The Association of System Performance Professionals

The **Computer Measurement Group**, commonly called **CMG**, is a not for profit, worldwide organization of data processing professionals committed to the measurement and management of computer systems. CMG members are primarily concerned with performance evaluation of existing systems to maximize performance (eg. response time, throughput, etc.) and with capacity management where planned enhancements to existing systems or the design of new systems are evaluated to find the necessary resources required to provide adequate performance at a reasonable cost.

This paper was originally published in the Proceedings of the Computer Measurement Group's 2004 International Conference.

For more information on CMG please visit http://www.cmg.org

Seeing the Forest AND the Trees: Capacity Planning for a Large Number of Servers

Linwood Merritt
Bank of America

This author's capacity planning learning expertise was built on a limited number of servers. Specifically, it involved analyzing and reporting workloads on a few mainframe footprints and, eventually, a few dozen distributed (Unix and NT) servers. Today's environments are growing in complexity and may include hundreds of distributed servers along with larger mainframe footprints. The challenge in this increasingly common environment is to manage the complexity of capacity planning and reporting for this explosion in the numbers and types of data processing platforms. Although this discussion is oriented to large companies (with a lot of "trees"), these approaches can be used by large and small companies alike.

1. Introduction

A common characteristic of companies today is a proliferation in the number of servers and the need to enhance the approaches to capacity planning and reporting for this environment. This data processing environment is significantly more complex that the "traditional" model (this author's background) of a large central (mainframe) server with some distributed functionality across a few dozen servers. This new environment often combines a large mainframe server presence with hundreds of distributed applications and servers. The capacity planning and reporting issues surrounding an environment with a large number of servers are outlined in this paper and fall into the following categories:

- Business-oriented reporting and acquisition of business data information using Capacity Councils
- Business driver based forecasting
- Exception detection / analysis
- Data capture analysis
- Bulk Capacity Planning
- Organization and reporting of server structure
- Visualization and reporting techniques

2. Business-Oriented Reporting and Acquisition of Business Data Information Using Capacity Councils

Business-oriented capacity reporting generally requires that contacts be made with business (or application) experts on a regular basis. This fact leads to a natural approach for organizing the Capacity Planning function within the application structure: the use of Capacity Councils [C99, C01, M02]. Capacity Councils are cross-organizational teams that bring together business and technical views of applications and supporting technologies. The maturity of this process may lend itself to a single Capacity Council, or a Capacity Council for each major business unit. There may also be platform-oriented types of Capacity Councils that bring together analysts who are responsible for given platform types. These technical Capacity Councils allow analysts to share common information independent of the business-oriented councils. This author's past experience with a robust Capacity Council structure is driving the potential formation of an *informal* (but documented) Capacity Council structure at his current company, with a similar deliverable: a quarterly capacity report for each line of business. You may already have an informal Capacity Council process in place if your company is currently producing capacity reports for lines of business (or is considering it).

Business-oriented Capacity Councils may consist of the following:

- Capacity Council leader, potentially from the business area
- IT Capacity Planner
- Business area representatives for primary applications and platforms
- Administrators for databases and platforms. For mainframes, this includes performance analysts.

The primary Capacity Council deliverable is a regular report (monthly or quarterly) that discusses the capacity status of each major business application, an analysis of significant variances from the previous report, and an overall capacity status of

the Capacity Council's environment. The stoplight rating of applications as "green," "yellow," and "red" is also a common reporting tool for evaluation of individual servers or workloads.

Each business Capacity Council can be chaired by a council leader from the relevant business area. The council leader schedules the meetings (or initiates e-mail or teleconferencing contacts), is responsible for the Council report, and is a conduit for capacity issues which arise within the council. In a multi-council environment, the final component of the Capacity Council process may be a regular Capacity Council Leaders meeting or conference call where council status ratings are discussed. The monthly council reports can be combined into a single Capacity Council report or web site for each period.

A significant outcome of the Capacity Council structure is that Capacity Planners' responsibilities may be driven by their Capacity Council assignments. This includes individual platform responsibilities as well as project assignments - when formal development projects identify a Capacity Planning need, the Capacity Planner who is assigned to the corresponding Capacity Council is the appropriate resource to address this need. In addition to reporting and project assignments, the Capacity Council is an invaluable tool for the Capacity Planner to gather and distribute capacity related data.

3. Business Driver Based Forecasting

Another significant output from the Capacity Council process is the identification of underlying business drivers (such as "number of accounts") and their mapping to applications and platforms, to enable the Capacity Planner to directly map server resources to business resources. The "demand side" forecasts from the business areas can then be used to build "supply side" capacity projections. The business driver data is generally produced by the business area and must include *historical* and *projected* values of the business driver(s). This data is generally most useful in a monthly format and, in this author's experience, is often contained in a set of spreadsheets that must be refreshed on a regular basis with new actual and projected data.

To use business driver data for capacity planning, it must be mapped to each relevant workload and platform. Early efforts by this author involved spreadsheet regressions that matched one or more business drivers to a resource such as CPU utilization on a server. In *Figure 1*, the "forecast" function derives a linear "least squares" trend of actual values in column B vs. business driver values in column C. Projections are then calculated from business driver projections that begin in row 27.

	A	B	C
	Month	**Actual**	**BusDriver**
23	Sep-02	6049	23133576
24	Oct-02	6310	23421420
25	Nov-02	6195	23645813
26	Dec-02	6406	23829750
27	Jan-03		23966066
28	Feb-03		24246086

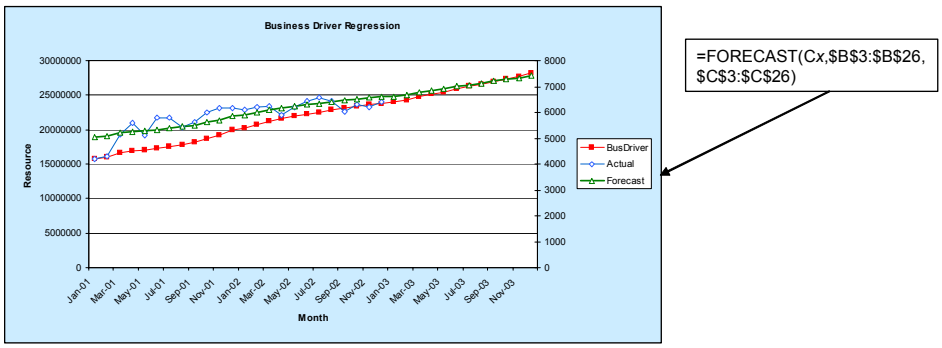

Figure 1 – Spreadsheet Based Business Driver Forecasting

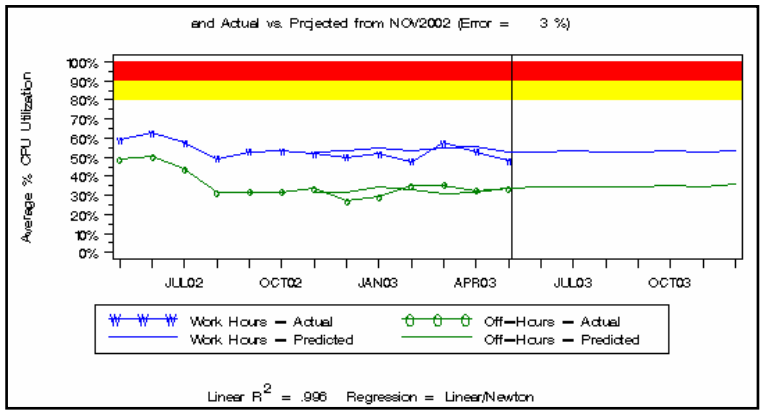

Figure 2 – Automated Business Driver Forecasting

```
Exception Detection Report for 05/03

_____________________________________________________
CPU_Utilization exception unix/unisys/tandem/MVS 5 boxes list: ServerA
ServerC ServerE ServerL ServerZ

_____________________________________________________
CPU_Utilization NULL DATA unix/unisys/tandem/MVS 0 boxes list:

_____________________________________________________
CPU_Utilization insufficient DATA unix/unisys/tandem/MVS 1 boxes list:
ServerG ==============================================
CPU utilization was greater than 50% yesterday for:
ServerA ServerD
```

Figure 3 – Automated Exception Detection E-Mail Notification

Spreadsheet based forecasting is an excellent approach to analysis of a few targeted servers or workloads. To forecast resource utilization for a large number of instances, an automated approach should be considered. This author used the SAS language to automate the production of resource projections from business driver inputs (see *Figure 2*). If projections are stored in a database, automatic "actual vs. projected" comparisons can also be displayed.

This projection system utilized spreadsheets to configure servers, applications, business drivers, and date-specific effects such as upgrades. This system also utilized spreadsheet versions of business driver data (historical and forecast), with a future view of directly accessing databases with these figures. The set of spreadsheets was read by a SAS program and converted to SAS datasets residing on the Capacity Planning server. These datasets were read by scheduled programs to perform multivariate regressions and derive linear coefficients from business driver and CPU utilization inputs. These coefficients were then multiplied by business driver forecasts to derive CPU forecasts.

4. Exception Analysis / Analysis

As the number of servers within each Capacity Planner's responsibility increases, it becomes more time-consuming to detect capacity issues as they arise. Although this may be considered a *performance* versus *capacity* issue and theoretically outside the responsibility of the Capacity Planner, the planner should be aware of resource utilization and be "in the loop" for such issues. There are two approaches to automate the detection of capacity and performance issues – the coding of symptoms to trap and report [K03], and the use of a "Statistical Process Control" mechanism [T01, T02, T03] to detect anomalies of resource measurements when compared to historical measurements (for example, when measurements exceed two standard deviations from a calculated mean). When problems are detected, notification of Capacity Planners via e-mail can detail exceptions on servers assigned to him/her (see *Figure 3*).

5. Data Capture

There is an operational side to Capacity Planning that involves the automation of data collection, performance database population, and report creation. A common approach is to use a script-based data pull to acquire capacity data for processing on a central server. For data-capture from a large server farm, an automated process should be implemented to check the successful completion of each step of the process and the completeness of data for each server. Output from the data-capture analysis can be integrated into the e-mail exception reporting mechanism described in the previous section.

6. Bulk Capacity Planning

When a large number of servers could be impacted by business initiatives, a high-level “bulk capacity planning” spreadsheet [M01] can be used for analysis of the entire set of servers in a single pass. This author has utilized this approach several times (for multiple companies). This approach takes imported utilization data for each server and calculates the effects of the business scenario. Wholesale impacts can be viewed with color-coded “conditional formatting.” In one such study (see *Figure 4*), an upgrade date was calculated from growth rates using the formula derived below, and color-coded "red” or “yellow” if the upgrade date occurred before a configured point in time.

$$\text{Threshold} = \text{Baseline} * (1+\text{AnnualGrowth})^{(\#\text{Years})}$$

$$\text{Log}(\text{Threshold}) = \text{Log}(\text{Baseline}) + (\#\text{Years}) * \text{log}(1+\text{AnnualGrowth})$$

$$(\#\text{Years}) = (\text{Log}(\text{Threshold}) - \text{Log}(\text{Baseline})) / \text{Log}(1+\text{AnnualGrowth})$$

7. Organization and Reporting of Server Structure

The organization of information about a large number of servers can be greatly assisted by a database of server characteristics and assignments. On approach is a server matrix that contains such information as business unit classifications, applications, Capacity Planner assignments, configuration details, and status color codes. This database system can reside on a spreadsheet or a database product, utilized as a central repository of capacity information, and used as a data source to build browser pages on the company’s Intranet. The organization of the browser reporting structure presents a color-coded view of servers by business area and application.

Business Unit	Server	Upgrade%	Base Upgd	Scen Upgd	Trend Upgd	Notes	Business Driver	Base Time	Base CPU%
DeptA	ServerA	75.0%	Nov-07	Oct-06	Jan-05		Widgets	Oct-01	31.6%
	ServerB	75.0%	Mar-04	Sep-25	**Sep-03**	Not impacted by Scenario	Widgets	Oct-01	45.0%
	ServerD	75.0%	Aug-04	**Nov-02**	**Apr-02**		Widgets	Oct-01	65.0%
DeptB	ServerC	75.0%	**Sep-02**	May-11	Oct-05		Gadgets	Oct-01	19.0%
	ServerE	75.0%	May-06	Mar-04	**Dec-03**		Gadgets	Oct-01	35.6%
	ServerF	75.0%	**Dec-03**	Jan-11	**May-02**		Gadgets	Oct-01	61.0%
DeptC	ServerG	75.0%	**May-02**	Mar-19	Dec-06	Scen Upgd Feb-03	Things	Oct-01	29.0%
DeptD	ServerI	75.0%	Oct-05	Nov-04	Sep-07		Things	Oct-01	32.6%

Figure 4 – Bulk Capacity Planning Spreadsheet

8. Visualization and reporting techniques

There are two approaches to organizing the reporting of the "forests" (business areas or large servers) along with the "trees" (servers, LPARs or workloads). A matrix reporting structure (such as discussed in the previous section) can display business and application level status information with drilldown (or thumbnail) reporting of individual "trees" (see *Figure 5*). In this example, the color logic rolls up the "worst" status color ("yellow" in this case) to the application and business fields.

To display a large number of entities as color-coded rectangles in a single chart, the Treemap concept, developed and published by Ben Shneiderman [S98] and the University of Maryland, can be used to display servers or workloads as rectangles, with the size of each rectangle representing a server's capacity rating and a color code representing its "green/yellow/red" capacity status or its CPU utilization. Another example of the Treemap concept can be found at the Smart Money web site [S02]. The "Treemap" concept can be used to group and display all mapped servers in the company (see *Figures 6* and *7*).

Dept A Systems Capacity Status

Click on the underlined text for charts

Business	Application	Database	Server	CPU	Memory	Disk I/O
Dept A	Application 1		MVS1	CPU Trend		Disk I/O Trend
	Application 2		MVS2	CPU Trend		Disk I/O Trend
		ProdDB1	Unix1	(5) CPU Trend Forecast Analysis	(8) Memory Trend	Disk I/O Trend

Figure 5 – Matrix Based Reporting

Figure 6 – Treemap Chart of Capacity Status

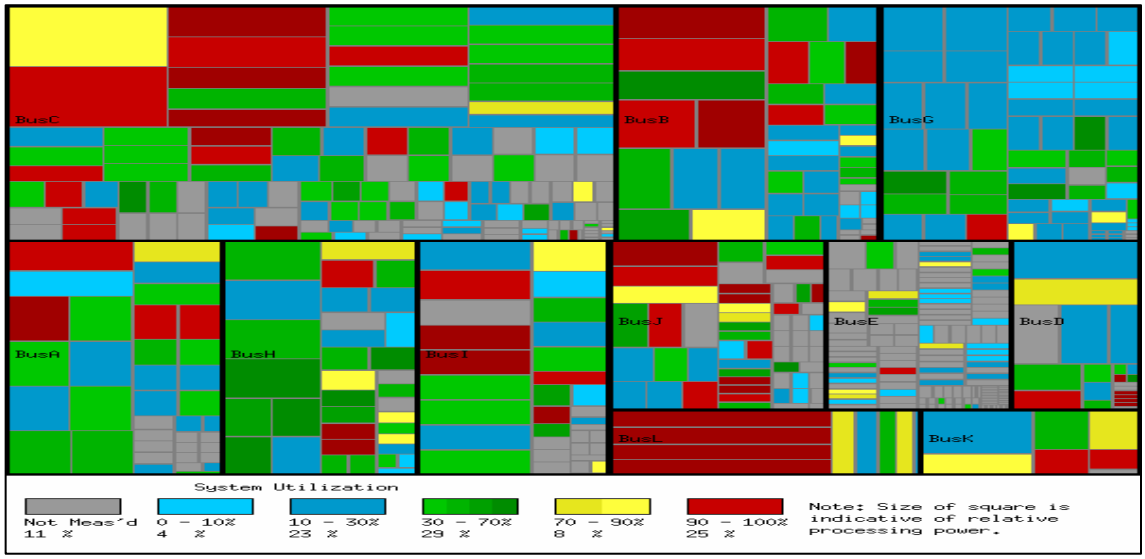

Figure 7 – Treemap Chart of Server Utilization

The Java application code for the Treemaps was designed to use a delimited text file as input. This text file was created from the performance database and server matrix data, and could be tailored according to server type ("production" versus "development") or subset into business areas. Drilldown capabilities were designed into the application for server-specific displays. A "Production" server version of these charts was used in a monthly Capacity Council summary to upper management.

9. Summary

For many companies, the role of Capacity Planning has become significantly more complicated as the number of servers has grown. The amount of effort for planning and reporting does not grow linearly with the number of servers. The Capacity Planner must be able to analyze and track many more servers at a time within a given business context. The operational aspects around processing and reporting data must be refined. A connection to the business becomes increasingly important and more complicated, particularly as applications cross business units and servers. The identification of trends and exceptions becomes more difficult and potentially more time-consuming. Finally, the difficulty of reporting a large number of servers in a concise and meaningful way increases significantly.

This author recommends that his readers actively consider the impacts of a growing number of servers in their own companies and consider implementing such measures as the following:

- One or more Capacity Councils, informal at first
- Some automation of exception detection and business driver-based forecasting
- A "matrix" database of relevant server information
- Redesigned approaches for processing, analyzing, and reporting data for large numbers of servers

10. References

[C01] Chaney, Bob, "Divide and Conquer: Implementing the Capacity Performance Council in Pieces," CMG2001 Proceedings

[C99] Chaney, Bob, "The Capacity Performance Council, Start Yours Today!" CMG99 Proceedings

[K03] Kaminski, Ron, "Automatic Process and Workload Pathology Detection," CMG2003 Proceedings

[M01] Merritt, Linwood, "Capacity Planning for the Newer Workloads," CMG2001 Proceedings

[M02] Merritt, Linwood, “A Capacity Planning Partnership with the Business,” CMG2002 Proceedings

[S02] SmartMoney Market Map, http://www.smartmoney.com/marketmap

[S98] Shneiderman, Ben, “Treemaps for space-constrained visualization of hierarchies,” http://www.cs.umd.edu/hcil/treemaps, December 26, 1998 and November 8, 2000

[T01] Trubin, Igor, Ph. D. and McLaughlin, Kevin, “Exception Detection System, Based on the Statistical Process Control Concept,” CMG2001 Proceedings

[T02] Trubin, Igor, Ph. D., “Global and Application Level Exception Detection System, Based on MASF Technique,” CMG2002 Proceedings

[T03] Trubin, Igor, Ph. D., “Disk Subsystem Capacity Management, Based on Business Drivers, I/O Performance Metrics and MASF,” CMG2003 Proceedings