology
- Demo of New IBM Tivoli Modeling Tool

Capacity Planning - Tips & Tools

- Definitions

 - Ensuring Service Levels (SLAs) are met

 - Timely and Proactive

- Drivers

 - Costs connected to available capacity

 - Despite shrinking hardware costs, software costs are increasing

- Difficulties:

 - Performance is the key metric (not utilization)

 - Performance is difficult to predict

Workload Sizing - the Easy Part

- The three Cs

 - Workload Characterization
 - Workload Correlation
 - Data Collection

- Workload Characterization

 - Based on computer metrics (i.e. Service/Reporting Classes)
 - Based on Key Applications
 - Prioritize by importance

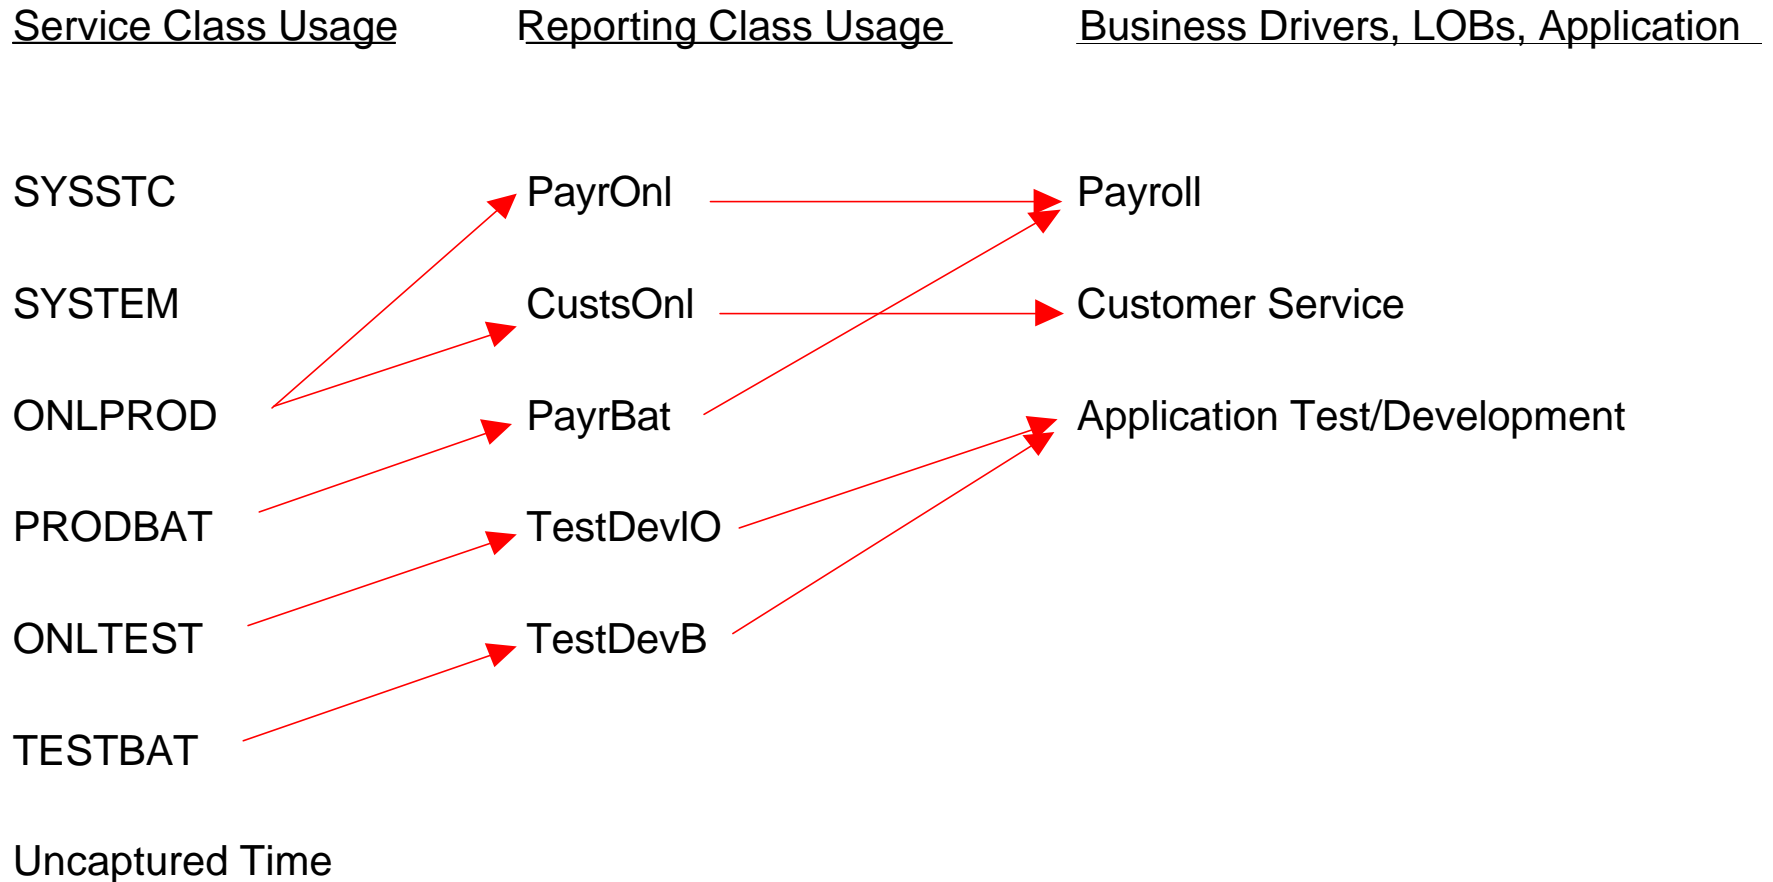
- Workload Correlation to business metrics

 - X amount of Widgets = Y CICS transactions/sec
 - X amount of customer calls/day = Y MIPS in Reporting Class Z

- Prediction based on past history plus planned changes

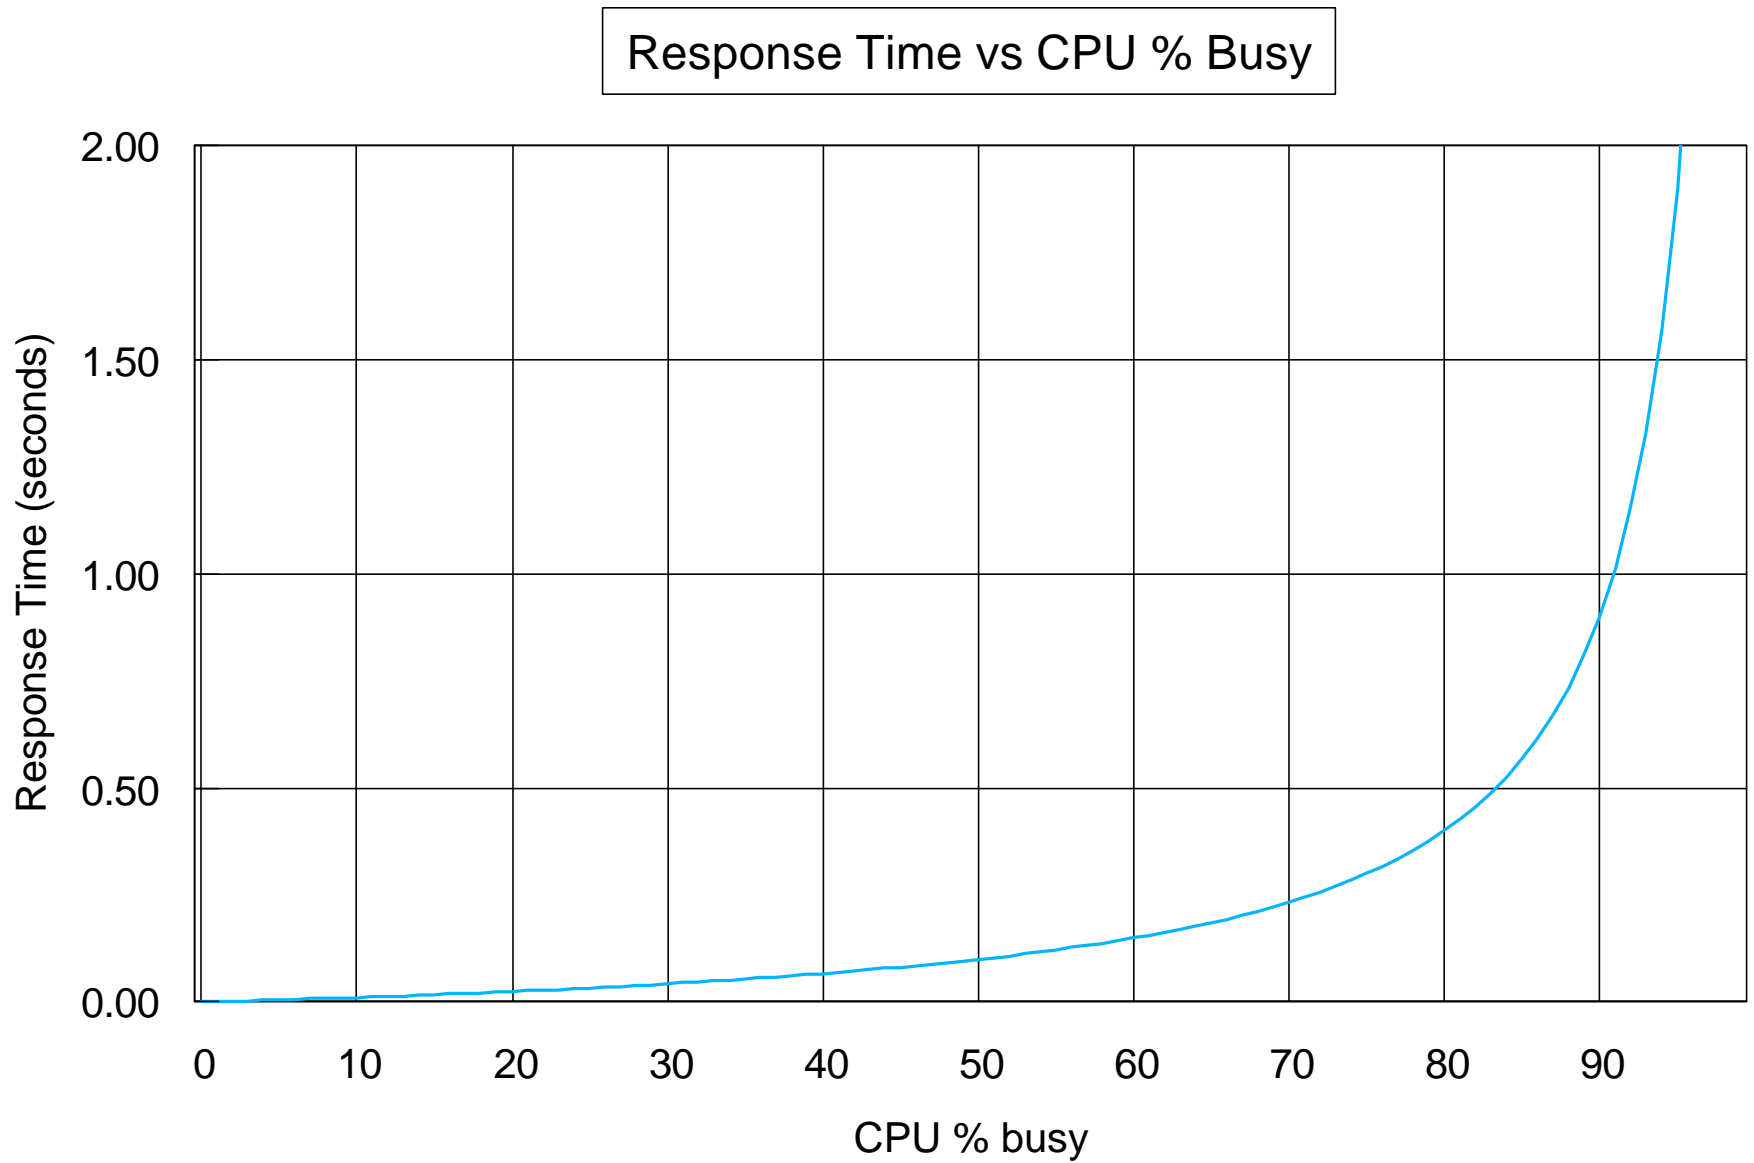
 - Navigating forward by watching your wake
 - Historical Data is key

Tracking System & Business Metrics:



Sum of all Service Class Usage = Total Utilization for this Image

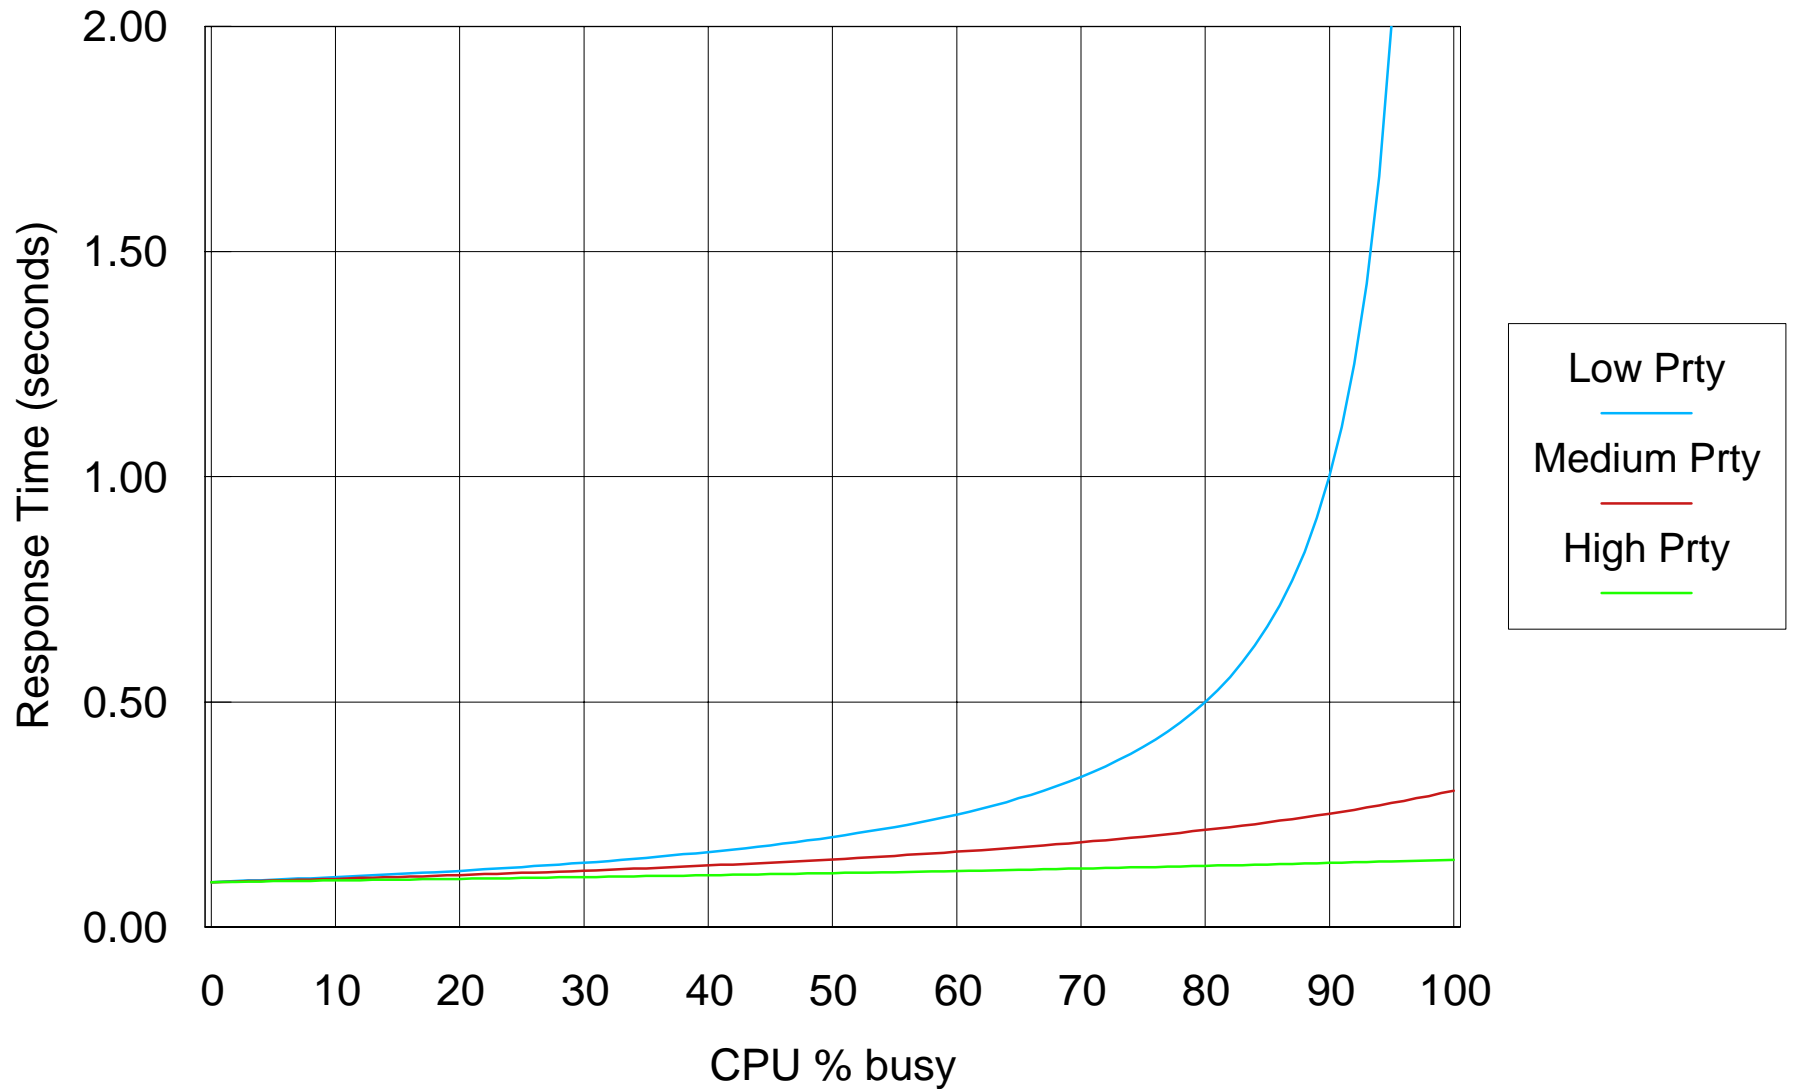
Performance Characteristics -The Hard Part



Performance has a non-linear relationship with processor utilization

Performance for Mixed Workloads

Response Time vs CPU (%) Busy
Compares Different Prty Workloads



Performance

- Workload Performance depends on many factors:

- Workload Priority
- Utilization of Higher Priority Work
- Number and speed of CPs
- Paging (to disk / to E-Storage)
- I/O Subsystem constraints
- LPAR effects
- Single CP constraints

- Without adequate tools, Capacity Planning dissolves into more art and less science relying on "Rules of Thumb"
- i. e. When average utilization reaches some threshold

And by the way ...

- Upper management sometimes ask tough questions

How bad will things get if we don't upgrade?
How much better will it get if we do upgrade?
How long will it last before next upgrade?

- Tools are needed

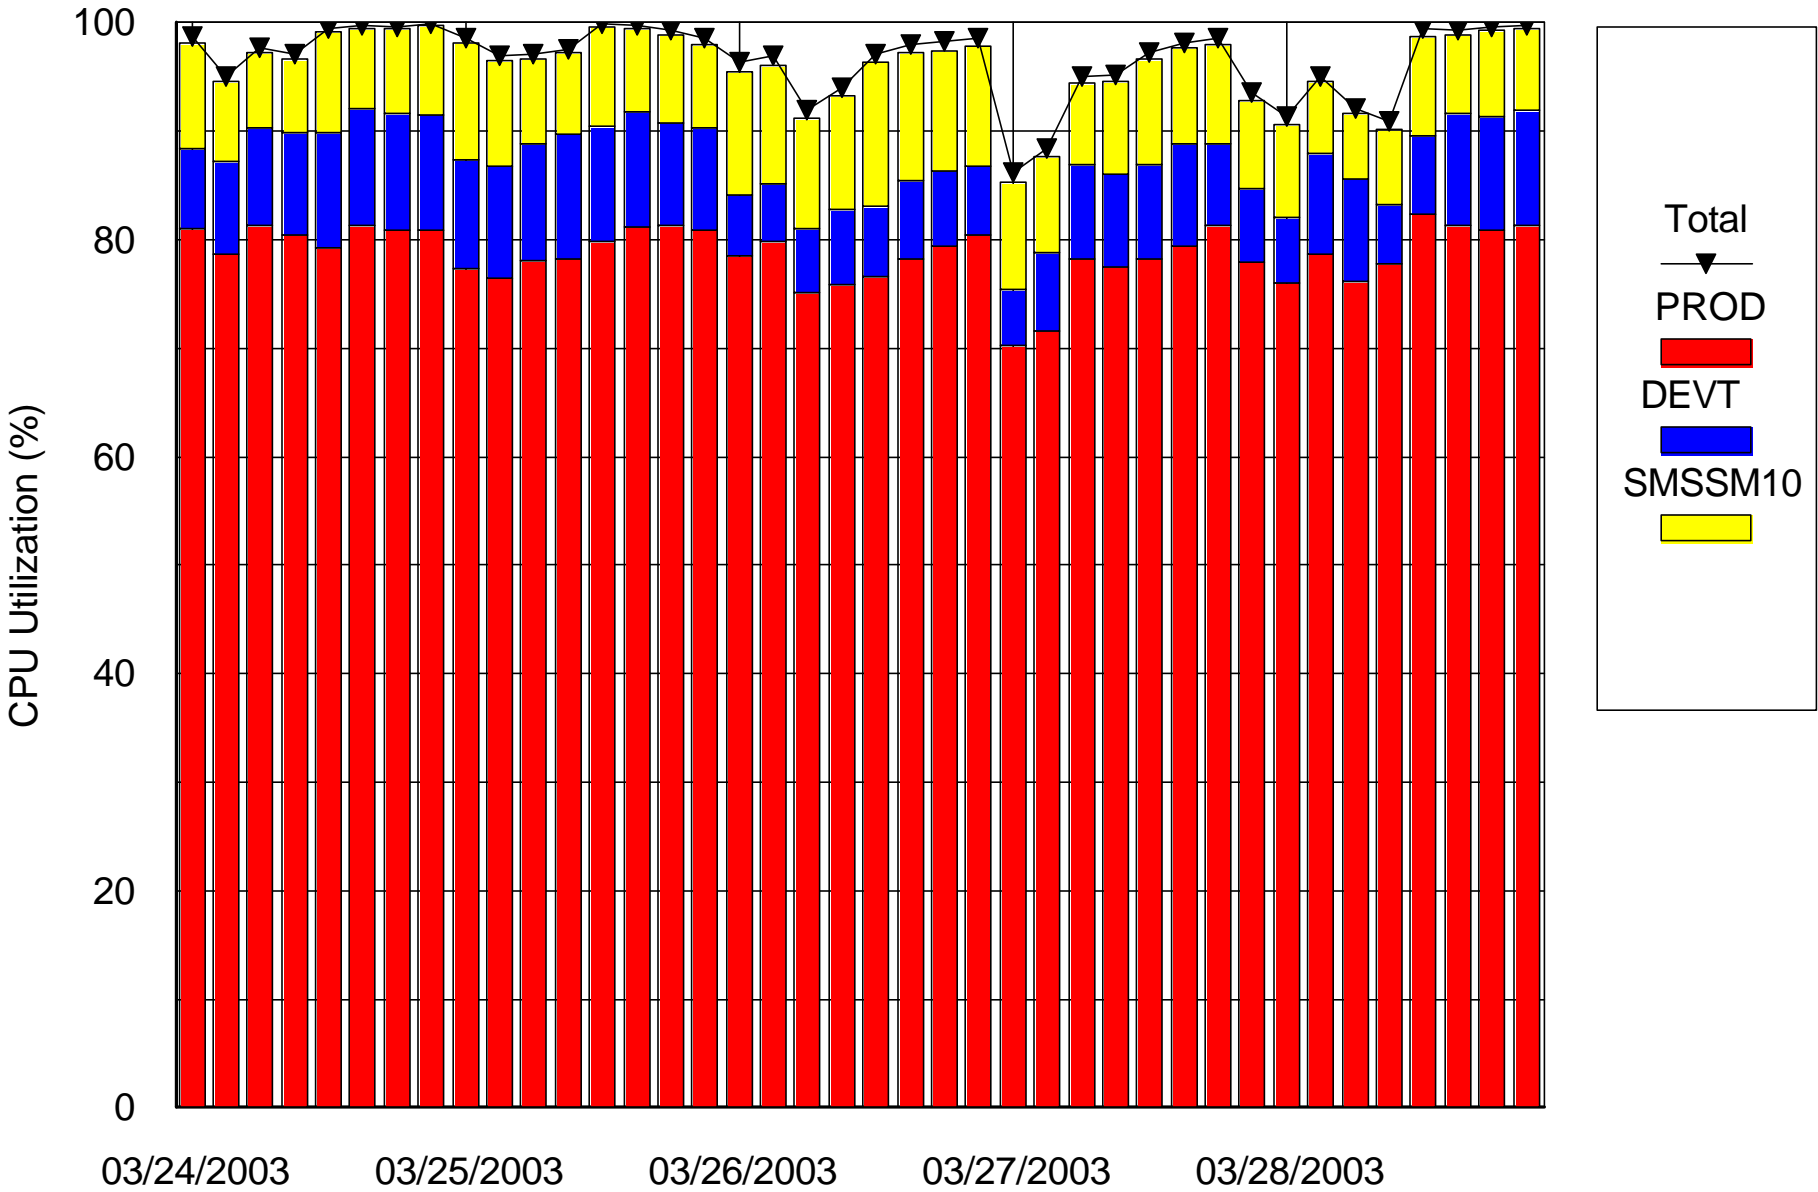
Tools that help predict changes to Workload Volumes
Tools that can predict changes to Workload Performance

- Performance Modeling Capability

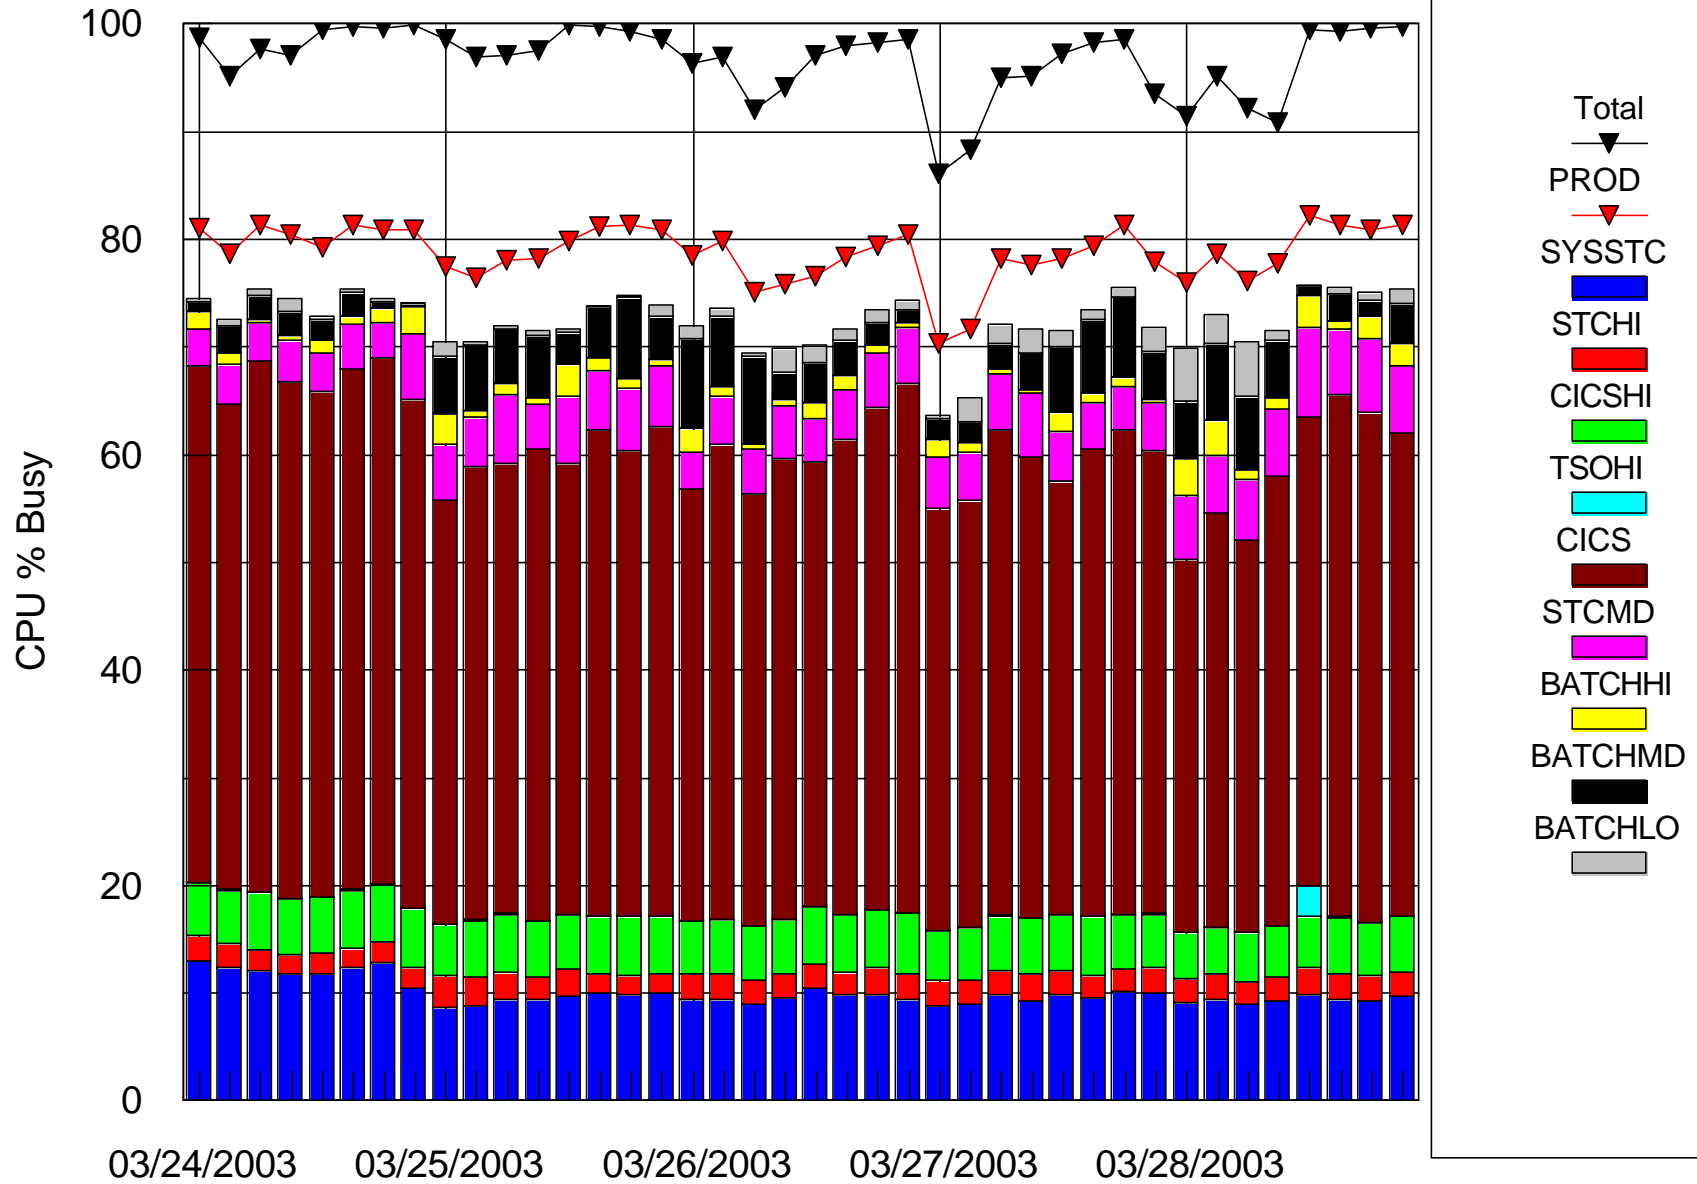
IBM Tivoli Performance Modeler for z/OS Demo (5698-A18)

- New Performance Modeling Tool from IBM/Tivoli.
- Version 2 GA on Aug/2003
- Windows/PC Based Simulation Model

2064-105 CPU Utilization by LPAR
for 9 AM thru 11 AM each day in 15 minute intervals



2064-105 Utilization by Workload for PROD LPAR
for 9 AM thru 11 AM each day in 15 minute intervals



Workload Analysis Report for PROD LPAR on 03/26/2003 at 10.30.00

<u>Workload</u>	<u>Serv Class</u>	<u>Period</u>	<u>Priority</u>	<u>CPU (%)</u>	<u>Total (%)</u>	<u>Cumulative</u>	<u>MPL</u>	<u>Xrate</u>	<u>Velocity</u>	<u>PI</u>	<u>Goal</u>	<u>Resp</u>
						4.34						
SYSTEM	SYSSTC	1	0	49.50	9.90	14.24	35.02	0.00	51.30	0.00	0.00	0.003
SYSTEM	SYSTEM	1	0	5.40	1.08	15.32	16.31	0.00	60.10	0.00	0.00	0.002
STC	STCHI	1	160	12.40	2.48	17.80	18.02	0.00	46.60	0.90	40.00	0.002
TSO	TSOHI	1	190	0.40	0.08	17.88	0.13	0.44	64.70	0.20	10.00	0.006
CICS	CICSHI	1	250	27.00	5.40	23.28	2.00	0.00	28.40	1.80	50.00	0.004
WEB	WEBHI	1	250	0.20	0.04	23.32	1.00	0.00	12.80	3.90	50.00	0.000
DDF	DDF	2	300	0.50	0.10	23.42	12.94	0.02	56.10	0.00	0.00	0.007
DDF	DDF	1	300	0.60	0.12	23.54	33.79	0.47	0.00	2.90	0.00	0.005
TSO	TSO	1	300	0.30	0.06	23.60	0.10	0.37	39.10	0.50	0.00	0.004
TSO	TSO	2	300	0.20	0.04	23.64	0.01	0.00	60.00	2.00	0.00	0.005
WEB	WEBHTML	1	300	0.10	0.02	23.66	0.30	0.72	0.00	0.40	0.00	0.004
CICS	CICS	1	350	232.90	46.58	70.24	9.00	0.00	26.60	1.90	50.00	0.010
WEB	WEB	1	350	0.20	0.04	70.28	1.00	0.00	0.00	1662.00	50.00	0.000
STC	STCMD	1	370	25.30	5.06	75.34	14.46	0.00	71.10	0.40	30.00	0.005
BATCH	BATCHHI	1	393	3.90	0.78	76.12	0.34	0.04	50.30	0.10	7.00	0.008
TSO	TSO	3	470	0.10	0.02	76.14	0.02	0.00	11.50	2.60	30.00	0.004
BATCH	BATCHMD	1	496	10.60	2.12	78.26	3.98	0.00	12.00	0.30	4.00	0.006
BATCH	BATCHLO	1	596	5.50	1.10	79.36	1.35	0.00	6.70	0.60	4.00	0.007
				Total		75.02						
				Actual		79.36						
				Uncaptured		4.34						
				C.R. (%)		94.53						
				All LPARs		98.24						

Simulator Results Summary for PROD LPAR (based on 3/26/03 at 10:30 AM - 10:45 AM)

<u>Run Description</u>	<u>Total MIPS</u>	<u>MIPS per CPU</u>	<u>Total CPU % Busy</u>	<u>PROD % Busy</u>	<u>CICSHI</u>	<u>CICS</u>	<u>BATCHMD</u>	<u>BATCHLO</u>
Base Run on 105	1055.00	211.00	98.03	79.16	0.02	0.06	17.97	14.24
Plus 10% Growth on 105	1055.00	211.00	99.91	79.14	0.02	0.25	594.20	366.46
% change=>	0.00%	0.00%	1.91%	-0.02%	0.96%	343.40%	3207.15%	2473.87%
Plus 20% Growth on 105	1055.00	211.00	99.97	77.31	0.02	4.95	2815.37	1302.65
% change=>	0.00%	0.00%	1.97%	-2.33%	3.03%	8745.99%	15569.59%	9049.24%
Base Run on 107	1399.02	199.86	79.68	65.44	0.02	0.05	5.78	3.36
% change=>	32.61%	-5.28%	-18.73%	-17.33%	-9.40%	-17.50%	-67.84%	-76.41%
Plus 10% Growth on 107	1399.02	199.86	86.44	70.78	0.02	0.05	6.34	3.84
% change=>	32.61%	-5.28%	-11.83%	-10.58%	-8.52%	-13.55%	-64.71%	-73.01%
Plus 20% Growth on 107	1399.02	199.86	92.84	75.76	0.02	0.05	7.98	5.27
% change=>	32.61%	-5.28%	-5.30%	-4.29%	-0.62%	-8.58%	-55.56%	-62.97%

Capacity Planning Methodology

- Planning based on formalized SLAs great if.....

They are documented, understood, tracked and modeled
but

Most SLAs are less formal "i.e. performance no worse than today"

- Real issue is -

When do we hit the knee of the curve
Each workload behaves differently
Can we quantify it

- Test/Development first to feel the pain and usually runs Discretionary

Consider impact on productivity