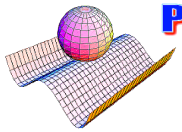

Capacity Planning Boot Camp

Part I: Getting Started

Dr. Neil J. Gunther
Performance Dynamics Company
Castro Valley, California
www.perfdynamics.com

CMG-T Session 405
CMG 2008 Conference, Las Vegas, Nevada

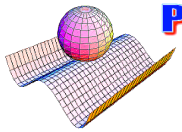

Table of Contents

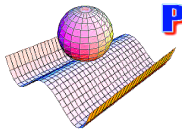

Modern Motivations

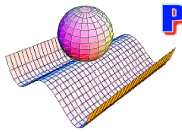

Minding the Marketplace

It's important to keep an eye on the marketplace.

It will ultimately impact your capacity planning.

Some examples:

1. General trend toward mobile devices will appear in your shop.
2. Rootkit wars. Virtualization performance and security.
3. Intel's very aggressive scale-down in VLSI processors.

Guerrilla Mantra 1.8 **Squeezing Capacity**: *Capacity planning is not just about the future anymore.* In today's economic climate, you need to squeeze more out of your current capital equipment.

Performance Mgmt vs. Systems Mgmt

A pile of shrink-wrap or COTS solutions.

- **Systems Management:**
  - Backup/recovery
  - Chargeback
  - Security
  - Distribution of software
  - Performance management

But, in reality, performance management is a complex set of disciplines.

- **Performance Management:**
  - Monitoring
  - Analysis
  - Planning

Performance Management Spectrum

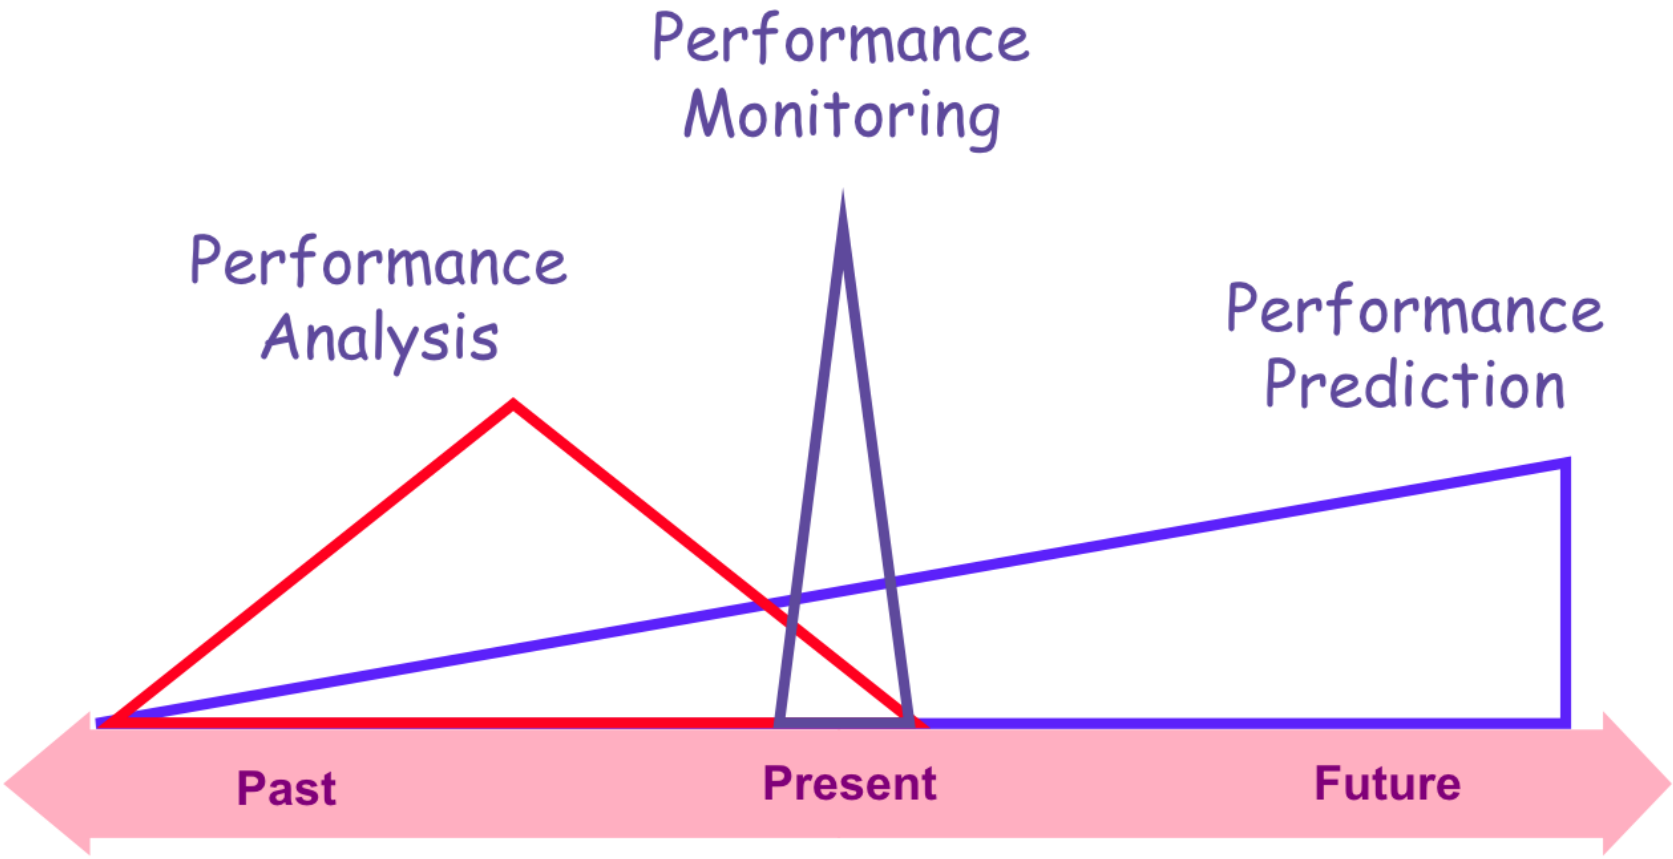

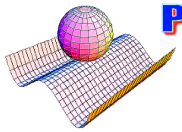

Performance Monitoring

Requires:

- Metric capture (sampled data)
- Metric display
- Alarming (from asynchronous events)
- Logging of metrics with timestamps

Guerrilla Mantra 2.1 **Monitoring vs. Modeling**: *The difference between performance modeling and performance monitoring is like the difference between weather prediction and simply watching a weather-vane twist in the wind.*

Performance Analysis

Requires:

- Metric with timestamps
- Data visualization
- Regression analysis (ANOVA)
- Time series analysis (ARIMA)

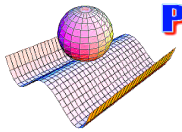

Performance Prediction

AKA capacity planning. Inevitably involves models with which to make predictions.

> **Guerrilla Mantra 2.4 The Big Picture**: *Unlike most aspects of computer technology, performance modeling is about deciding how much detail can be ignored!*

But in today's business environment management is generally only looking for a *sense of direction*. Anything else takes too long and will change next week, anyway.

Therefore (unlike the days when "big iron" cost real money), there is no need to produce the exquisitely precise compass bearing. In fact, taking the time to do that is probably a bad career move.

A Performance Model is Not ...

... like a model train set.

It's More Like a Map

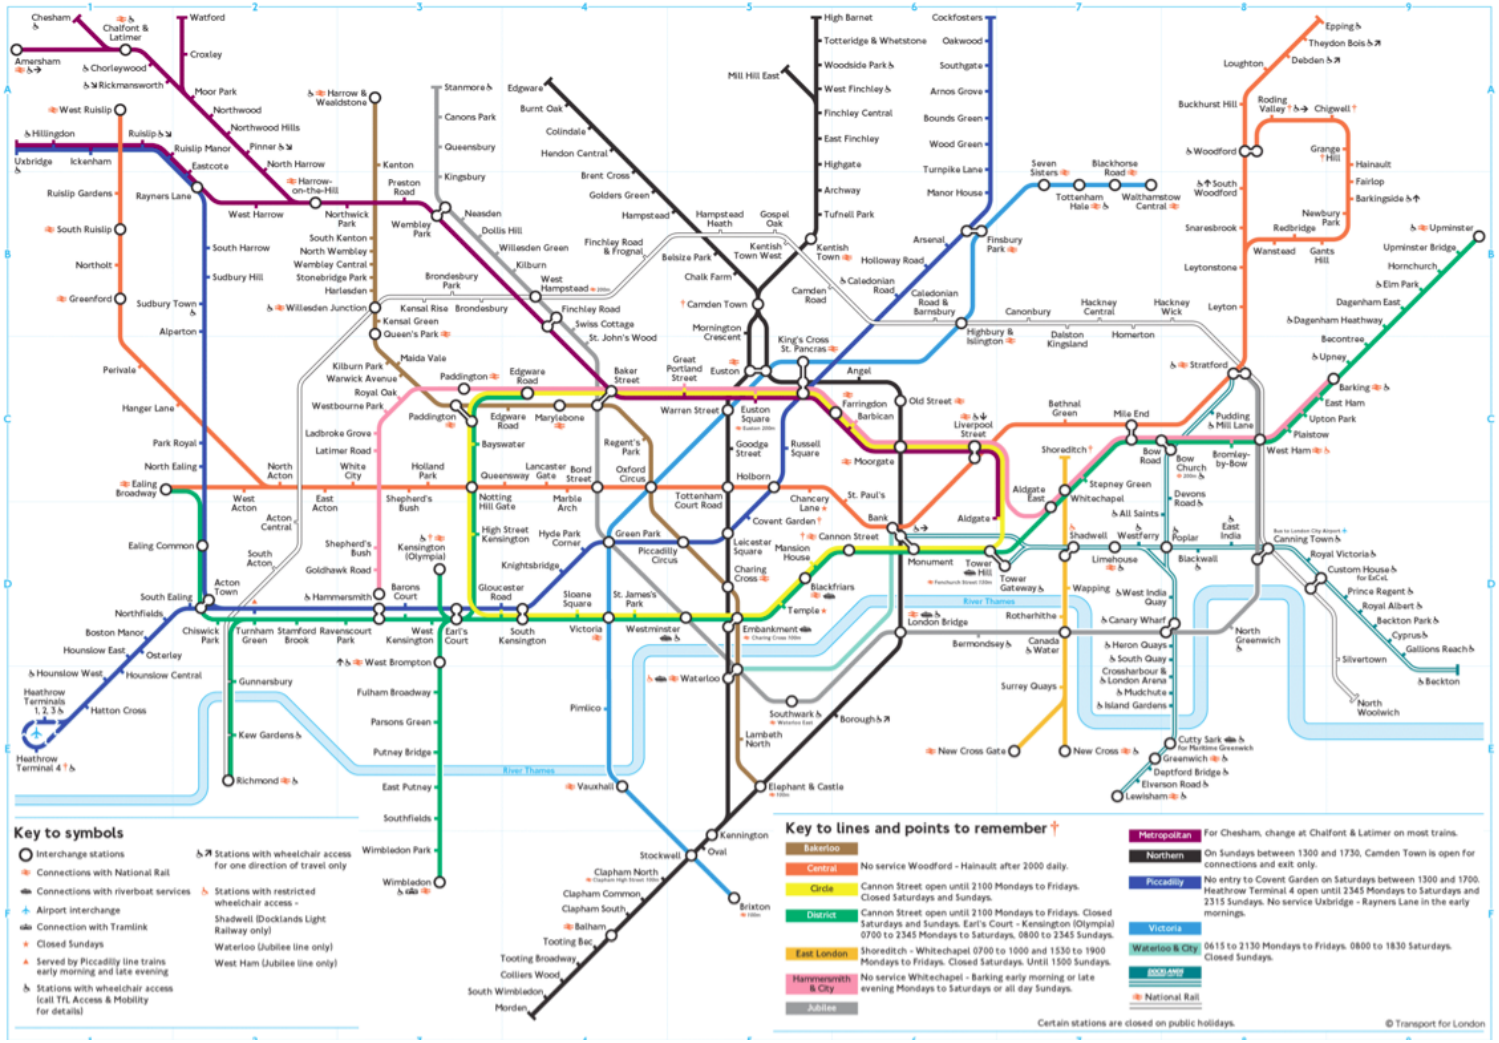

Risk Management vs. Risk Perception

Public Web 2.0 Embarrassments

In 2008 alone, we saw performance failures at:

- Twitter.com
- Amazon Elastic Cloud
- Cuil.com (supposed Google killer)
- Apple iStore
- Google Gmail

All due to lack of proper capacity management.

Why do companies keep doing this? It's not good for business. Are they insane?

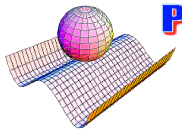

Pick a Number

My favorite example:

> California DMV wanted an IT system that would integrate the License and Registration databases (using image files, etc.)
>
> After 5 years and $40 million[a], the project was declared unsuccessful and terminated.

If the world were sane (which apparently it is NOT), you'd think that every CEO, CIO, CFO, etc., would be petrified of similar disasters befalling them.

Therefore, they should treat us performance analysts and capacity planners like the Gods (that we are) instead of someone just one notch above the janitor.

---

[a] When I told this story to another client of mine, they replied: "Pfft! We regularly lose that amount every quarter!"

What's wrong with this picture?

The CEO (from the previous slide) is getting ready to fly to an important company branch meeting.

He has the TV on and hears that a commercial airliner just crashed and hundreds of people were killed.

This event plays on his mind and he continues to worry about the safety of *his* flight while driving to the airport.

Remember that even if we could tell the CEO that statistically he is 30-40 times more likely to die in a car accident, he will still continue to worry about his flight.

Statistics Can't Beat Perception

Preaching to management about the goodness and benefits of CaP usually falls on deaf ears.

Even if you tell the CEO that his best strategy is to fly to the airport. ☺

Therefore, don't do it. There is another way, which we present in the next section.

The Guerrilla Approach

Guerrilla Manual

Sometimes very difficult to get management support for CaP.

Difficult to convey validity of CaP concepts.

Useful to have an authoritative reference to cite [1]. That's the purpose of the Guerrilla Manual (in the back jacket)

Online version of Guerrilla Manual is always being updated at:
www.perfdynamics.com/Manifesto/gcaprules.html

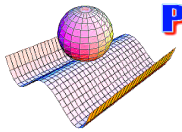

Tiger Team Tactics

Asking the right questions is often 90% of the battle.

Listen for opportunities to ask good performance questions.

Use ROT (rules-of-thumb)

Guerrilla Mantra 1.9 When Wrong is Right: *Capacity planning is about setting expectations. Even wrong expectations are better than no expectations.*

Collecting ROT

Conversing with more knowledgeable people e.g., engineers, architects.

- Keeping a detailed notebook with dates
- Reading trade rags
- SPEC.org for CPU and *in situ* benchmarks
- TPC.org for database benchmarks
- slashdot.org for general tech gossip
- arstechnica.com for PC-related technologies

Put it all in a database. Your own performance database.

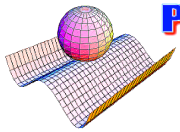

Example ROT: Pareto's Principle

Source: *"Extensive analysis of user traffic on Gnutella shows a significant amount of free riding in the system. Sampling messages on the Gnutella network over a 24-hour period, we established that almost 70% of Gnutella users share no files. We argue that free riding leads to degradation of the system performance and adds vulnerability to the system."* [2]

ROT: Too hard to recall exact percentages. Convert to Pareto 80:20 (or 75:25) rule: *Any distributed application which involves file sharing between uploaders and downloaders, expect only 25% will be writers, while 75% will readers.*

VIP if capacity planning for web sites like YouTube or pod-casting apps.

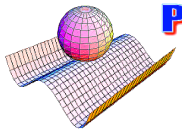

Tracking Rumors and the News

Source: Intel, AMD, IBM, Sun, et al., have heavily promoted *multicores*, i.e., no more single CPUs. The reason is, they run too hot! The power density (Watts) is linearly related to clock speed (GHz).

In Jan. 2007, however, Intel and IBM made a joint announcement that they would produce single CPU parts using 45 nanometer (nm) technology. Intel claimed they would produce *penryn* by the end of 2007.

News Watch: Did Intel deliver? Yes they did! And now (in 2008) they are heading for 32 nm.

What does this mean for the future of multicores?

We know that multicores are going to have concurrent programming issues.

VIP for your capacity planning for procurement of new servers.

Top Ten TPC Performance Tricks

www.tpc.org/tpcc/results/tpcc_perf_results.asp

- Home
- Results
  - TPC-App
  - TPC-C
  - TPC-E
  - TPC-H
- Benchmarks
  - TPC-App
  - TPC-C
    - **Results**
    - Description
    - FAQ
  - TPC-E
  - TPC-H
  - Pricing Spec
  - Obsolete
    - TPC-A
    - TPC-B
    - TPC-D
    - TPC-R
    - TPC-W
- Technical Articles
- Related Links

Top Ten TPC-C by Performance
Version 5 Results As of 8-Sept-2007 10:16 PM [GMT]

Note 1: The TPC believes it is not valid to compare prices or price/performance of results in differe

(•) All Results ( ) Clustered Results ( ) Non-Clustered Results Currency All

Rank	Company	System	tpmC	Price/tpmC	System Availability	Database	Operating System	TI Mon
1	hp invent	HP Integrity Superdome-Itanium2/1.6GHz/24MB iL3	4,092,799	2.93 US $	08/06/07	Oracle Database 10g R2 Enterprise Edt w/Partitioning	HP-UX 11i v3	BE Tux 8.0
2	IBM	IBM System p5 595	4,033,378	2.97 US $	01/22/07	IBM DB2 9	IBM AIX 5L V5.3	Micro COM
3	IBM	IBM eServer p5 595	3,210,540	5.07 US $	05/14/05	IBM DB2 UDB 8.2	IBM AIX 5L V5.3	Micro COM
4	IBM	IBM System p 570	1,616,162	3.54 US $	11/21/07	IBM DB2	IBM AIX	Micro

TPC Exec Summary: HP Superdome

	HP Integrity Superdome - Itanium2/1.6GHz/24MB iL3 - 64p/128c			TPC-C Revision 5.8 Report Date: August 21, 2007
Total System Cost	TPC Throughput	Price/Performance		Availability Date
USD $11,978,134	**4,092,799 tpmC**	**USD $2.93/tpmC**		**August 6, 2007**
Server Processors/Cores/Threads	Database Manager	Operating System	Other Software	Number of Users
64/128/256 Intel Itanium2 1.6GHz	**Oracle Database 10g Release 2 Enterprise Edition with Partitioning**	**HP-UX 11i v3**	**TUXEDO 8.1**	**3,226,200**

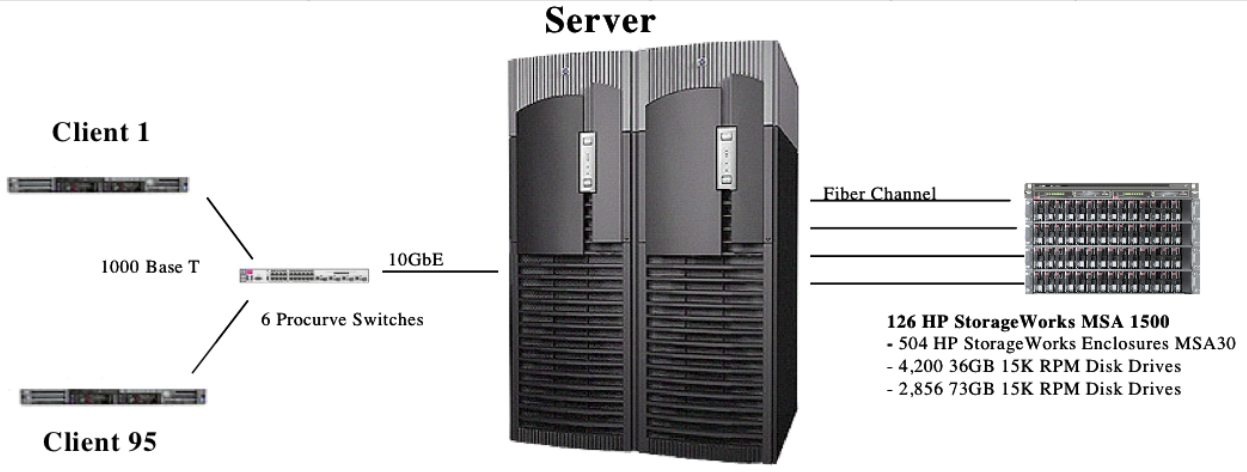

TPC Exec Summary: IBM System P5

IBM	IBM System p5 595 Model 9119-595	TPC-C Rev. 5.7 Report Date: January 22, 2007		
Total System Cost	**TPC-C Throughput**	**Price/Performance**	**Availability Date**	
$11,980,244 USD	**4,033,378**	$2.97 USD	**January 22, 2007**	
Database Server **Processor Chip/Core/Thread**	**Database Manager**	**Operating System**	**Other Software**	**No. Users**
32/64/128 POWER5+ 2.3GHz	DB2 9	AIX 5L V5.3	Microsoft Visual C++ Microsoft COM+	3,200,000

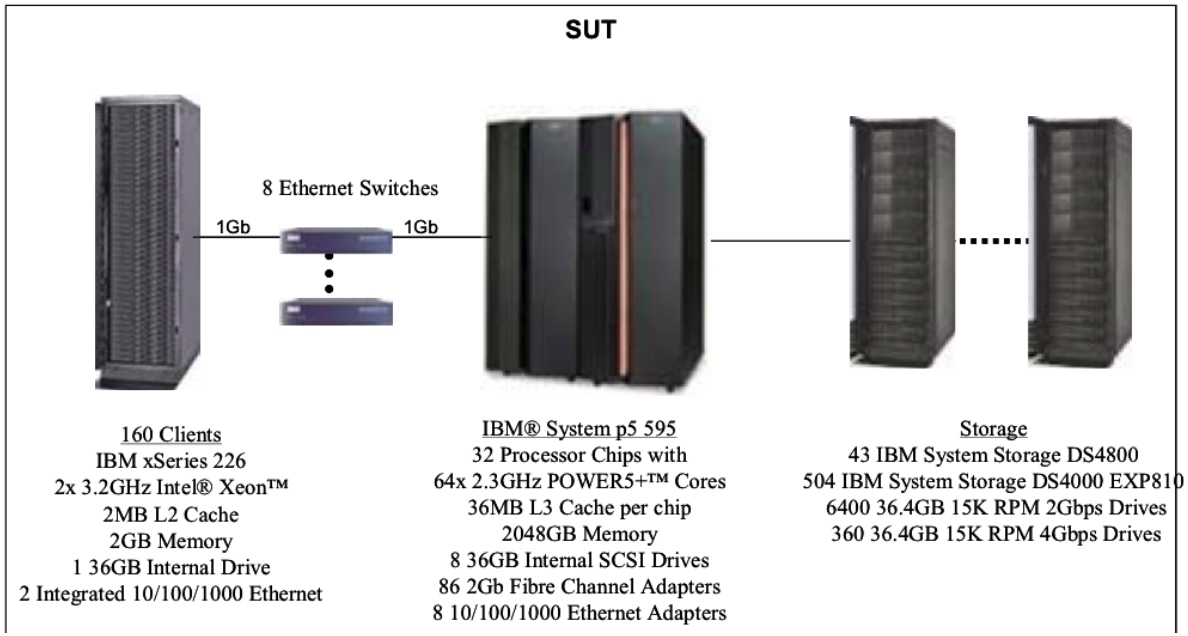

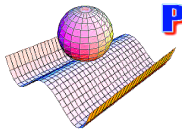

But On Closer Scrutiny ...

The TPMC rates (throughput in Transactions/Minute for the TPC-C benchmark) are within spitting distance of each other.

HP beats IBM by less than 1.5%. But HP uses 64-way 1.6GHz Itanium2 processors (128 cores), while IBM uses only 32-way 2.3GHz P5 processors (64 cores) or half the number of processors.

Even if we allow for the $\frac{2.3}{1.6} = 1.4375$ factor in clock frequency, HP would still need $1.4375 \times 32 = 46$ Itanium2 ("2.3GHz") processors.

Definitely IBM might have more tuning opportunities because the own DB2, but what else could be going on here?

Guerrilla Mantra 1.19 **Benchmarking**: All benchmarking is institutionalized cheating.

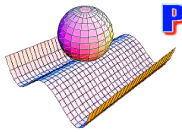

Tools and Techniques

Performance Prediction Methods

Primarily two basic methods:

1. **Statistical forecasting:**
   - Apply to raw data
   - Basically a form of trend analysis
   - No deeper abstraction
   - Cannot predict bottlenecks
2. **Queueing analysis:**
   - Must extract queueing parameters
   - Must create underlying abstraction
   - Solve "analytically" or by simulation
   - Can predict bottlenecks

We focus on queueing analysis using PDQ.

Statistical Modeling Tools

1. **Commercial:**
   - Microsoft Office *Excel*, www.microsoft.com
   - SAS *IT Resource Management*, www.sas.com
   - Insightful *S-Plus*, www.insightful.com
   - Minitab Inc. *Minitab*, www.minitab.com
   - SPSS *Trends*; www.spss.com
2. **Open Freeware:**
   - *R* subset of S-Plus, http://www.r-project.org/
   - *DataPlot* from NIST, www.itl.nist.gov/div898/software/dataplot/homepage.htm
   - *MacAnova* for DOE, www.stat.umn.edu/macanova/

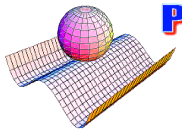

Queueing Analysis Tools

1. **Commercial:**
   - BMC*Perform-Predict*, www.bmc.com
   - TeamQuest *Model*, www.teamquest.com
   - Hyperformix *Performance Designer*, www.hyperformix.com
2. **Open Freeware:**
   - *SimPy* queueing simulator written in python.
     simpy.sourceforge.net/
   - *PDQ* queueing solver.
     http://www.perfdynamics.com/Tools/PDQcode.html

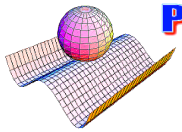

Simple Forecasting Model

Collect performance data for:

1. Aggregate number of IOs issued
2. Cumulative CPU busy time

over some measurement period.

Number IOs	CPU Time
14	2
16	5
27	7
42	9
39	10
50	13
83	20

Excel Scatter Plot

Simplest starting point is a scatter plot [3].

Too hard to see patterns in raw numbers.

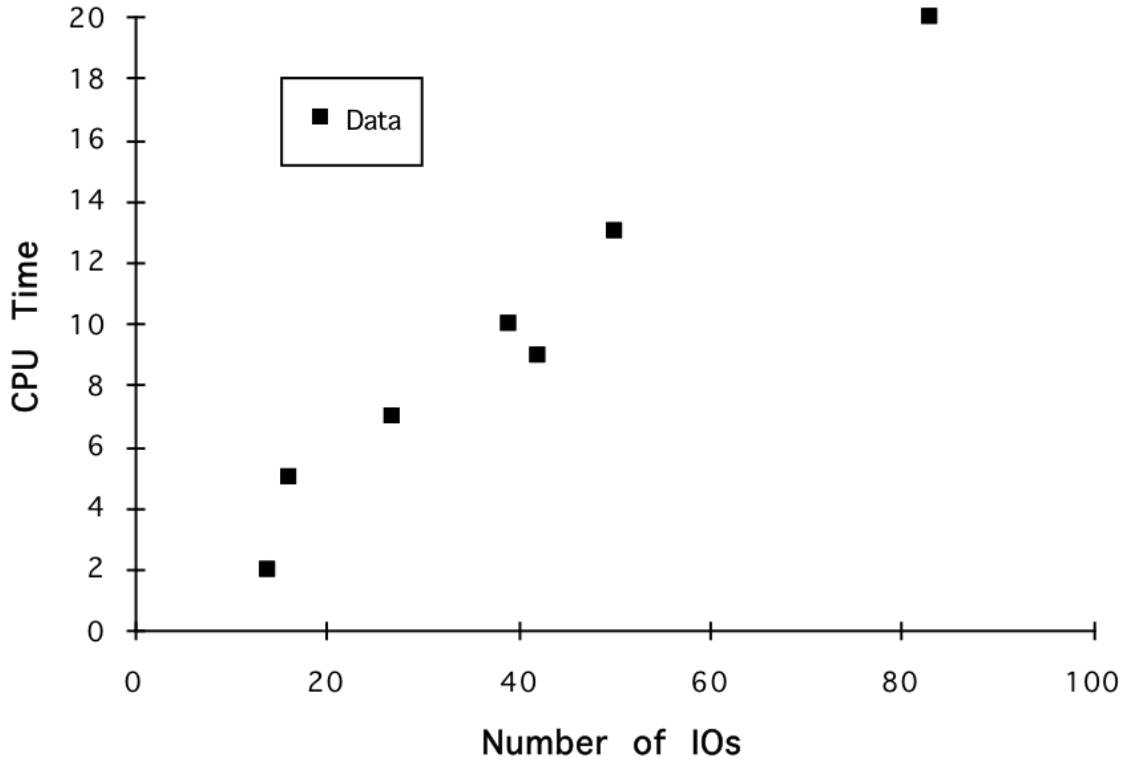

But is it a *linear* relationship or not?

Excel Regression Tool

Must have the Excel *Analysis ToolPak* [sic] add-in installed.

Go to: **Tools** menu → **Data Analysis ...** → **Regression**

Regression Statistics							
	Multiple R	0.98563213					
	R Square	0.9714707					
	Adjusted R Squ	0.96576484					
	Standard Error	1.08340978					
	Observations	7					
Analysis of Variance							
		df	*Sum of Squares*	*Mean Square*	*F*	*Significance F*	
	Regression	1	199.845402	199.845402	170.258442	4.7161E-05	
	Residual	5	5.86888379	1.17377676			
	Total	6	205.714286				
		Coefficients	*Standard Error*	*t Statistic*	*P-value*	*Lower 95%*	*Upper 95%*
	Intercept	-0.0082824	0.831105	-0.0099655	0.99237191	-2.1447023	2.12813755
	x1	0.24375637	0.01868107	13.0483118	1.2488E-05	0.19573524	0.2917775

Assumes a linear equation: $y = mx + c$ where intercept $c = -0.0082824$ and slope $m = 0.24375637$.

Linear Forecasting Model

$$\text{CPU time} = 0.24375637\ (\text{Number IOs}) - 0.0082824 \qquad (1)$$

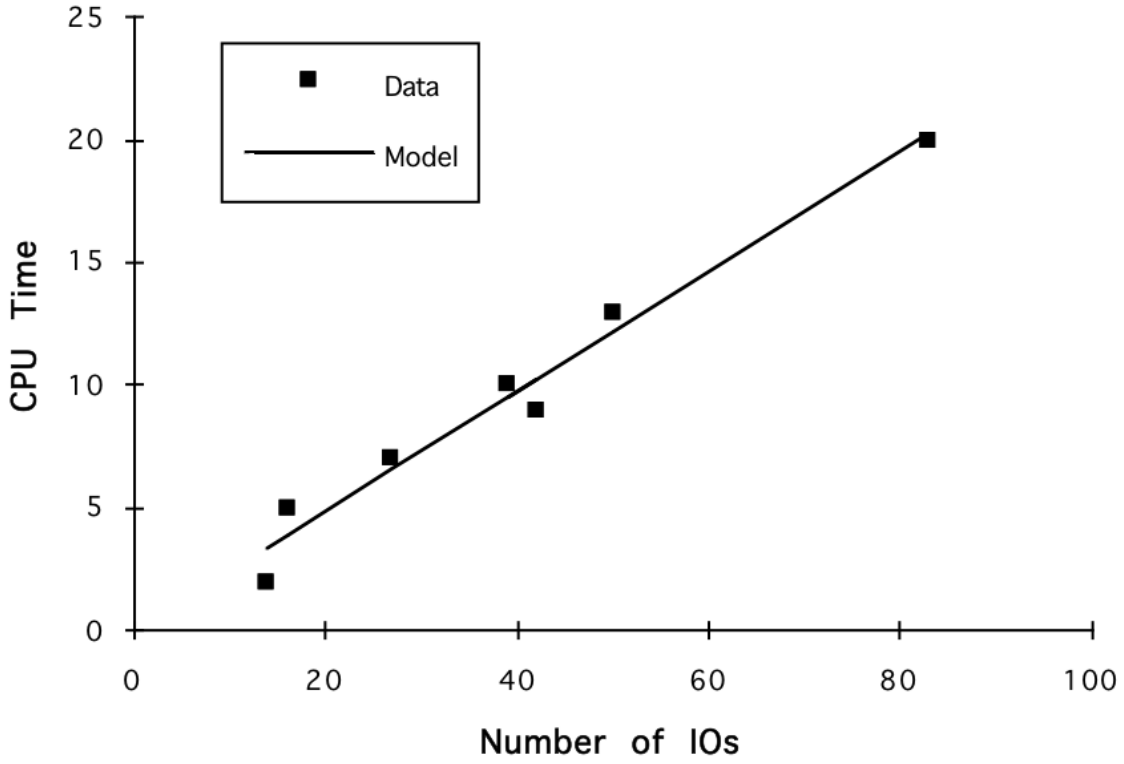

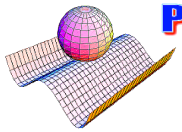

Summary

Reasons for doing CaP are very different from those of 40 years ago.

MIPS, memory and storage are all relatively cheap today (even for mainframes).

Recognizing the need for CAPACITY is easy. It's the *PLANNING* part that takes work.

Today, the emphasis is more on *tactical planning* than strategic planning, and that's where GCaP fits in.

GCaP uses the stealth of opportunistic insertion of CaP concepts and methods.

GCaP gets management buy-in through asking insightful questions in team meetings, not by preaching to the great unwashed.

A lot can be accomplished with simple tools like Excel.

References

[1] N. J. Gunther, *Guerrilla Capacity Planning*, Springer, 2007

[2] E. Adar and B. Huberman, "Free Riding on Gnutella," (PDF) October 2000

[3] Raj Jain, *The Art of Computer Performance Analysis*, Wiley, 1990

