Issue
One Hundred Twenty-Seven
Autumn 2010

CMG

Journal of Computer Resource Management

A Publication of
The Computer Measurement Group, Inc.

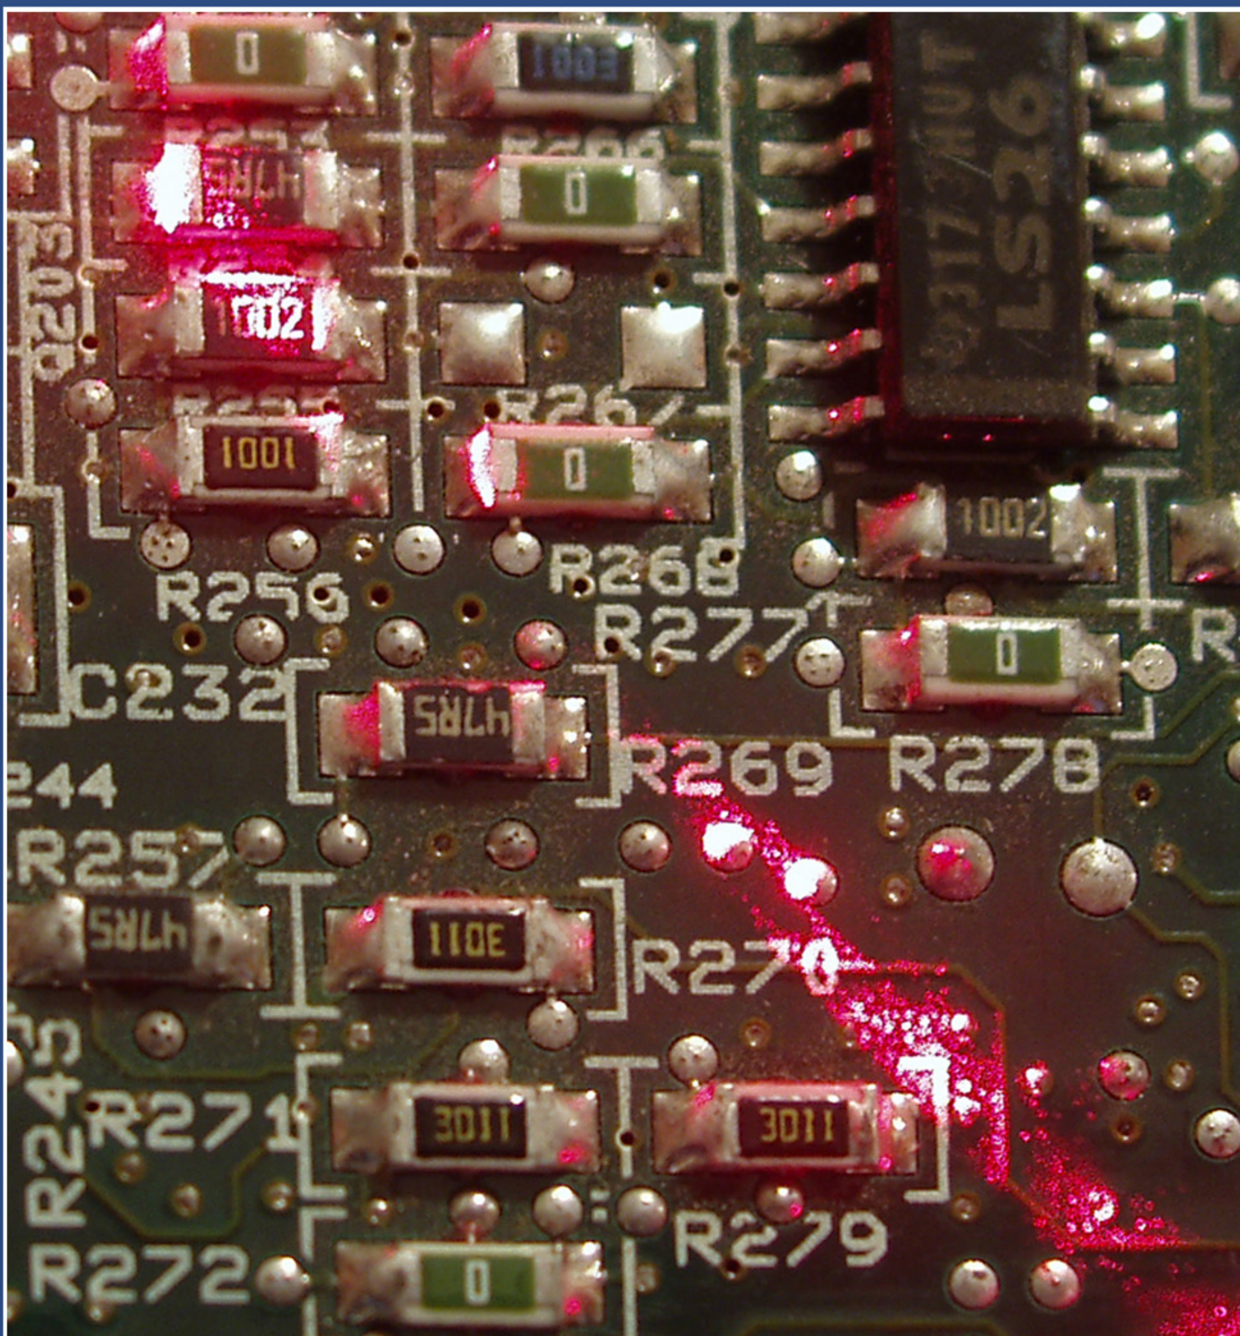

The Computer Measurement Group, Inc. (CMG) is the most influential organization worldwide for the exchange of information among Computer Performance Evaluation (CPE) professionals. To keep CPE professionals on top of their field, CMG is dedicated to providing its members the education, networking, and leadership opportunities vital to success in today's competitive industry.

PUBLICATIONS & MEMBERSHIP BENEFITS:

- The ***Proceedings*** contain all the information presented at the annual CMG conference. With more than one hundred papers on current hot topics, the *Proceedings* is your best source of the latest information in the expanding field of CPE. Some Late Breaking papers may not be in the *Proceedings;* they appear in a *Journal* issue.
- The ***Journal*** provides an additional source of information with emphasis on special interest topics and/or focus on specific CPE industry issues. Members receive the *Journal* at least three times a year.
- The ***Bulletin*** is a newsletter written by members for members. This publication is available in electronic format and includes vital information on CMG's national, regional, and international groups.
- **Various Internet Services.** CMG provides various online membership only services including a directory of all CMG members.
- CMG's electronic newsletter, **MeasureIT,** sent to all those who subscribe.

INFORMATION FOR AUTHORS

The *Journal of Computer Resource Management (CMG Journal)* is distributed only to members of the Computer Measurement Group, Inc. (CMG). CMG membership consists primarily of practicing professionals in the area of computer performance, analysis, and capacity planning. The objective of the *CMG Journal* is to bring the membership information of current importance involving:

- performance characteristics of computer systems of all sizes and architectures
- techniques to measure, analyze, compare, predict, and report performance
- techniques for managing computer performance, including capacity management, computer cost accounting, and computer performance reporting
- integration of disparate components into hardware and software systems to meet organization needs.

The intent of CMG is to publish relevant papers as rapidly as possible, preferably within six months of completion. Authors can assist CMG in achieving this objective by submitting papers in a uniform format.

To submit your paper for consideration in the *CMG Journal,* it must follow a format consistent with CMG guidelines. For complete instructions on paper submittal you should download two documents from http://www.cmg.org/national/journal.html located on the CMG Website. These documents are 'Paper Guidelines and Instructions' and 'CMG License to Publish.' The first document will provide guidelines on format and text; the second document is a required form that must be signed/returned to CMG that grants your permission for CMG to publish your paper.

Here are some highlights on what is expected from authors:

- Camera-Ready Papers, in Microsoft Word, should be sent via email to journal@cmg.org
- Electronic copies of the paper are preferred; send two hardcopies of the paper in the required format (see 'Instructions' document for details) if an electronic copy is not possible. Hardcopies must be camera ready Papers should be submitted three months prior to the Journal publication date
- Please indicate if you wish to review the paper after editing, prior to publication
- Refer to all tables and figures by number

Authors must read and follow the 'Instructions for Preparing Papers' document. The above highlights do not include all the guidelines required for submitting your paper. Papers not complying with the instructions may be rejected for inclusion in the *CMG Journal.*

Journal of Computer Resource Management

A Publication of the Computer Measurement Group

Issue 127

Table of Contents

CMG Journal #127: Letter from the Editor

Welcome to CMG Journal number 127, our third issue for 2010. What was originally intended to be the Summer 2010 issue and published in September has become the early Autumn issue. Thomas Jefferson once said "Never put off until tomorrow what you can do today." When I was a US Navy officer we also had a metaphorical saying about excuses for being late: just like a part of our anatomy, everyone has excuses and they all stink. So, all I will say is better late than never, good things come to those who wait. We have four very interesting papers in this issue. Thanks for waiting. Growing up in Wisconsin, when we had a period in October with nice warm sunny weather we called it "Indian Summer." Welcome to the Indian Summer issue of 2010.

Leading this issue off is *Transaction Modeling for Performance Tuning and Capacity Planning* authored by John Meisenbacher and Karl Steger. This paper describes a straight forward and effective method for capacity planning and performance management which does not require expensive, commercial software. The authors discuss how the technique of correlating system resources with transaction rates can be applied to production and test systems with equal ease. They also describe how significant savings in development, testing, and production expenses have been obtained in multiple systems and environments.

Our second paper, *e-Governance for Local Government Administration in Nigeria: Benefits and Implementation Challenges* is from Akinnuwesi Boluwaji Ade, Ezike Joseph, and Shakirat Raji. This paper discusses how the use of information and communication technology leads to more effective planning, policy formulation, decision making and forecasting by local governments. The authors review how electronic governance (e-Governance) in local government administration (LGA) is carried out in Nigeria, along with its benefits as well as the implementation challenges.

Our third paper, *High Level Capability Assessment Aligned to Business Metrics* is from Dr. Abhijit S Ranjekar and Swati Dorge. This paper illustrates the methodology to have the capacity assessment done in terms of the business metrics giving various advantages over conventional capacity planning techniques. The simple yet powerful methodology yields accurate results and is applicable to a wide variety of applications – be it enterprise applications or services. The authors' techniques enable coupling the capacity planning decisions with the business growth and can quantify the risks in terms of financials relating to business. The new processes described also lead to superior alignment of capacity planning with the objectives as specified within ITIL V3.

Our fourth paper this issue is *Managing Processor Usage in a ClearPath MCP Metering Environment Using Multi-level Detailed Management Reports* from Wim Te Lintum. In order to proactively manage performance and capacity and to improve predictability of their mission critical ClearPath MCP mainframe systems, Senior IT Management of an end user was looking for easy accessible high level information about capacity and performance of these systems. At the same time there was a need for their Capacity Manager to zoom in on possible bottlenecks and to follow the results of tuning and optimization efforts in applications and processes. To fulfill these requirements Wim built a reporting tool for the end user with both a high, medium and detailed level of information on system usage and especially on processor usage, because that is the basic driver for metering costs and metering balance. Because of the different levels of information, the tool gives the end user the opportunity to zoom in (drill down) on usage per system, per brand, per workload type, per application and even per individual process for more detailed information on system usage.

Thanks to everyone who contributed to this CMG Journal. We plan on publishing one more issue (late Fall) this year, which would be four issues in 2010. We are always looking for good ideas for papers. Please consider writing a paper for the CMG Journal. You can submit your papers, as well as feeback to us at cmgjournal@cmg.org.

CMG 2010 is now only two months away. Have you registered yet? December in Orlando, Florida is always magical (thanks Mr. Disney) and CMG 2010 has a great lineup of speakers this year. Please look at http://www.cmg.org/conference/cmg2010/ for more information and to register.

Thanks again for reading, and we hope you enjoy this issue.

Stephen R. Guendert, PhD

Transaction Modeling for Performance Tuning and Capacity Planning

John Meisenbacher
john_meisenbacher@ieee.org

Karl Steger
karl_steger@sbcglobal.net

This paper describes a straightforward and effective method for capacity planning and performance management that does not require expensive commercial software. The technique of correlating system resources with transaction rates can be applied to production and test systems with equal ease. Significant savings in development, testing, and production expenses have been obtained in multiple systems and environments.

Introduction

Capacity planning and performance management are critical functions to properly size large scale transaction processing systems. An oversized system will waste money and reduce profits. An undersized system will result in poor customer service and harm the business. An effective solution to find the balance is available with data already collected.

Background

Large scale global online transaction processing systems are becoming common. Such systems support the needs of merchants, consumers, governments, and corporations throughout the world. Many of these systems require the highest quality and level of service 24 hours a day, 365 days a year and are expected to process transactions in the shortest time practical. This need for high quality, reliability, and availability with high transactions rates place intense performance demands on these systems. Service levels suffer if these systems are undersized. However, oversized systems wastes resources that would be better applied elsewhere.

Some companies rely on stress testing to determine the correct size for their system hardware. The stress test requires a test environment similar to production and involves sending test transactions near the maximum rate and then increasing the rate until the system breaks. The analogy is a racecar tested on the salt flats at full speed until the engine blows up. This technique is useful to understand the theoretical top speed. However, the results are not suited to forecast city and highway driving. Similarly, the company could spend time and money to maintain the test environment and develop the test transactions but not understand normal production performance. In addition, if the test transactions or the environment are not correctly configured then test results could be misleading.

Organizations use system monitoring software, such as Best/1®, to identify the peak hardware demand periods and plan their capacity on that load. However, these products generally are unaware of the underlying transaction model that drives the system. Even when this information is provided, the generic tool typically does not correlate the resource usage against the transactions or other workload metric.

A simple low cost technique to improve performance management decisions is available. The technique works in the test and production environments. It can be applied to a wide range of transaction systems even if the organization does not have access to a full scale performance test environment or dedicated capacity planning personnel.

The first part of this paper describes the technique in a production environment. A goal of this paper is to help organizations that are currently dealing with

® Best/1 is a registered trademark of BMC Software, Inc.

performance issues in their production systems. The paper also describes how a test environment can be more effectively used to forecast production performance.

Transaction processing cycle – the curse and the key

Transaction processing systems typically have daily, weekly, and annual transaction cycles. The peak time of the day on the peak day of the peak week of the year is a common target for capacity planning for such systems.

The transaction cycles ebb and flow in the day and over the year but the system must be sized to meet peak demands to assure high customer satisfaction. These systems tend to be underutilized during the lows in the cycle. Predicting these cycles and managing the systems to support the peaks is a major challenge. Fortunately, these cycles also provide the best way to understand the performance of the system.

These transaction cycles provide the information necessary to create an accurate performance model of how the system acts across the full workload range. Transactions per second (TPS) is a convenient measure for transaction processing system workload. A daily TPS cycle for a system can be seen in the graph in figure 1. The curves vary from day to day, but generally follow more or less a bell shape. In most systems this information can be obtained from the transaction logs in an extracted data warehouse. The production system database can be used directly during periods of low workloads to avoid impact on transaction processing.

All graphs presented in this paper are taken from production systems. The graphs have been generalized to remove sensitive information.

This paper focuses on CPU performance. However, the approach can be used to model network bandwidth, disk I/O, and other system resources that vary with workload and may constrain the system. Transaction data and workload data, i.e. CPU utilization, are collected for the same time periods and analyzed over several daily cycles.

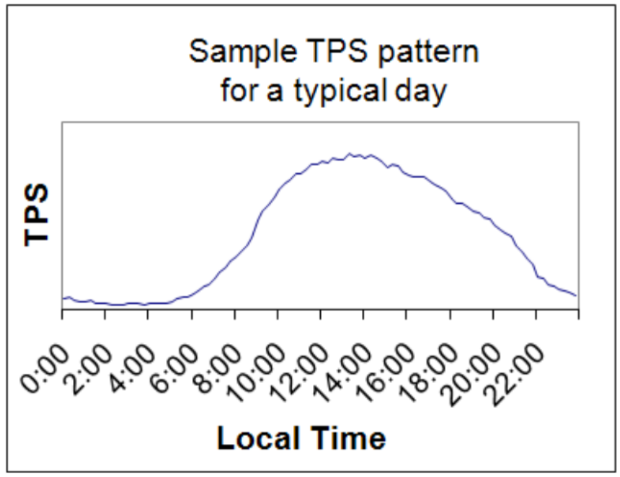

Figure 1. Daily transaction rates tend to follow a bell curve

CPU costs can be obtained from capacity planning tools, such as Best/1®, or from standard operating system utilities. This information is typically collected by the organization anyway, so the collection process may not add any additional burden on the system.

Most commercial tools tend to correlate CPU usage with time, as in the chart above. Figure 2 plots the previously diagramed TPS rates on the right Y-axis with CPU utilization on the left Y-axis.

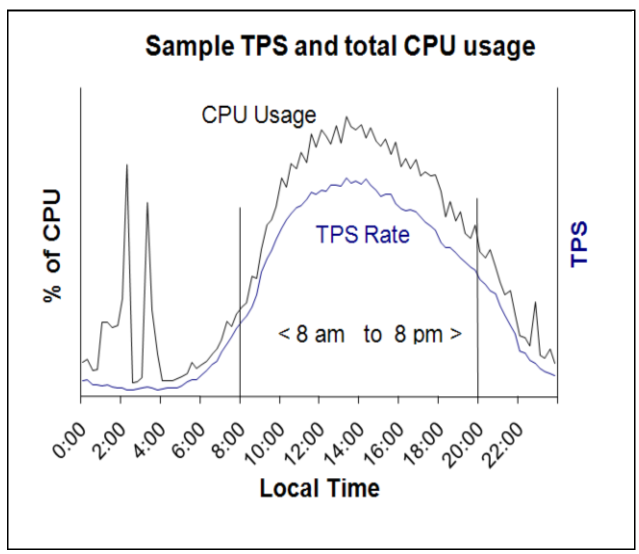

Figure 2. Transaction rates and CPU costs

The system performs batch functions during the early morning and evening, as seen by the processing spikes during those periods in the graph above. These periods tend to have low transaction rates so the additional CPU demands need not impact the transactions.

Figure 2 shows the ebb and flow of transactions and CPU utilization through the day. However, this graph does little to help understand and predict performance.

A correlation of TPS and CPU usage can be used to understand and predict performance. A preliminary step is to identify the data points to be used in the correlation. The graph above can identify the timeframe that includes the peak period but excludes periods where CPU is not dedicated to transaction processing. The time between 8 am to 8 pm in the graph above provides a good sample of TPS rates. This sample provides twelve hours, or forty-eight points if data is collected in fifteen minute periods. These data points of TPS and CPU within this sample can now be correlated. Multiple days can be merged to a single plot.

The transaction rates and processing costs are correlated and regression analysis is used to derive a model of the production system. Spreadsheet packages, such as Excel®, are generally sufficient for this analysis once the data has been collected.

Analysis

High level analysis

Transaction processing systems typically involve many different types of transactions each with multiple processing paths. High level analysis, using total TPS, ignores any differences in the transaction types and processing paths. Similarly, total CPU utilization is used in a high level analysis rather than the CPU usage of the various application components. This information is generally good enough to begin to better understand the system performance and predict capacity.

TPS rates are treated as the independent variable and plotted on the X-axis. CPU usage is treated as the dependent variable and plotted on the Y-axis. This can be done in a spreadsheet with an X and Y graph or in a statistical tool. Each point in the graph in figure 3 represents an observed value for TPS and CPU. The line in the diagram is the trend line generated by Excel®.

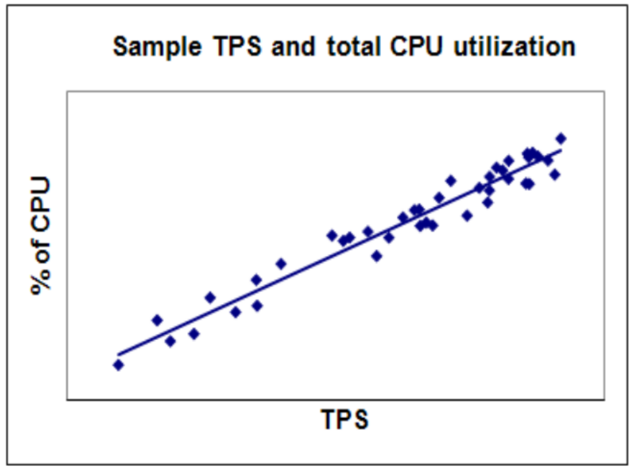

Figure 3. TPS and CPU trend analysis

The trend line may be linear, exponential, logarithmic, or a more complicated function. Experience has shown that fairly complicated transaction processing systems tend to have linear performance models. The slope of the line, in the graph above, represents the percent of CPU used per TPS. The fixed processing costs include processing that does not vary with transaction rate, and are represented by the intercept. The performance model in this case reverts to total CPU = (TPS * CPU per TPS) + fixed CPU costs, $y = ax + b$. The system should be observed in production or tested over a sufficiently wide TPS rate to verify the linear, logarithmic, or exponential nature of the correlation remains true throughout the range of the system.

During the analysis, it is also useful to determine the “accuracy” of the regression model. Various methods exist to determine the “variability” of the model’s calculated data set versus the observations. There are several methods to do this: the correlation coefficient, R-Square, sums of squares, and the Pearson product-moment correlation coefficient, to name but a few. MS-Excel provides two functions: correl() and rsq(). Note that these methods do NOT tell whether:

- The independent variable is a true cause of the changes in the dependent variable. They do not prove causality, only correlation.
- The correct regression was used.
- The most appropriate set of independent variables has been chosen.

Human understanding of the system is necessary to appropriately use the statistical approach.

® Excel is a registered trademark of Microsoft, Inc

Detailed analysis

A more detailed analysis can be performed either by tracking different processing costs and or different types of transactions. Detailed analysis of transaction types would identify the costs associated with each type of transaction.

Decomposition of total CPU is also useful. Typically, of the dozens of processes running on the system, the top ten processing activities can be correlated and the remaining tallied with “other costs”. Figure 4 on the following page was generated from a production system a few years ago.

The major activities of the system are database operations, transaction processing, format conversion, and communication. The database operations and transaction processing components have the highest transaction costs, slope. Format conversion costs more than communication processing, which is trivial from a CPU perspective.

Everything else, the “other processing”, appears to be moderately expensive but does not vary by TPS. If the X-axis covers the full TPS range up to annual peak periods then it would be a big factor in the total TPS capacity of the system. In such a case, it would be a good candidate to analyze and reduce. However, the graph above actually covers only the lower range of TPS capacity. “Other processing” might still be analyzed and reduced, but its impact on TPS capacity is overshadowed by database operations and transaction processing at higher TPS rates.

Identify and evaluate performance improvements

One use of this approach is to identify where to focus process improvement activities and to quantify the finished improvement. This activity was performed on the system modeled in figure 4 and database processing costs were targeted for improvement.

Modifications to the database subsystem were undertaken as part of the periodic system release. The activities of database storage and replication were reengineered. Figure 5, on the following page, shows the performance model after the system upgrade. The model clearly shows a significant drop in database costs, which are now below the cost line for format conversion at higher TPS rates. The quantitative nature of the model allows a clear measure of the change in system costs in production - the change in the slope of the cost line for database processing. The performance management activities can be clearly judged and evaluated in production against the defined performance goals before peak season.

The model can also be used to quantify the increased transaction capacity of the system. In this case, it was determined that the increased transaction capacity gained by reduced database and total system costs would enable the company to avoid an expensive system upgrade that would have otherwise been required to support the increased transaction demands for the next peak season.

Interval Selection

The interval, 15 minute in production, was the company standard collection interval for Best1 at the start of this analysis.

The 15 minute interval provided the necessary data points to model several transaction processing systems with sufficient precision to make useful performance and capacity planning decisions.

A longer interval in production, 30 minutes or longer, tends to smooth out and therefore hide performance anomalies. Models generated with such data may not provide sufficient precision to make effective performance improvement decisions.

A shorter interval, 1 to 5 minutes, provides a finer granularity model that may improve precision and accuracy of the analysis. This finer granularity would be warranted when there is a large variation in the minute by minute data.

In order to test the 15 minutes interval the minute by minute transaction rates were analyzed. The rates tended to vary within 5% of the 15 minute average for the systems reported in this article. It should be noted that the variation in the minute data tended to increase to 10% swings during the peak hour of the peak day in some systems.

Even though minute TPS data tracked well with the 15 minute average a Best1 collection change was made to collect at minute intervals. The finer granularity 1 minute model was comparable to the coarser 15 minute model. Therefore, a change in Best1 collection policies was not warranted. At this time this remains the standard, following the rule of thumb to keep the model as simple as possible.

As the large performance issues are tuned out of the system, a finer granularity model may be warranted if additional improvements are required.

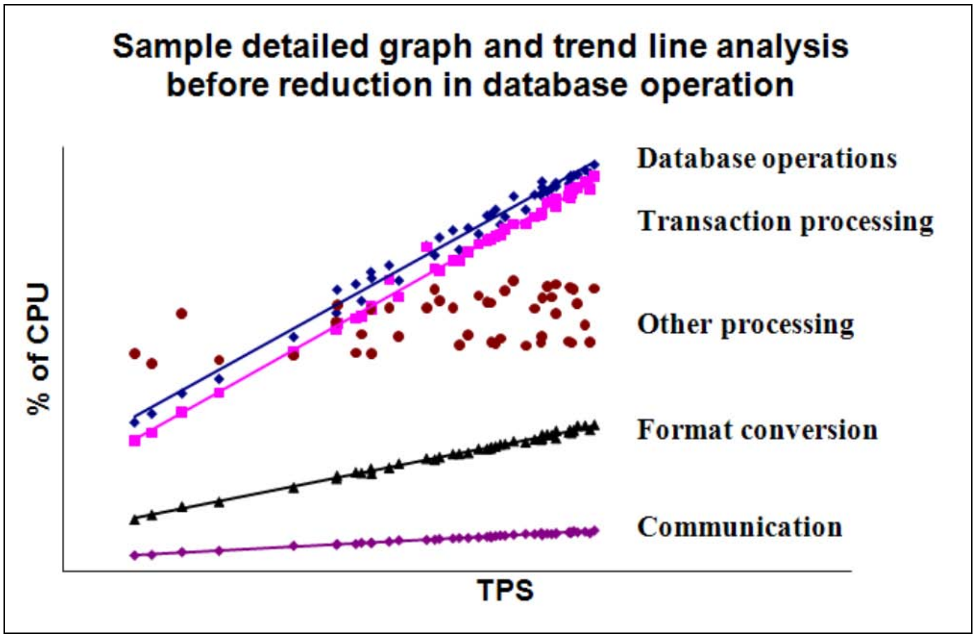

Figure 4. Detailed performance model with high transaction and database processing costs

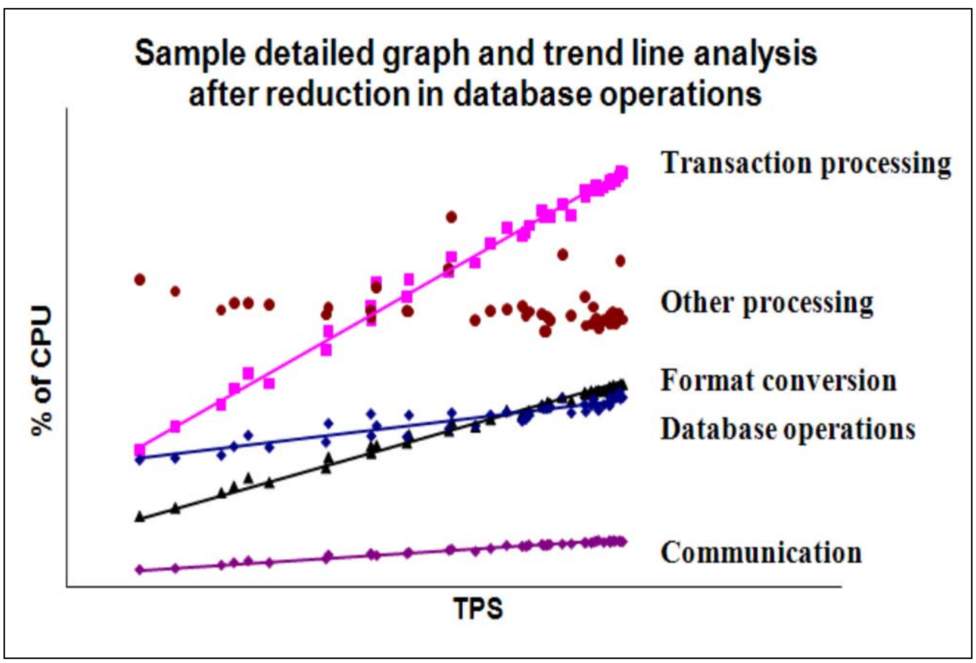

Figure 5. Detailed performance model showing reduced database processing costs

So, what interval should you use?

Basically, start with the data you already collect. If that is 1 minute data then by all means use the data to create your model. If you collect at the 10 or 15 or even 20 minute periods then create the model and determine if it is sufficiently precise for your analysis before changing to a shorter interval. If you collect at the 30 minute or longer intervals then you may need to consider a change in data collection to effectively model you system.

Trending the capacity of a system

Regression models can be used to determine the maximum capacity of a system. Solving the original equation [CPU = (TPS * CPU per TPS) + fixed CPU costs] for TPS, gives us the new equation [TPS = (CPU – fixed CPU costs) / CPU per TPS]. The maximum TPS rate for the system can be found by applying the maximum desired CPU.

For this system, the authors determined a maximum CPU utilization of 80% was optimal, which is in line with capacity planning policy. This provided sufficient excess capacity for the following:

- Inaccuracies in the business volume forecasts.
- Volatility of the TPS rates beyond the model's 15-minute interval.
- Reduced operational efficiencies at high CPU utilization. Testing revealed that this application's performance remained linear as the system neared 90% utilization. Above 90%, some systems broke down and TPS rates dropped significantly. Another way of looking at this – the knee of the response time curve was 90% for this application.

Additional operational monitoring overhead during peak times, and other anomalies also increase the risk of service issues above 80% CPU system utilization.

This 80% threshold may not be appropriate for all applications. For example, one of the authors generally uses 12% for web proxy servers, 50% for application servers, and 80% for database servers.

Figure 6 shows the CPU utilization per TPS for another system with the linear trend model as a solid line. Notice that the model does not encompass all of the observed data points. Production observations revealed that the outlying data points were an occasional change in the transaction mix and had to be included in the model. Transaction costs and rates vary dramatically for this system.

The authors used a more conservative approach to the model. By modifying the model to encompass all of the data points, a new trend appears as the dashed line. This modified trend produced a maximum TPS rate that was 7% lower than the unmodified model and was proved accurate by production observations.

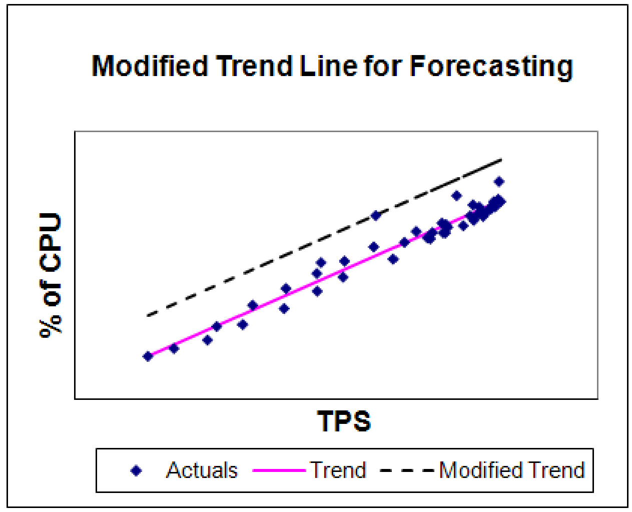

Figure 6. Modified trend line for forecasting

Comparison with test environment

Performance testing is critical during product evaluation and prior to the first deployment of the system. However, large scale performance testing can become less important for systems already deployed since a performance model can be generate days after system release into production.

Laboratory testing is useful to determine the shape of the correlation function: linear, logarithmic, exponential, or other before the system is first released into production. If production data samples of TPS and CPU usage across the full TPS range of the system are not available, then testing would also be worthwhile to determine the performance of the system near capacity. The production data at the low TPS range may appear linear, but could turn logarithmic or exponential at the high end. Peak season is the wrong time to be surprised with an exponential curve.

The test cases and test environment should be designed to be close enough to production to verify the shape of the correlation function. The stress test need not be so well designed as to provide an exact estimate for production slopes, if the system can be deployed and then sized and tuned as required prior to the peak period.

A realistically tuned test setup is useful to evaluate alternative design and implementation approaches during the software development cycle.

Once the system has been deployed in production, the technique allows the test setup to be evaluated against production. The model from test can be compared with the model from production to determine if there is a significant need to tune the test model.

Tuning the test model may involve changing the test system or the transaction test generator. This may require a significant investment. Since this is a quantitative mathematical model one may also simply apply a simple adjustment factor to the test model to conform to production and avoid additional development and costs associated with the test environment.

Stress testing periodic releases of the integrated system becomes less important with this technique is regularly used in production. In most cases, the current system and next release can be analyzed and compared in production far more accurately and at less cost than stress testing. Since such testing would be performed near the end of integration testing the time between a stress test of the integrated system and the performance model taken from the production system is typically short.

Release evaluation

As stated earlier, this technique can be used to track performance improvements from one release to another. Figure 7 depicts two releases for a transaction processing system. Each release has its own linear trend line. It is easy to see that significant performance improvements have been made – “release 2” has a smaller slope as TPS increases, meaning that for high TPS rates the amount of CPU used is less than the baseline. This indicates that the system’s capacity has significantly increased: more transactions in the same amount of CPU time and more unused CPU time.

This concept could easily be extended to include more releases, showing steady improvements as well as performance losses.

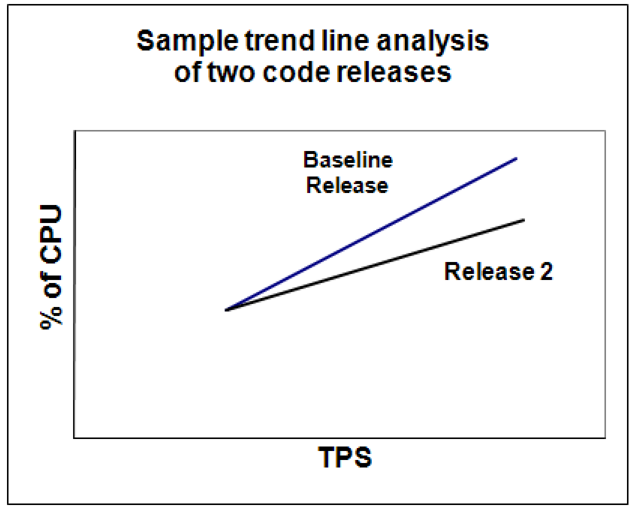

Figure 7. Performance improvements for different releases

Periodic performance monitoring

Production systems should be periodically reviewed. This approach is useful before and after each major release to evaluate the performance targets of the release. The feedback from production can be helpful in planning future development projects as well as final tuning before peak season.

Transaction shifts and “minor” changes that should not have an impact on performance can be quantified by periodically collecting production information and running the regression. Monitoring at this level could occur once a quarter and serve as a safety check.

Focused monitoring is also be used for systems that will be operating near capacity. One option for a system expected to have capacity issues in the coming peak period would be to add hardware, just to be sure. In some cases this would be a good business decision. In other cases, a better business decision would be to manage performance during the peak period and either delay the upgrade several months or wait for the next release to improve performance.

However, monitoring the system could also create problems. Increasing the number of operators and support personnel monitoring the system would increase the processing demands of the system. For example, the following scenario was observed during a critical processing time for a production system. A large number of support personnel were running certain operating system commands, such as TOP, from multiple terminals. This resulted in a majority of CPU time being used for monitoring versus transaction processing breaking our SLAs. A company can avoid this problem by using tools to collect processing data and moving that data to a location for display and analysis off the transaction processing system, either in near real-time or historically. Again, even with these tools, simply watching CPU rise and fall is of limited value for performance management.

Peak season monitoring example

A peak season example can be used to demonstrate the benefit of this model. A particular system had been developed and tested in the traditional performance environment and monitored in production by the capacity department. The technique presented in this article had not been previously used on this particular system.

The capacity planning department predicted that the system would be close to its capacity in the upcoming peak season. The development group was requested to insure proper system operation without additional hardware.

The following general plan was developed. Collect data from weekly peak transaction days, Friday and Saturday. Analyze the results on the following Monday. Investigate and resolve any issues by Thursday. Data from other days were collected and analyzed to provide additional information as needed.

A baseline performance regression model was generated prior to the last software release before peak season. The model suggested slightly more capacity compared to the commercial capacity modeling tool in general use at the time. In many cases, the commercial tools tend to "aim high" because they do not include transaction modeling, and therefore cannot precisely predict performance.

Still both models suggested that the system was going to be operating close to its capacity during peak season.

The decision was made to use the company standard commercial tool and test this approach in parallel.

During a minor release roll-out, the performance model identified unexpected performance degradation in production. The commercial monitoring system did not flag the error as the system was operating well below peak at this time. Two additional days were analyzed. Both days confirmed the performance drop.

Processing logs were more closely examined to prove or disprove the performance issue identified by the model. The culprit was identified as a script that was added late in the cycle. The impact of the script was amplified by security software running on the system.

This can be seen in figure 8 below. The flat line is the expected peak season capacity required. The solid line is the projected capacity by the model leading up to peak season. The points represent an observed data point. The first point was the baseline system and the first large dip below the line was due to the script discussed above.

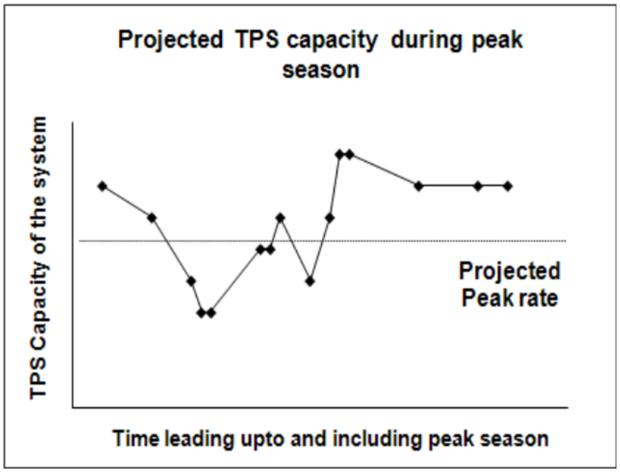

Figure 8. Changes in system performance monitored and managed

The script was adjusted to improve its performance. The next performance check verified that performance had been improved and capacity was back to expected levels. Observe the rise to just above expected peak requirements.

A later performance check identified another unexpected performance loss. Processing logs were again reviewed. In this case the problem was tracked to an operator console function that was abruptly terminated. The problem was caused when procedures to close a session were not followed. Again the commercial capacity planning tool did not flag the performance issue.

The issue was resolved by reinforcing proper procedures. A software fix was applied in the next system release to resolve the issue programmatically. By this time the model was the new "trusted" source for performance predictions.

A decision was made at this time to reduce the cost of some automated built-in-test and fault-isolation-test functions during peak season. As seen in the graph this further improved performance and provided additional reserve capacity during the peak period without significant loss of system reliability.

The final checks with the performance model of the system provided feedback that the changes had the desired effect and that the system would support peak demands.

Figure 8 shows the swings in projected capacity over several weeks as the system was adjusted and controlled. The graph also shows that projected capacity stabilized leading up to the peak period. A potential worrisome peak season was managed with relative calm by optimizing the application and system environment rather than adding hardware and increasing associated software license costs.

Conclusion

The approach described in this paper provides an effective method for capacity planning and performance management of large-scale transaction processing systems. The method is effective in both test and production systems. Its application to production systems reduces the need for complex performance testing in many cases. This technique can be used to augment commercial capacity planning tools or it can be used as the primary method for capacity planning and performance management.

Suggested Readings

Linwood Merritt, “A Finger in the Wind: Forecasting Techniques for Capacity Planning”, CMG Archive, (2004)

e-Governance for Local Government Administration in Nigeria: Benefits and Implementation Challenges

Akinnuwesi Boluwaji Ade (akinboluade@yahoo.com)
Ezike J.O.J. (josephezike@yahoo.com) and
Shakirat Oluwatosin Raji (tosinraj@yahoo.com)

Department of Information Technology, College of Information and Communications Technology, Bells University of Technology, Ota, Ogun State, Nigeria.

Abstract

The administration of a community has to do with the collection, storage and processing of data for effective planning, policy formulation, decision-making and forecasting. The continuous growth in national development calls for an effective and efficient tool to carry out proper administration with the view to ensuring good service delivery. Thus the interest in Information and Communication Technology (ICT) arose to assist in the administration of activities in both the government and private sectors of the economy. Today ICT products and services make up the fastest growing major industries in the world. ICT has become the information management tool of our age. It is currently expanding man's general capacity to manage the information available in the complex world. ICT has not only been used to structure, organize and access operational data of corporate organizations, but apparently it has been used to synthesize data in ways that sometimes create additional data or generate information. In this paper, a review of electronic governance (e-Governance) in Local Government Administration is carried out with the view of highlighting its benefits and implementation challenges in the Local Government Areas (LGA) in Nigeria.

Keywords: ICT, LGA, e-Governance, Administration

1.1 Introduction

The impetus for thinking about online dimensions to public sector operations came during the 1990s when the mainstream advent of the Internet began to translate into dramatic declines in the cost of both communicating and processing information. Consistent in large manner with the re-engineering movement of the preceding decade, public sector organizations sought new ways to control costs and improve organizational efficiencies. New and better approaches to managing information and the emergence of online channels of service-delivery promised significant financial savings [1][2][3].

In this era of global technological advancement, a new kind of rationalization has been introduced in the public sector by the use of modern ICT tools. Increasingly the use of ICT tools and applications is leading to transformational shifts in public policy, processes and functions. e-Governance is being deployed not only to provide citizen services but for public sector efficiency purposes, improving transparency and accountability in government functions and allowing for cost savings in government administration. ICT is changing the way the government does business for the people. In this context, e-Governance is seen to be a lever for the transformation of government.

Most governments around the world started their e-government initiatives with a focus on providing information and services to the citizen while service delivery platforms remained separate and parallel across various government agencies. In this case, service delivery was built around individual agency functions, structures, information, systems and capabilities. Figure 1.1 presents the evolving approach to public service delivery.

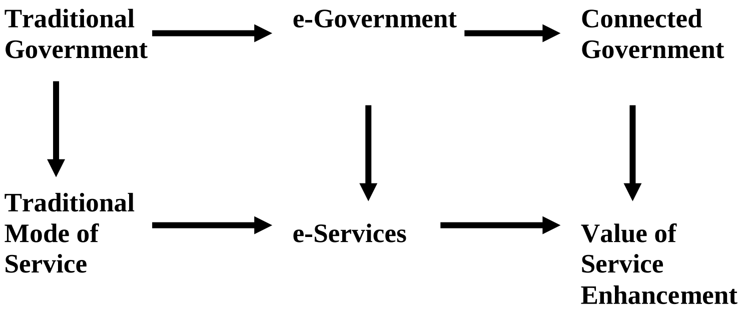

Figure 1.1 Evolving Approaches to Public Service Delivery

With the private sector leading the way, advances in accessibility and a greater use of technology have allowed an expansion of innovative ICT solutions. Now citizens and businesses around the world are increasingly demanding that their governments follow suit. Citizen groups have come to expect a 24/7 convenient communication with the government via a user friendly interface and a language that the user understands.

1.2 What is Governance and e-Governance?

The World Bank defines governance as:

The exercise of political authority and the use of institutional resources to manage society's problems and affairs [7].

The Worldwide Governance Indicators project of the World Bank defines governance as:

The traditions and institutions by which authority in a country is exercised [8]. This considers the process by which governments are selected, monitored and replaced; the capacity of the government to effectively formulate and implement sound policies and respect the views of citizens.

An alternate definition sees governance as:

The use of institutions, structures of authority and even collaboration to allocate resources and coordinate or control activity in society or the economy [9].

Therefore a good government, following these definitions, could consist of a set of inter-related positions exercising coercive power that assures, on behalf of those governed, a worthwhile pattern of good results while avoiding an undesirable pattern of bad circumstances, by making decisions that define expectations, grant power, and verify performance.

E-Governance is the use of ICT by different actors of the society with the aim to improve their access to information and to build their capacities. It is the public sector's use of ICT tools with the aim of improving information and service delivery, encouraging citizen participation in the decision-making process and making government more accountable, transparent and effective.

E-governance is more than just a government website on the Internet. It is a form of e-business in governance and refers to the processes and structures pertinent to the delivery of electronic services to the public (citizens and businesses), collaborating with business partners, and conducting electronic transactions within an organizational entity. That is the application of electronic means in: *the interaction between government and citizens and government and businesses, as well as in internal government operations.*

1.3 Objectives of e-Governance

Backus [12] highlighted the following objectives of e-Governance:

a. To support and simplify governance for all parties (government, citizens and businesses) with the view of connecting all the three parties and stimulating good governance.
b. To provide citizens access to information and knowledge about the political process, services and choices available.
c. To enable the transition from passive information access to active citizen participation by: *informing the citizen, representing the citizen, encouraging the citizen to vote, consulting the citizen and involving the citizen.*
d. To fulfill the needs of the public and expectations satisfactory on the front-office side, by simplifying the interaction with various on-line services.
e. To facilitate speedy, transparent, accountable, efficient and effective interaction with the public, citizens, businesses and other agencies.
f. To facilitate speedy, transparent, accountable, efficient and effective process for performing government administration activities in the back-office.

1.4 Delivery Models of e-Governance

The primary delivery models of e-Governance can be divided into:

a. Government-to-Citizen or Government-to-Customer (G2C)
b. Government-to-Business (G2B)
c. Government-to-Government (G2G)
d. Government-to-Employees (G2E)

Figure 1.2 presents G2C, G2B, G2G and G2E interactions. Figure 1.3 presents the model for the interactions

Government-to-Citizen (G2C) is the online interaction between government (local, state and federal government) and private individuals. For example, government sectors become visibly open to the public domain via a Web Portal, thus making public services and information accessible to all.

Government-to-Business (G2B) is the online interaction between government (local, state and federal government) and the commercial business sector. For example, http://www.dti.gov.uk/ is a government web site where businesses can get information and advice on e-business 'best practices'. http://g2b.perm.ru/ is another example.

Government-to-Government (G2G) is the online interaction between government organizations, departments, and authorities and other government organizations, departments, and authorities. Its use is common in the UK. G2G systems generally come in one of two types: ***Internal facing*** - joining up a single Governments departments, agencies, organizations and authorities and - ***External facing*** - joining up multiple Governments.

Government-to-Employees (G2E) is the on-line interaction between government (local, state and federal government) and the civil servants (government employees).

1.5 Three Perceptions of e-Governance

Citizens' Perception:- Citizens increasingly expect governments to perform effectively like private entities. They want convenient and instant access to public services 24/7. They want to access public services from home, work or remote geographical location. Citizens do not want any limitation on how they can access services.

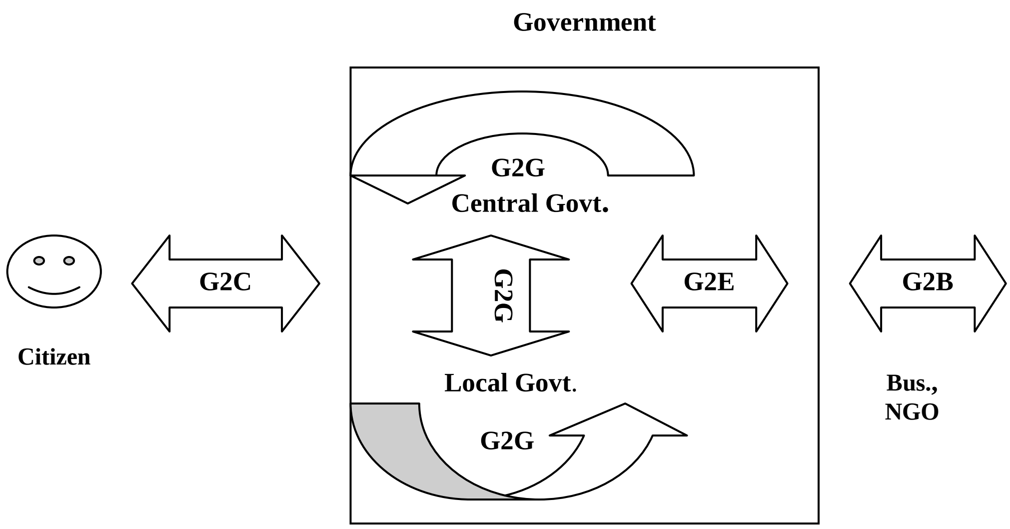

Figure 1.2 G2C, G2G, G2E, G2B Interactions [12]

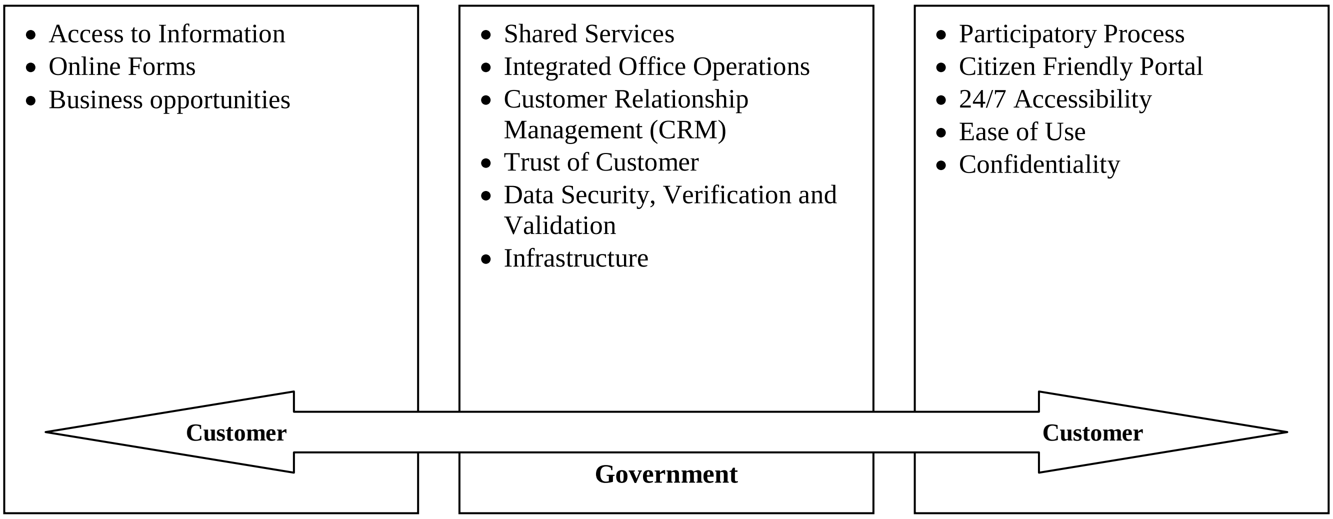

Figure 1.3 e-Governance Interaction Model [11]

Business Perception- is the expectation of the private companies to interact well with the government and thus minimize the physical bureaucracy and cost. Private companies want to conduct business transactions with government online. Thus, instead of appearing physically in government offices to complete paper forms, a contractor will find it easier to bid for a contract on-line from any remote geographical location.

Government Perception:- Government is expected to fulfill the needs of the public and facilitate speedy, transparent, accountable, efficient and effective process for performing government administration. Government is meant to provide easier public access to services, increase service volume and provide necessary infrastructures for its employees in order to render better service to the public.

1.6 Stake Holders in e-Governance

a. ***Politicians:-***–They enact / legislate a law. They are the suppliers of the e-Government systems

b. ***Public Administrator:-*** They define the process for realizing a law.

c. ***Programmers:-*** They design and implement the e-System for realizing the law.

d. ***End-Users:-*** They use the e-Government services. They are the customers to the government.

1.7 The Four Phases of e-Governance

Figure 1.4 presents the developing phases of e-Governance

a. In the *first phase*, e-Governance means being present on the web, providing the public (G2C & G2B) with relevant information. The format of the early government websites was similar to that of a brochure or leaflet. The value to the public is that government information is publicly accessible; processes are described and become more transparent, which improves democracy and service. Internally (G2G) the government can also disseminate static information with electronic means, such as the Internet.

b. In the *second phase*, the interaction between government and the public (G2C & G2B) is stimulated with various applications. People can ask questions via e-mail, use search engines, and download forms and documents. These save time. In fact the complete intake of (simple) applications can be done online 24 hours per day.

c. With *phase three*, the complexity of the technology is increasing, but customer (G2C & G2B) value is also higher. Complete transactions can be done without going to an office. Examples of online services are filing income tax, filing property tax, extending/renewal of licenses, visa and passports applications and online voting. Phase three is made complex because of security and personalization issues. For example digital (electronic) signatures will be necessary to enable legal transfer of services. On the business side, the government is starting with e-procurement applications.

In this phase, internal (G2G) processes have to be redesigned to provide good service. Government needs new laws and legislation to enable paperless transactions.

d. The *fourth phase* is when all information systems are integrated and the public can get G2C & G2B services at one (virtual) counter. Government employees in different departments have to work together in a smooth and seamless way. In this phase cost savings, efficiency and customer satisfaction are reaching highest possible levels.

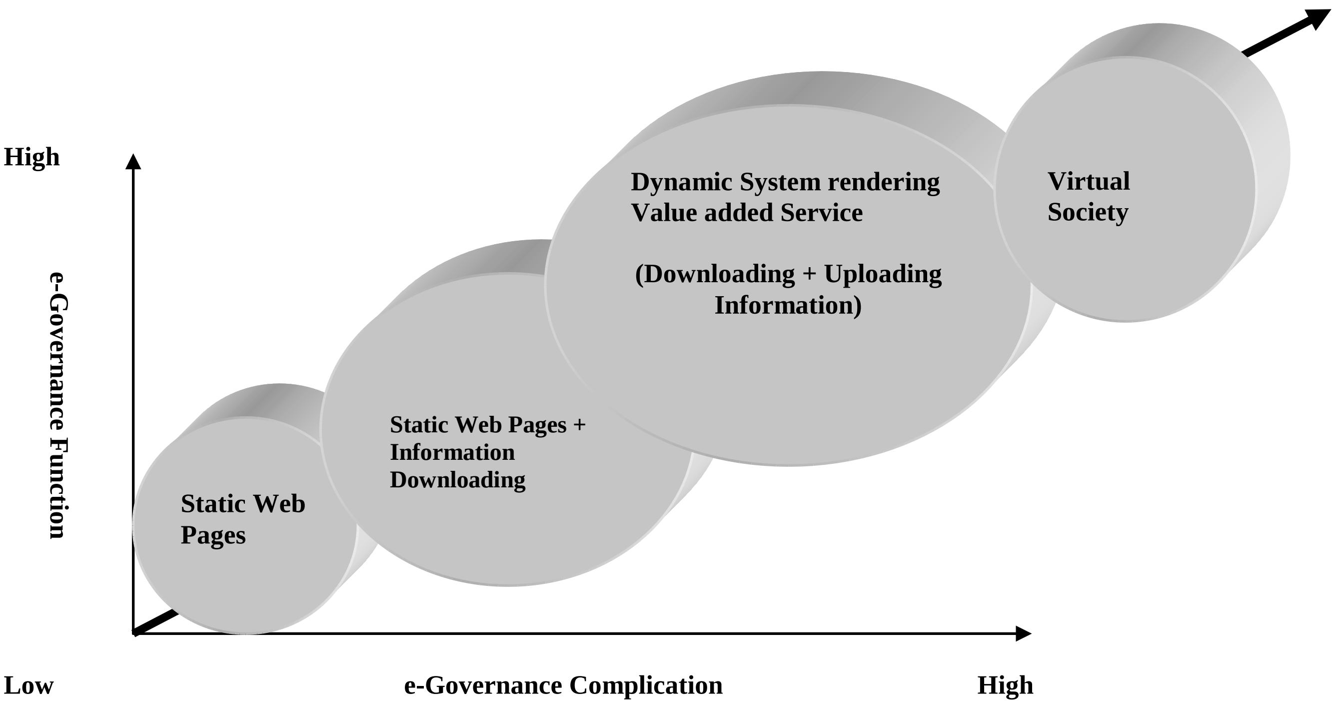

Figure 1.4 The Developing Phases of e-Governance [11][12]

The conceptual diagram of e-Governance system is presented in Figure 1.5.

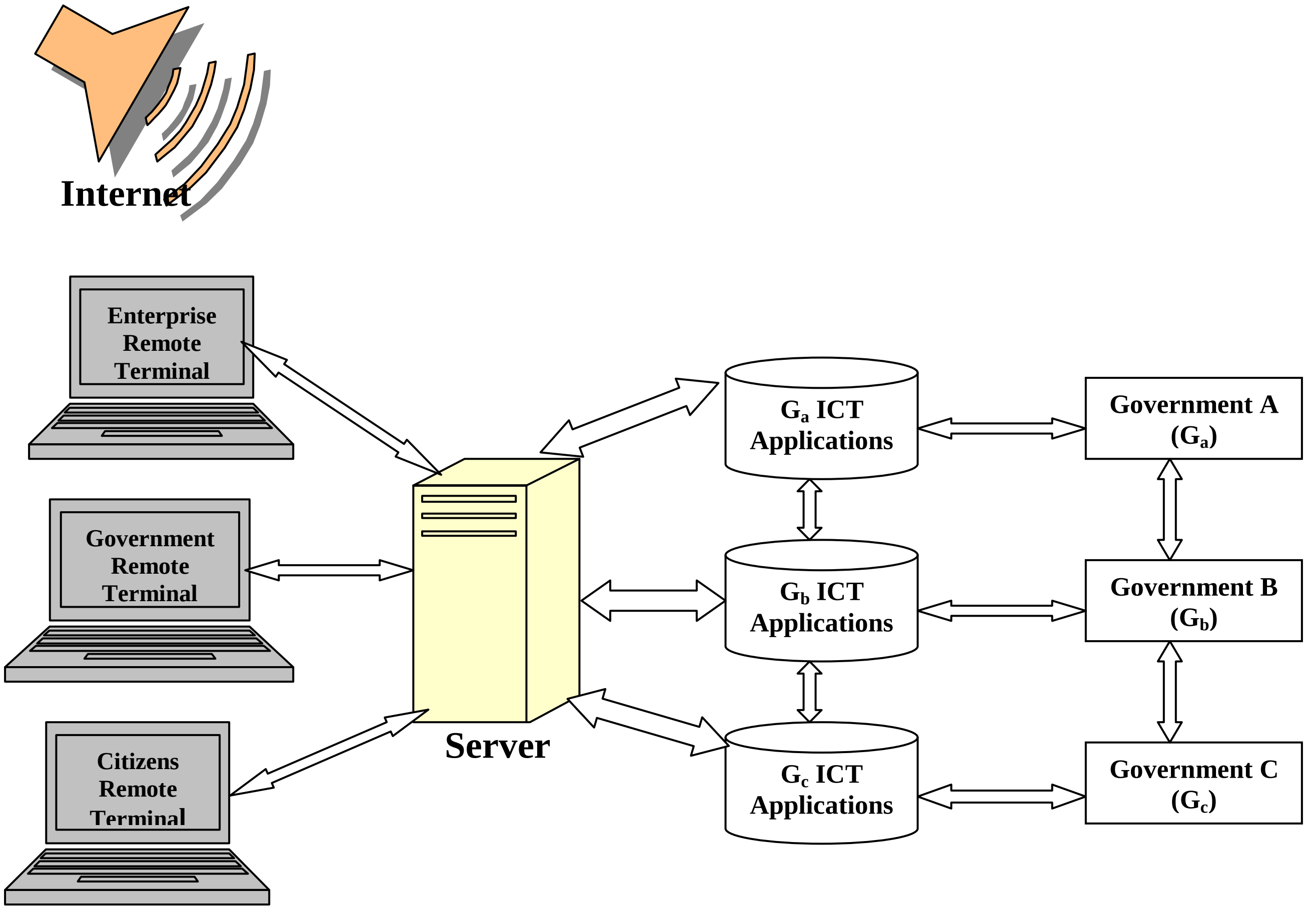

Figure 1.5 Conceptual Diagram of e-Governance System

2.1 Accounts of e-Governance Practices Nationally and Internationally

In [5], an account of e-Governance implementation in Lagos State University Nigeria (LASU) was given. LASU was ranked, within the first three leading universities, out of over 80 tertiary institutions in Nigeria. This was achieved because LASU Administration recognized the benefits of e-Governance.

Using questionnaires, interviews and observations, [6] did assessment of e-Governance resource use in South-Western Nigeria. The perspectives of the government employees (GEs) and non-government employees (NGEs) were assessed with the view of establishing their level of awareness of e-Governance, computer literacy level, access to the Internet and proficiency about the use of e-Governance resources. It was found that 68% of both GEs and NGEs are aware of e-Governance in the states, and the governments have achieved this high rate of awareness by means of mass media. Computer literacy amongst employees is put at 19 users out of every 20 (that is 95%), and 52% of those that are computer literate have over 3yrs experience in using the computer. Also, 76.6% are reported to have access to the Internet and only 31.3% of those that have access to the Internet have access points in their offices. It was also found that about 50% of the users are proficient and about 35% of them use the Internet on a daily basis.

[10] set out to examine how the public-private partnership being implemented in the pension administration system has helped to promote the culture of electronic public service delivery in the country. It was found that the ICT-based system has eliminated resource wastages, fraud, corruption, rent seeking and even loss of life

associated with the traditional system. It also provides opportunities for the private sector, especially banks and investment houses, by creating wholly new financial services providers– Pension Fund Administrators (PFAs) and Pension Fund Custodians (PFCs), with attendant growth in that industry and multiplier effects on the larger economy.

Internationally, e-Governance is well embraced. This is evident in the UN e-Government survey carried out in 2008 [11]. The United Nations conducts an annual e-Government survey which includes a section titled *e-Government Readiness*. It is a comparative ranking of the countries of the world according to two primary indicators: (i) the state of e-government readiness; and (ii) the extent of e-participation. Constructing a model for the measurement of digitized services, the survey assesses the 191 member states of the UN according to a quantitative composite index of e-government readiness based on website assessment; telecommunication infrastructure and human resource endowment. The results of the survey indicate that governments are moving forward in e-Governance development around the world.

It is worth noting that in the 2008 Survey, there are no countries in the top 35 from the African, Caribbean, Central American, Central Asian, South American and Southern Asian regions. The high cost of deploying a robust infrastructure capable of handling e-Governance applications is one reason for this discrepancy. In addition, many developing countries have been unable to fully implement their e-Government policies, mainly due to other competing pressing social issues that need to be dealt with in the context of tight budget constraints, such as: health, education and employment, to name a few. In the breakdown of the 2008 e-Government readiness index, Sweden was ranked number 1 with e-Governance readiness index of 0.9157. South Africa was ranked 61 with e-Governance readiness index of 0.5115. Nigeria was ranked 136 with e-Government readiness index of 0.3063.

It is obvious from the few accounts of e-Governance, that Nigerian government is yet to make e-Governance a priority and commence full implementation despites all the benefits associated with it. The traditional approach to governance is still in vogue in the 3 tiers of government. Few ministries and government parastatals that embrace e-Governance are still in the first phase of e-Governance and struggling to migrate to the second phase.

2.2 e-Governance Applications

In Nigeria today, e-passport, e-drivers license, e-vehicle registration, e-vehicle license production and renewal, e-tax payment, e-company registration, e-custom duty payment, e-job recruitment are some of the e-Government applications available.

3.1 Benefits of e-Governance Implementation in Local Government Areas (LGAs)

The benefits of e-Governance can be classified as internal and external benefits. Figure 3.1 summarizes these benefits. In addition, e-Governance brings urban and rural communities together and breaks the barrier of distance and thus leads to efficient administration.

3.2 Factors of a Successful e-Governance Implementation in LGAs

The following should be in place for a successful implementation of e-Governance:

a. Robust and cordial relationship among the intuitions and government. This will ensure good interconnectivity among the institutions and the government.
b. Development of an effective and friendly legal framework with the view of propounding laws and formulating policies that will serve as firewalls for the stakeholders.
c. Long-term investment in ICT infrastructure which is the platform for e-Governance to operate.
d. Political stability.
e. High level of trust in government.
f. Good economic structure.
g. Good government structure (centralized or decentralized).

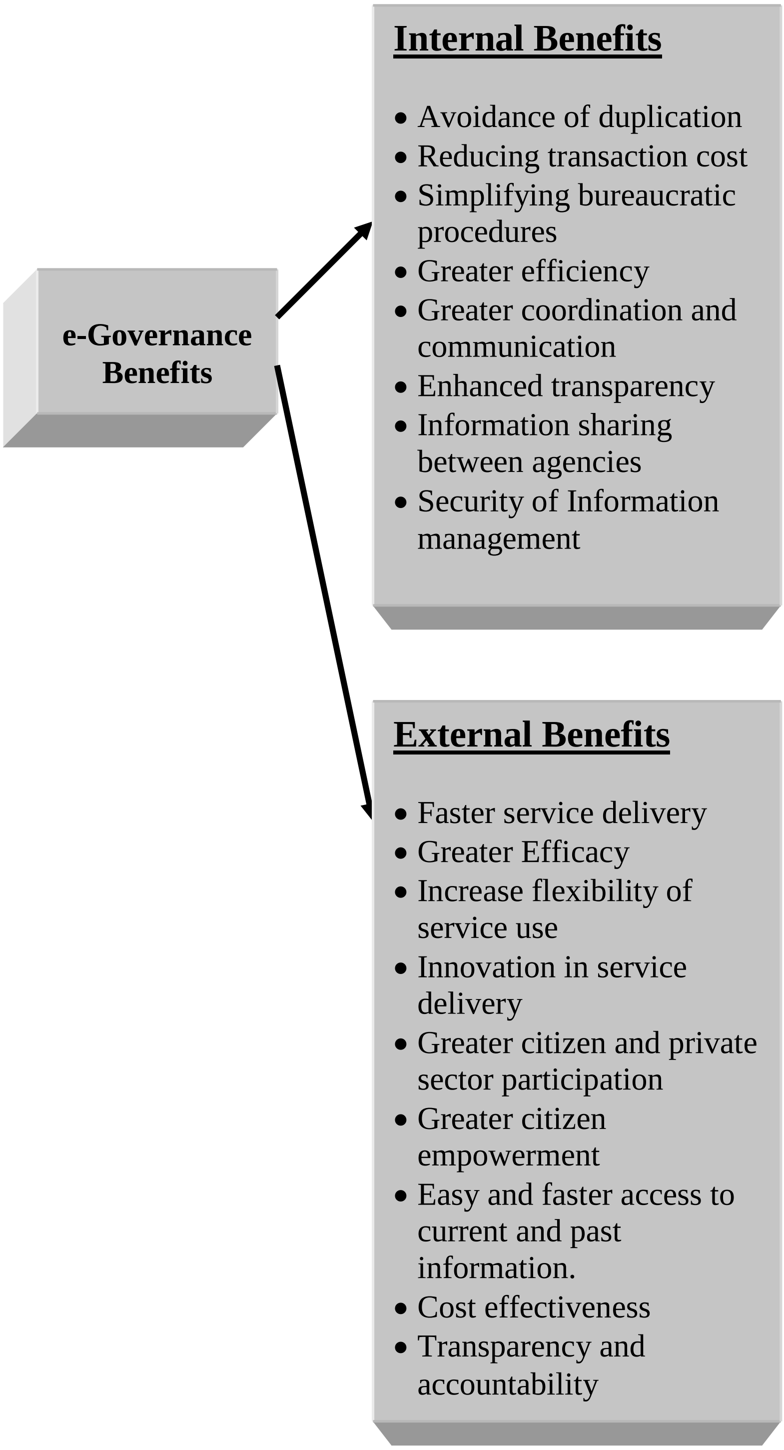

Figure 3.1 Benefits of e-Governance.

4. SWOT Analyses of e-Governance in LGAs

Four SWOT analyses are presented and focus is on political, social, economic and technological sectors.

Political Sector	
Strength	***Weakness***
❖ Existence of Strategies and policies ❖ Provision of laws and legislation	❖ Lack of cyber laws. ❖ Slow decision making process. ❖ Short term approach due to elections. ❖ Lack of project continuity as regime changes
Opportunities	***Threat***
❖ Raise government standard ❖ Transparency in governance ❖ Raise internal and external interests in community development ❖ Exposes community to national and international political growth	❖ Bureaucracy ❖ Corruption among the politicians ❖ Lack of total transparency ❖ Political instability ❖ Government resistance to ICT development

Social Sector	
Strength	***Weakness***
❖ People eager to learn IT skills ❖ Enhances the social relationship of people electronically ❖ Tourism enhancement	❖ Poor basic education ❖ Low level of IT literacy ❖ High competition with private sector ❖ Language barrier ❖ Problem of general public acceptance of ICT deployment
Opportunities	***Threat***
❖ Better education system ❖ Job availability ❖ Employment increases ❖ Open communities up for tourist attractions	❖ Resistance and hostility among stakeholders ❖ Loss of IT skilled personnel after training ❖ Presence of IT digital divide

❖ Promotion of cultural and social events ❖ Provision of linkages socially	❖ Adverse cultural influences ❖ Insecurity of life and properties in communities ❖ Abuse of use by youth that may concentrate more on social events than other aspect of growth

Economic Sector	
Strength	***Weakness***
❖ Enhance revenue collection ❖ Provision of tranparency in business ❖ Enhance eceonomic relationship between government and private sector ❖ Enhance accurate keeping and processing of financial records ❖ Reducing cost of Internet access	❖ Insufficient fund for ICT deployment ❖ Poor budget control ❖ Inadequate linkage and partnership of local government with the industries ❖ Poor investment in indigenous software development
Opportunities	***Threat***
❖ Availability of e-transaction ❖ Development and sales of software customized for local government services ❖ Business re-engineering using ICT	❖ Increasing siphoning of public fund ❖ Global economic meltdown ❖ e-stealing using cloned credit cards ❖ Stakeholders losing confidence in e-government technology

Technological Sector	
Strength	***Weakness***
❖ Expansion of Internet technology ❖ Shipment of telecommunication devices at minimum cost ❖ Availability and affordability of telecommunication devices ❖ Plug and play features of telecommunication for easy installation ❖ Enforcement and use of common IT standard	❖ Shortage of IT skills ❖ High cost of bandwidth ❖ Heterogeneous presentation and representation of data ❖ Fluctuation in IT standard ❖ Lack of standard software ❖ Inadequate performance of Internet Service Providers ❖ Inadequate performance of GSM operators ❖ Insufficient telecommunication infrastructure
Opportunities	***Threat***
❖ Easy networking of systems ❖ Increasing access to the web ❖ Increasing acquisition of telecommunication devices by stakeholders ❖ Increasing global use of e-government technology for service delivery	❖ Technology failure ❖ Piracy ❖ Presence of hackers and crackers ❖ Inadequate firewalls ❖ Poor supply of electricity

4.2 Challenges of e-Governance Implementation in LGA

a. Poor acceptance of e-Governance by the personnel of the Local Government Areas. Some of the personnel feel threatened with the computerization of their services. Instead of accepting ICT tools and applications as artificial partners, they are afraid that it is a way of relieving them of their jobs. As a result, they fail to fully support e-governance and prefer using the conventional manual/mechanical data management method.

b. The complexity of e-Governance is due to the complexity of administrative management. The purpose of future e-Government is to realize “one-stop online service” without time and space limits, which needs the interconnection between the departments to deal with public service affairs. However, management affair in each department is a relatively separated system with quite different affairs. So it’s a difficult and complex problem for the various departments to realize interconnection between them. Uniform programming and standard is vital to the development of e-government.

c. The low informatization level in the whole society and the slow development of electronic commerce restrict to some extent the development of e-governance in LGAs. Government informatization is closely related with enterprise, society and individual informatization. The current situation in our LGAs is that not only the informatization level in the whole society is low, but the development of e-commerce in enterprises is also slow, which necessarily restrict the development of e-Governance.

d. Security has become the key problem in government informatization, which influences greatly on the development of e-governance. Compared with e-commerce, the e-Governance has higher demands for the security of information. Therefore security has become the top issue in developing e-Government.

e. The capability in research and development of information technology in LGAs is relatively weak, which challenges to some extent the development of e-government.

f. The legislation in Nigerian e-Government lags relatively behind the developed countries, which to a great extent affects the development of e-government. This is evident in [11] where Nigeria ranked 136 among 191 UN member countries assessed for e-Government implementation. Many developed countries have already established a series of regulations, laws to improve their e-Government. Electronic signature is allowed, e-payment is acceptable legally, and some network security policies are published. There is need for laws and regulations related with e-Government to guide the electronic transactions and e-payment, and to protect the safety of databases.

5. Conclusion

Imagine a situation in which all interaction with government can be done through one virtual counter 24 hours a day, 7 days a week, without waiting in lines. This will be possible if governments are willing to decentralize responsibilities and processes, and if they start to use electronic means such as the Internet.

All in all, implementing e-Government is a systematic re-engineering. It cannot be achieved only by drafting a law or issuing an order from political leaders. It requires changing how officials think and act, how they view their jobs, how they share information between departments, with businesses and with citizens. It requires re-engineering the government’s business processes, both within individual agencies and across governments. A series of reform should be conducted to the current governmental system in order to create an ideal environment for developing e-government in all LGAs in Nigeria.

6. Reference

1. Nelson, M.R. (1998). Government and governance in the Networked World.

2. Lowy A. and Ticoll D (1998). *Blueprint to the Digital Economy: Creating Wealth in the Era of e-Business.* New York*:* McGraw-Hill 274-298.

3. Heeks, R. (1999). Reinventing Government in the Information Age – International Practice in IT-enabled Public Sector Reform, Routledge, London.

4. Delan Bu (2005). The Application Study of Chinese e-Government Construction. In proceedings of ICEC’05, August 15-17, Xi’an, China. Copyright 2005 ACM 1-59593-112-0/05/08. Pg 836 – 838.

5. Falabi .O. Jide (2007). The Benefits and Challenges of E-Governance Implementation in a Third World Country: Lagos State University As a Case Study. In proceedings of ICEGOV 2007, December 10-13, Macao. Copyright 2007 ACM 978-1-59593-822 -0/07/12

6. Awoleye Michael, Adeniran Oluwaranti, Willie Siyanbola and Adegunodo Rotimi

(2008). Assessment of e-Governance Resource Use in South-Western Nigeria. In proceedings of ICEGOV2008, December 1-4, Cairo, Egypt. Copyright 2008 ACM 978-1-60558-386-0/08/12.Pp 154 – 159.

7. World Bank, Managing Development - The Governance Dimension, 1991, Washington D.C .Available @ http://www-wds.worldbank.org/external/default/WDSContentServer/WDSP/IB/2006/03/07/000090341_20060307104630/Rendered/PDF/34899.pdf
8. A Decade of Measuring the Quality of Governance. Available @ http://info.worldbank.org.governance/wgi2007/pdf/booklet_decade_of measuring_governance.pdf
9. Bell Stephen (2002). Economic Governance and Institutional Dynamics, Oxford University Press, Melbourne, Australia.
10. Faniran o. Sanjo and Olaniyan M. Kayode (2007). Promoting E-Governance through Public-Private Partnership - A Study of Nigeria's Pension Administration. In proceedings of ICEGOV2007, December 10-13, Macao. Copyright 2007 ACM 978-1-59593-822 -0/07/12. Pp 301 – 304.
11. United Nations e-Government Survey 2008: From e-Government to Connected Government. Available @ http://unpan1.un.org/intradoc/groups/public/documents/UN/UNPAN028607.pdf
12. Michiel Backus (2001). E-governance in Developing Countries. IICD Research Brief - No 1.
13. Baoling Li (2005). On the Barriers to the Development of e-Government in China. In proceedings of ICEC'05, August 15–17, Xi'an, China. Copyright 2005 ACM 1-59593-112-0/05/08. Pp 549 - 552

High Level Capability Assessment Aligned to Business Metrics

Dr. Abhijit S Ranjekar **(abhijit.ranjekar@gmail.com)**

Swati Dorge **(dorge.swati@gmail.com)**

This paper illustrates the methodology to have the capacity assessment done in terms of the business metrics giving various advantages over conventional capacity planning techniques. The simple yet powerful methodology yields accurate results and is applicable to a wide variety of applications – be it enterprise applications or services. The techniques enable coupling the capacity planning decisions with the business growth and can quantify the risks in terms of financials relating to business. The new processes described also lead to superior alignment of capacity planning with the objectives as specified within ITIL V3.

1. Introduction

Capacity planning is a practice which has been in place for many years and has seen many changes over this time. . The conventional techniques employ the creation of non-functional application model(s) and evaluating the CPU-memory requirements. This requires the mapping of the application process flows -their execution steps and the CPU-memory associated with the steps into the model. Many tools based on this methodology are already available and on the market. A major drawback of these tools 'methodology is the lack of accuracy of the outcome. The process (es) and the computations are not the culprits – it's the accuracy of the inputs that plays the role of the spoil sport. The outcome is as accurate as the inputs provided. The fundamental building block of the methodology is the CPU-memory associated with the execution of a sub-step within the process flow. The efforts involved in extracting the metrics and the accuracy of the final results are disproportionate. Hence the need to provide an alternative where the target is achieved by overcoming the known drawbacks.
The process-methodology further described is based on actual production scenarios and has been proven across a variety of applications. The authors have derived the process based on actual implementations which have established the accuracy of the methodology by comparing the actual with the forecasted results.

2. Capability assessment framework

The diagram below briefly elaborates the framework for a high level capability assessment aligned to business metrics:

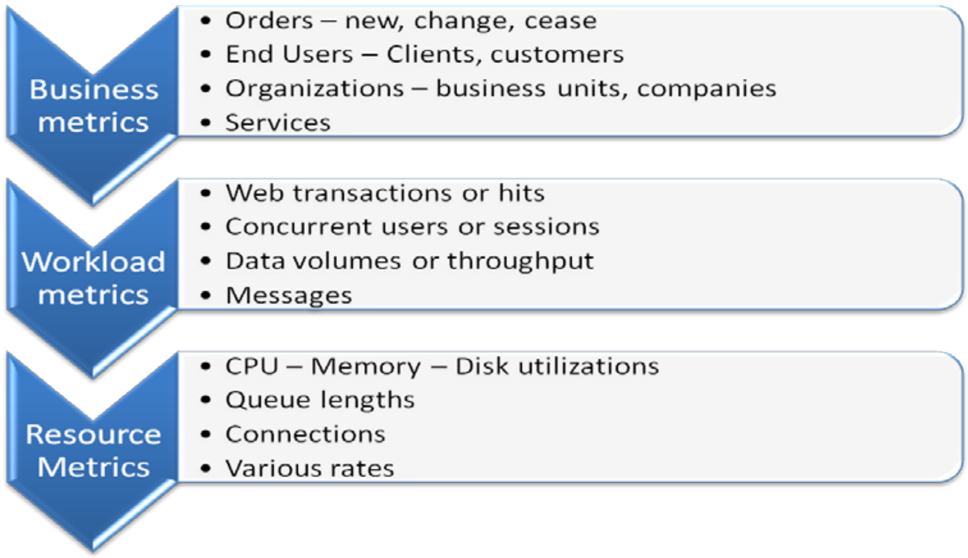

The approach works on 3 levels as illustrated above and involves the following:

- Fixation of the relevant metrics applicable to the given system under consideration
- Choice of a suitable time interval bucket – typically a 5 minute period is well suited for many applications
- Selected metrics are then either summed over or averaged over the intervals selected
- Analysis of the metrics - behavior profiles of the metrics shed light on the way the load comes into the system and how the system responds to it.
- The metrics are then correlated – extrapolated and capacity planning is done through the 3 layers giving the ability to inter-convert between the layers.
- The conversion into business metrics allows the stake holders from the business management to take appropriate decisions and facilitates the capacity planning as aimed in ITIL V3.

The framework is based on the principle as mapped in the diagram below.

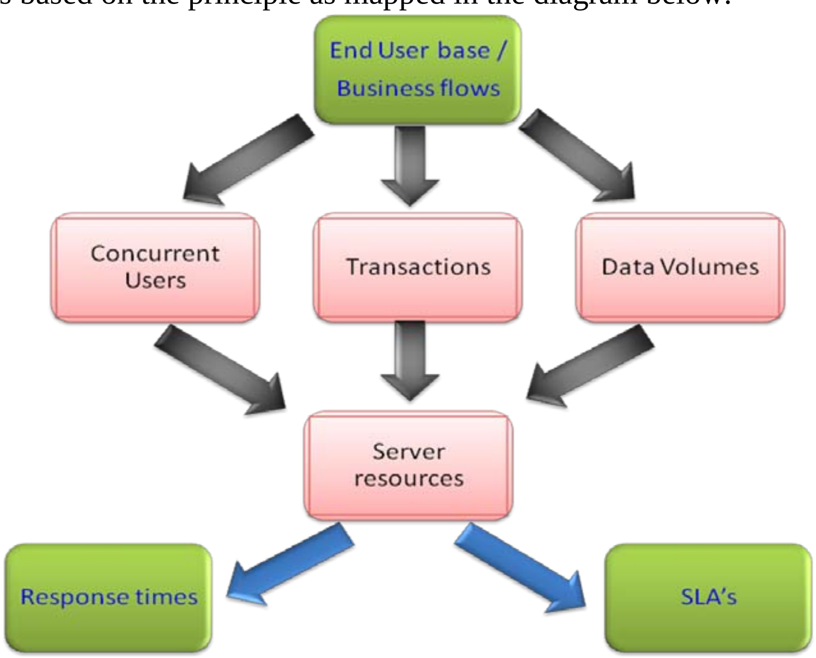

The business metrics cascade into various workload metrics which in turn consume the resources. The end user responses are driven by the resource utilizations-performance of the service on the given hardware-platform. Thus interlinking the 3 tiers-entities through the associated metrics achieves weaving of the application-service being delivered through the platform-architecture.
The approach is elaborated through the case study below which is one of the many successful implementations of the framework. The case can be regarded as a typical one showcasing the business growth – the current one being growth in terms of end users.

3. Case Study: Framework metrics

This section expands the methodology to one of the services being hosted-provided within the managed service offering to a major Health service organization. The application is used by professionals associated with the Health sector – doctors, nurses, admin staff etc. Service is rolled out to client organizations viz. hospitals, medical centers, individuals etc . This leads to addition of large groups of users consuming the hosted services. The services are bound to be offered with stringent SLA's attracting hefty penalties for violation, thus requiring very careful performance and capacity considerations. The application-service is hosted on a multi-tier architecture – with 2 major components crucial for capacity as below:

3.1. Framework Metrics: Business

Capacity needs to be aligned with the business planned roll-outs. The first level of the framework viz. business metrics is fixed to be the end user base. The charts below present the business aspects of the service under consideration. Load on the Service-application increases as the new roll-outs take place according to the time phased schedule. The charts show the plan for addition and the cumulative user base consuming the service. The X-axis shows the timelines along with the roll-out organizations while the Y-axis shows the user counts : individuals in the first graph while cumulative on the second graph. All users, even though from different organizations, would use the same service-application with each organization having an independent database. The resources are shared across organizations

– users in order to optimize the hardware. Thus the cumulative users are the driver for the business growth and the choice for the business metric in the framework.

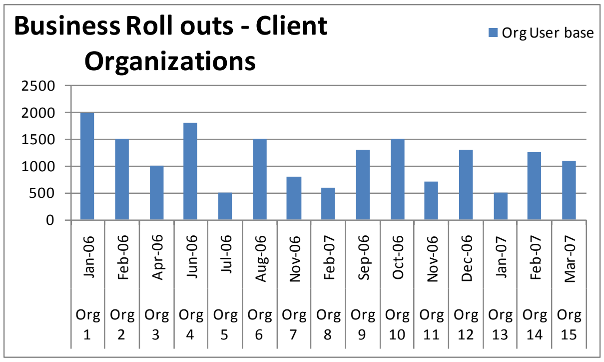

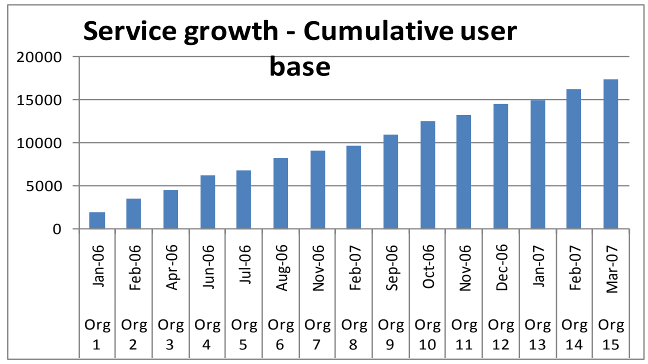

3.2. Framework Metrics: Workload

Note that for the workload and the resource tiers – there needs to be a uniform "time interval bucket" in which the metrics are either averaged over the bucket or aggregated/summed over the bucket. Typically a 5 minute interval works well across a majority of applications-loads. One can consider further smaller units down to seconds in case the requests come into the system at fast rate and/or high volumes. On the other hand you can expand the interval to 15 minutes to 1 hour subject to the load pattern being sparse/spread across. The smaller the time interval – the more data points and therefore better accuracy. The interval in the model is analogous to the "least count" of a physical instrument. The determination of the interval is based on the trade-off between "speed" at which the application needs to work, granularity of the Service Level Agreements and the expected "accuracy" of the model. In this case the time interval or the "bucket" was set to be 5 minutes.

Next is the turn of workload metrics – one for each of the significant tiers. The metric for the web-app tier was nailed down to be the number of http transactions in a 5 minute interval and for the DB tier it was the number of concurrent users over the 5 minute interval.

3.3. Framework Metrics: Resource

Based on the historical behavior and the way the application works – it turned out that the memory utilization is not a factor which varies linearly with the load. Thus the CPU was chosen as the resource metrics for both tiers. The charts below clearly bring out the justification for the corresponding choices. The X-axis has the time of day, the secondary Y-axis has the % CPU utilization, and the primary Y-axis in the first chart has the http transaction count while the second graph has the number of concurrent users (the workload metrics as defined above). The graphs show the variation in the workload and the resource metrics to be in synch – justifying the choice of both in corresponding tiers.

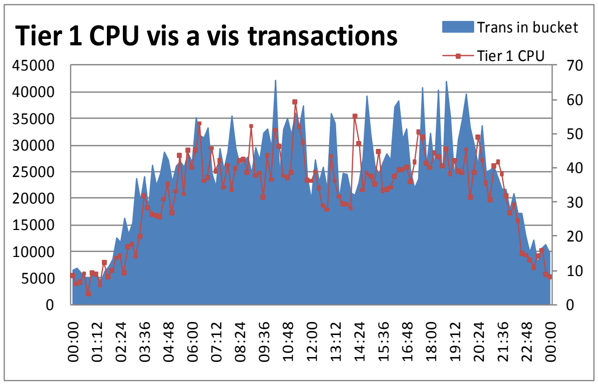

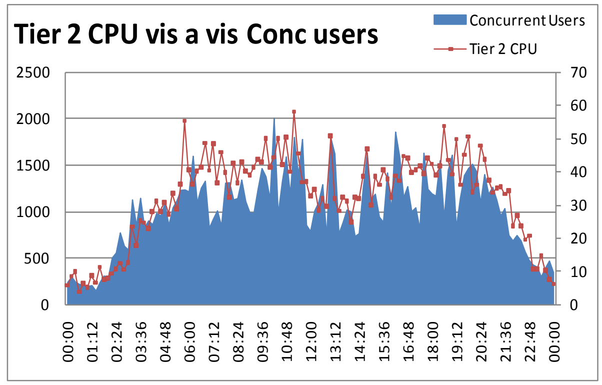

4. Case Study: Interlinking the metrics across tiers

The next step involves determination of the “correlation” factors across the workload and resource metrics. This is the crucial and the differentiating factor which leads to the advantage of this methodology over the conventional practices. The choice of the metrics – workload and resource has to be done in such a way that the variation is in synch. Different tiers turn out to have different metrics which affect the resource utilizations. The charts below bring out the correlation – note that there would always be a “band” type of variation – the correlations in practice would typically not have points lying on sharp curves. The X-axis has the workload metrics while the Y-axis plots the resource metrics: the first chart is http transaction count versus Web CPU and the second chart concurrent users versus DB CPU.

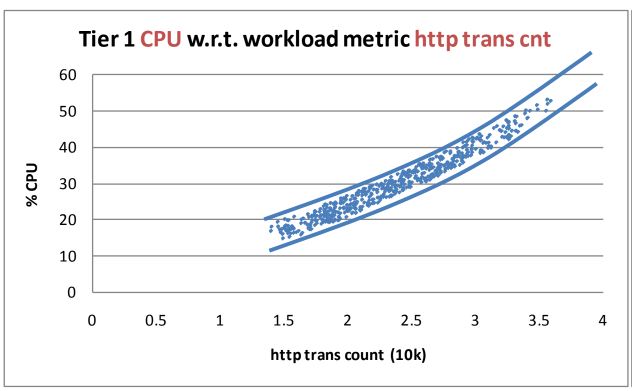

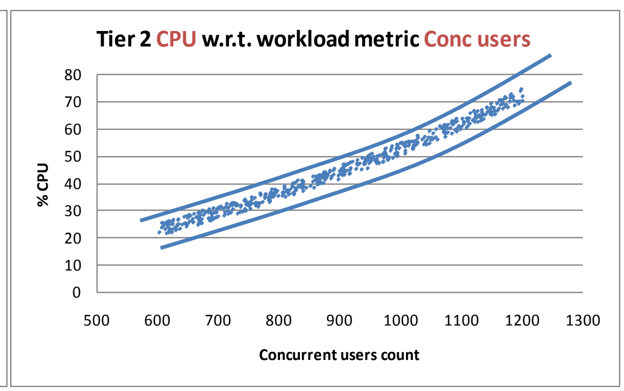

This step corresponds to the realization of the “lower part” in the aggregate philosophy of the framework.

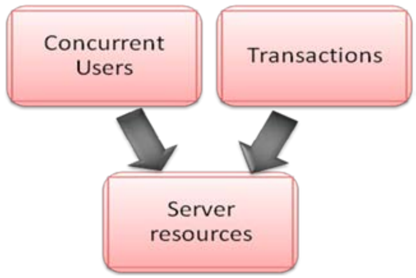

Relating lower tiers

Using correlation factors, one needs to "extrapolate" further within the bands. Finding the "correlation bands" and "extrapolation" is a straightforward mathematical exercise. The extrapolation then yields the capability of different tiers in terms of possibly different metrics subject to the thresholds. For example one can set the threshold to determine the capacity of the tier to be 80% CPU utilization or in some cases 100% memory utilization. (These are just indicative examples.) In the given case – the thresholds were used to be 80% CPU utilization for determination of the capacity. The capability thus turned out to be 50K http transactions per bucket for the web tier server and 2650 concurrent users for the DB tier server. Note that for individual servers in the same tier there can be different capacity figures depending on the design-deployment, load balancing, database usages etc.   Therefore, one can evaluate individual server capacities as the need may be subject to the variation in the utilizations shown.

The next step involves traversing one tier above to the business metrics. This is accomplished in a similar manner by finding the "correlation" factor across the "upper tiers". The business metrics being the cumulative user base – this needs to be translated into the workload metrics. Although the aggregate cumulative user base is a large number – the actual online users concurrently using the system are low and it's these concurrent users that "load" the system. Obviously, we had concurrent users as one of the workload metrics and these concurrent users in turn create the http transactions. The charts below elaborate the inter-conversion: first chart shows the time phased cumulative user base of the primary Y-axis while the secondary Y –axis gives the percentage of concurrent users among the aggregate user base at that time. The chart shows that the concurrency has been in the stable band of 7%-14% even if the number of registered users has increased.

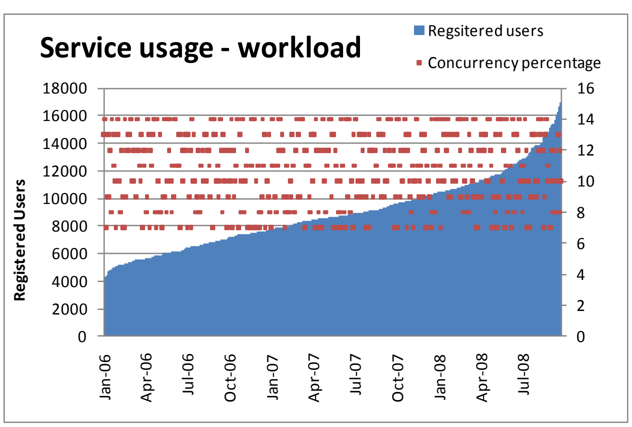

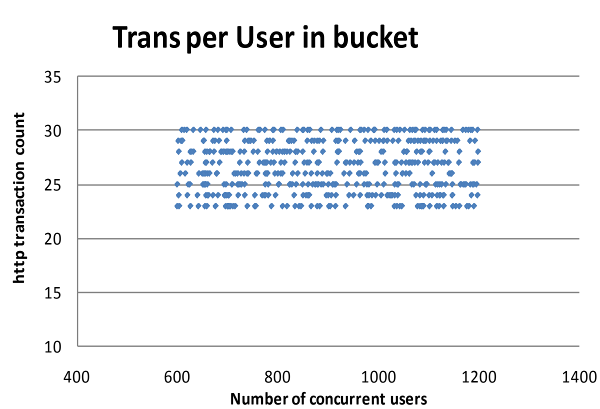

The second chart brings out the number of transactions the concurrent users are making – thus relating the workload metrics with one another. These charts (methodology involved) bring about the traversing from bottom resource to top business metrics.

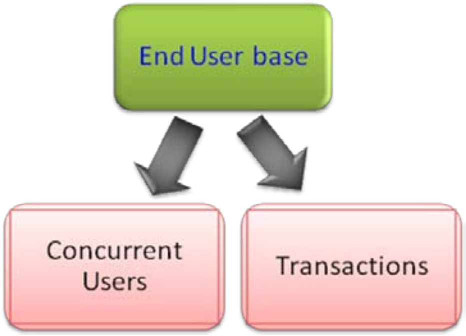

Relating the uppet tiers

The charts below give the capacity of the architecture in a single snapshot. The X-axis has the individual servers in different tiers, the primary Y-axis denotes the number of concurrent users and the secondary Y-axis gives the http transaction count.

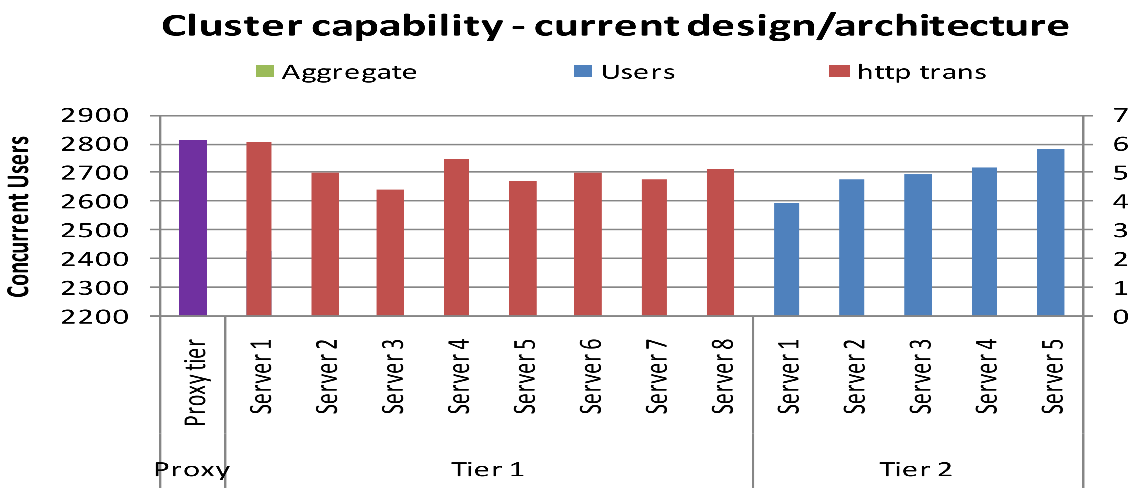

Note the variation is the capacities of individual servers within the same tier. The aggregate capability of the cluster-system is denoted using the green band. We now have the capability of the existing cluster.

5. Case Study: Business Decisions enabled by the framework

Turning to capacity planning in terms of business metrics – the growth plans give the anticipated business in terms of new roll-outs and corresponding user base rise. The chart below overlays the capacity in terms business metrics. The X-axis has the timelines for addition of clients – cumulative user base while the bands correspond to the cluster capability bringing out the life spans. It indicates the need for the 3rd cluster at the tail end of current horizon of business estimates.

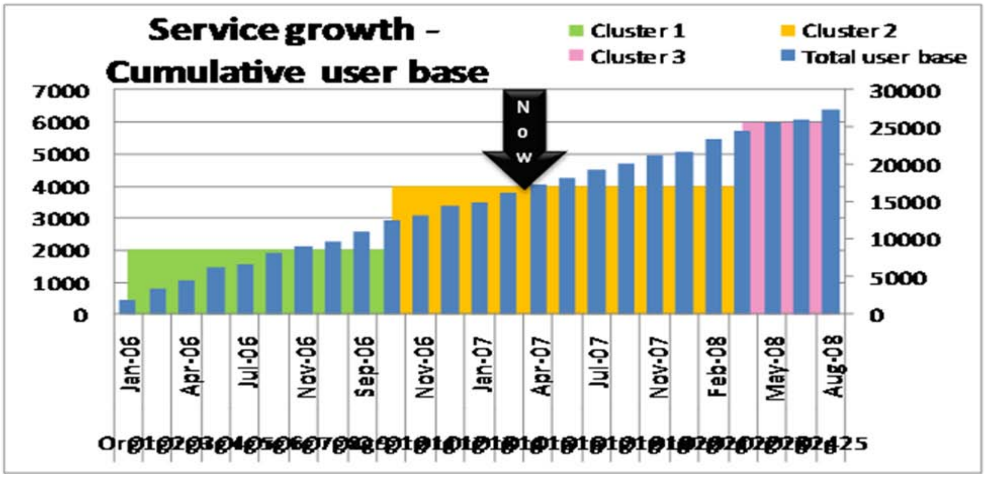

Deployment of 3rd cluster would involve significant CAPEX and efforts. One could "extend" the "shelf-life" of the existing cluster by enhancing their capacity so that the existing "2 cluster" pattern can stretch and include the "tail users". The solution would be enhancing the single core CPU's to dual core ones for the existing with appropriate optimal number of instances within the cluster. The chart below then gives the "extension" in capacity of individual clusters and how they will cater for the "tail part" – the "expansion" can be clearly seen vertical (in terms of users) as well as horizontal (in terms of timelines) thus accomplishing the capacity alignment with business goals.

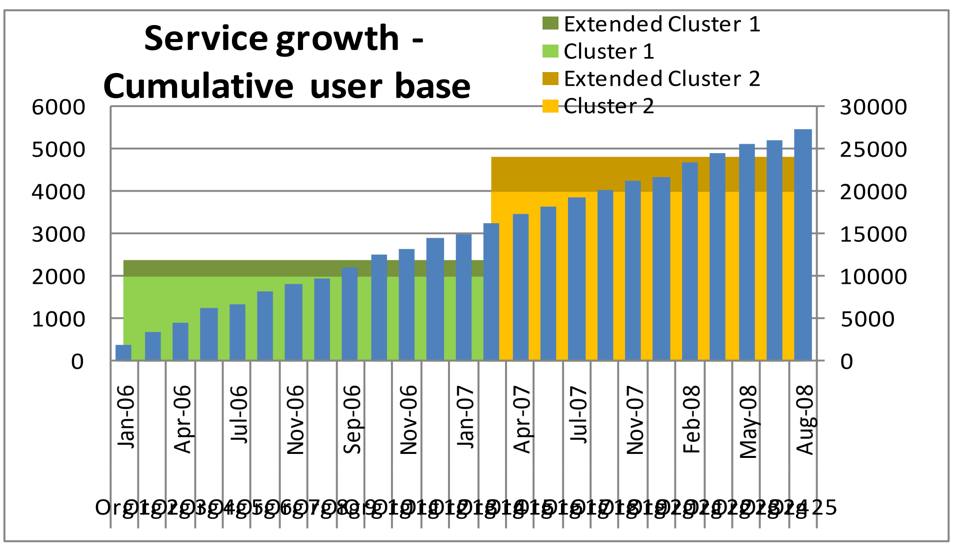

6. Summary

The chart below gives the pros and cons of the methodology as compared with the conventional approach:

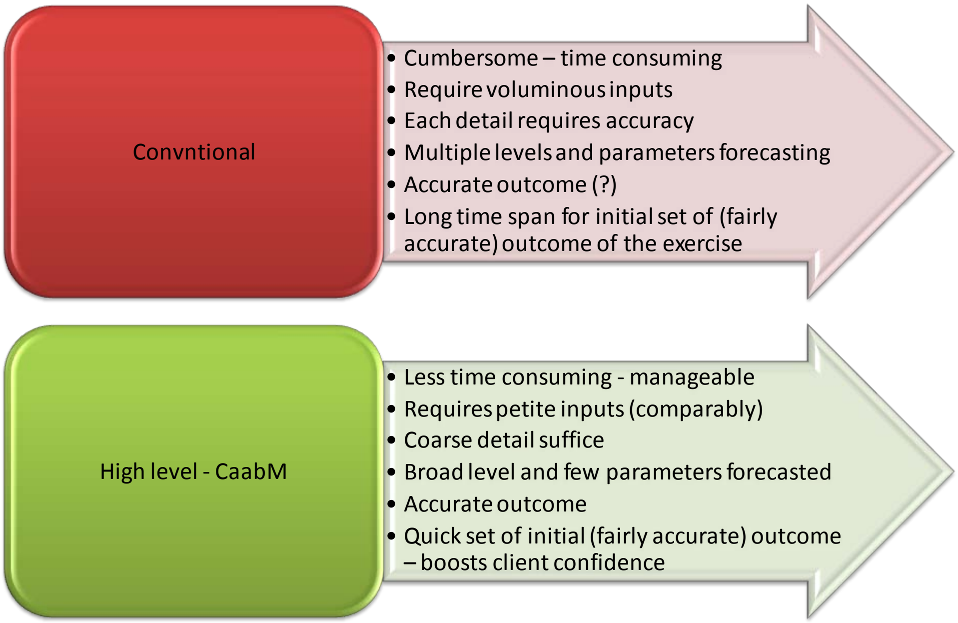

The proposed methodology-framework has in brief the following advantages:

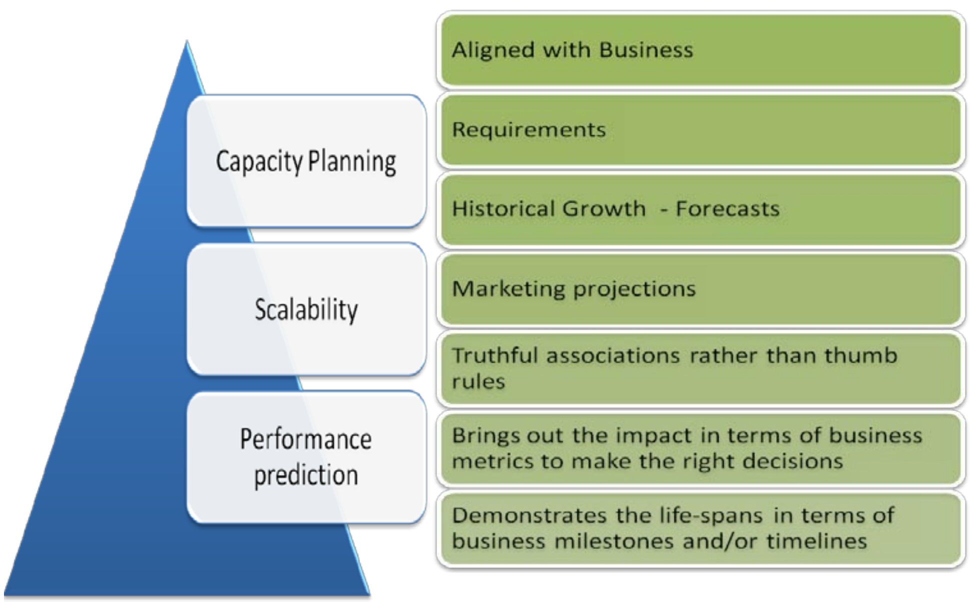

It addresses the following key issues which are the drivers behind the ITIL V3 by providing in part answers to :

- IT and Business strategic planning
- Integrating and aligning IT and Business goals
- Optimizing costs and the Total Cost of Ownership

- Demonstrating the business value of IT
- Delivering the required, business justified IT service(i.e. what is required, when and cost) etc

The framework has been applied across multiple platforms – architectures – applications by the authors. One can use a variety of metrics to suit the case under consideration – some already suggested in the framework. Further, the parameters for decision making can be subject to the goals as specified by the service being offered. The Case study demonstrated has been applied to Service provider in recent past and not only does it address the shortcomings of conventional practices but lead to the inclusion of Capacity planning expertise in further rounds of business expansion discussions and fine tuning agreements with clients on service level agreements.

7. References

[MA2000] D. Menasce and V. Almeida. Scaling for E-Business - Technologies, Models, Performance and Capacity Planning. Prentice Hall, Upper Saddle River, NJ, 2000.

[WA2002] Elaine J. Weyuker and Alberto Avritzer, "A metric for predicting the performance of an application. under a growing workload", IBM Systems Journal, 41(1), 2002, pp. 45-54.

[MAD1994] Menascé, D. A., V. A. F. Almeida, and L. W. Dowdy, Capacity Planning and Performance Modeling: from mainframes to client-server systems, Prentice Hall, Upper Saddle River, NJ, 1994.

[RJ1991] Raj Jain, "The Art of computer systems performance analysis", 1991; John Wiley & Sons, Inc.

[JRR2009] Shilpa Jain, Anil Rahate and Abhijit Ranjekar, "Capacity prediction based on Scaling and Growth Factors using Performance Test Data", CMG 09.

[AACIJS2007]Alison Cartlidge, Ashley Hanna, Colin Rudd, Ivor Macfarlane, John Windebank, Stuart Rance, "itSMF_An_Introductory_Overview_of_ITIL_V3", The UK Chapter of the itSMF, 2007

Managing Processor Usage in a ClearPath MCP Metering Environment Using Multi-level Detailed Management Reports

Wim te Lintum, Unisys Corporation

(wim.telintum@nl.unisys.com**)**

In order to proactively manage performance and capacity and to improve predictability of their mission critical ClearPath MCP mainframe systems, Senior IT Management of this customer was looking for easy accessible high level information about capacity and performance of these systems.

At the same time there was a need for their Capacity Manager to zoom in on possible bottlenecks and to follow the results of tuning and optimization efforts in applications and processes.

To fulfil these requirements we built a reporting tool for this customer with both a high, medium and detailed level of information on system usage and especially on processor usage, because that is the basic driver for metering costs and metering balance. Because of the different levels of information, the tool gives the customer the opportunity to zoom in (drill down) on usage per system, per brand, per workload type, per application and even per individual process for more detailed information on system usage.

The Challenge

The customer has a number of ClearPath Libra Model 690 MCP mainframes, running a health care insurance application for a number of brands. The contract for these ClearPath mainframes is based on the Unisys' pay-for-use business model, which makes use of the metering technology, which is included in the ClearPath MCP servers. In this pay-for-use model, "MIPS*months used" is the main chargeability characteristic. One of the key elements in this MIPS*month value is Processor usage. Part of this metering technology is a monthly metering report that is composed by and distributed from the MCP operating system automatically. Because he is charged for the number of MIPS*months used and processor usage is key in this MIPS*month value, it is very important for the customer to manage processor usage of the MCP systems. Therefore the customer asked Unisys to develop a detailed monthly metering report, which gives a much more detailed insight into which applications and processes are responsible for the processor usage.

The assignment

Design and build a detailed and fully automated monthly report that gives information on processor usage at different levels:

1. Companywide (for all ClearPath MCP based applications)
2. Per ClearPath MCP mainframe system
3. Per brand (on the Production system) or per environment (for the Test/Development system)
4. Per type of work
    a. online transaction processing, regular batch, sleeping reports (on the Production system)
    b. development, acceptance, testing (on the Test/Development system)
5. Per batch program
    a. For 20 named reports
    b. For the top-20 of reports with actual highest processor usage

The available Performance Management Tooling of the customer will provide the basic data for the monthly reports.

The design

The first phase of this activity was to:

- design a database that will be used to hold the required data for the monthly report;
- develop programs that use some standard reports out of the customers performance tool as input to fill the database with the required performance data;
- develop tooling that automatically creates the monthly detailed reports out of the database.

Originally we decided to call this central database the ‘dashboard’ database. Although this name of the database is somewhat confusing, we have not changed it.

The result of this design phase is visualized in the next two pictures. Figure 1 represents the flow of data, while Figure 2 gives a global overview of the structures in the Dashboard database:

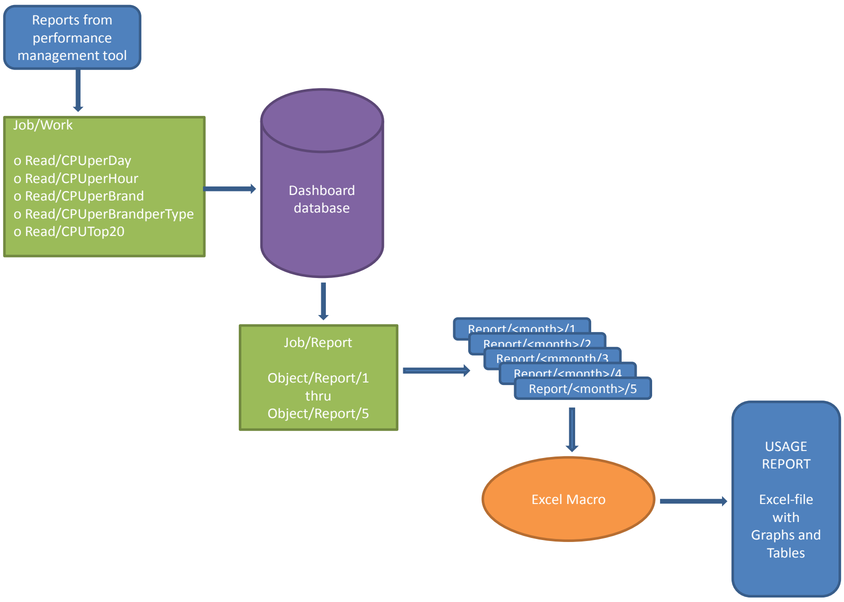

Figure 1 - Flow of data

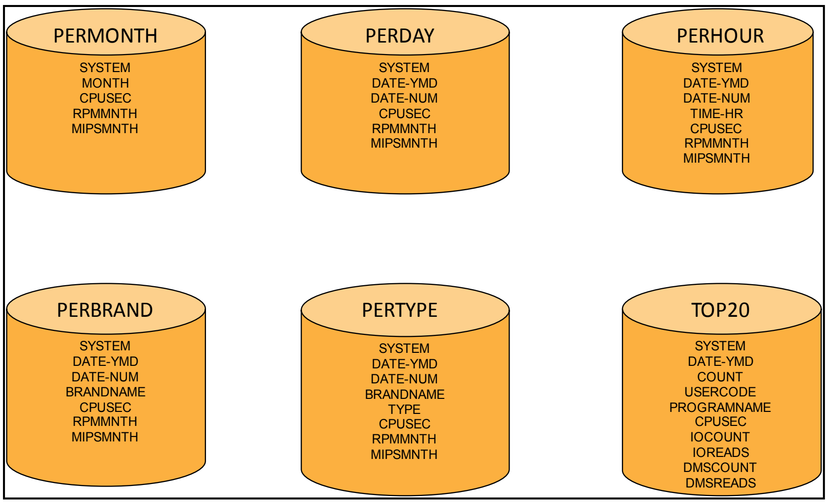

Figure 2 - Database model

The implementation

The second phase of this project can be divided into a number of steps:

1. Create reports using the existing performance management tool. Based on the defined requirements, a monthly set of basic reports (9 reports per day per ClearPath system) is created out of the customer's capacity management database. All reports are text files with a fixed name, fixed format and fixed layout.

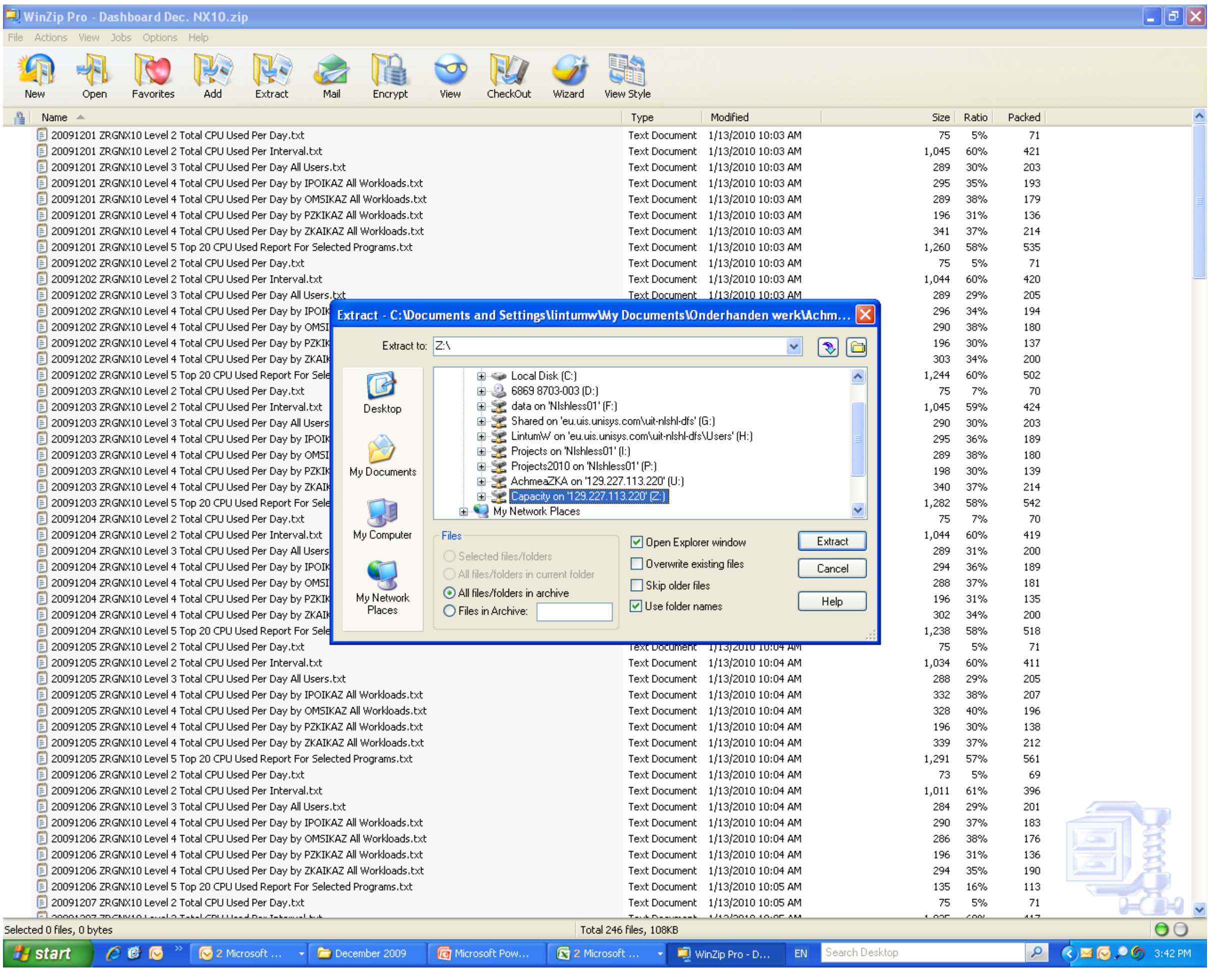

Figure 3 - Import standard performance reports

2. Develop (Algol) programs that read the text files (created in step 1), analyze the data in it and use it to fill the Dashboard database with information:
    a. Read/Cpuperday & Read/Cpuperhour ➔ Add records to data structures 'Permonth', 'Perday' and 'Perhour' (level 1 and 2)
    b. Read/Cpuperbrand ➔ Add records to structure 'Perbrand' (level 3)
    c. Read/Cpuperbrandpertype ➔ Add records to structure 'Pertype' (level 4)
    d. Read/Cputop20 ➔ Add records to structure 'Top20' (Level 5 – named reports)
    e. Read/Cputop30 ➔ Add records to structure 'Top20' (Level 5 – actual heavy-processing reports)
    f. A job (Dashboard/Job/Work) that executes these Algol programs

    All programs accept parameters that control the particular MCP mainframe and the date range for the analysis, for example: Run Object/Dashboard/Read/Cpuperday ("MCPProd1", 20100801,20100831)

```
START ACHMEA/DASHBOARD/JOB/VERWERK (PARTITION, DATE-FROM, DATE-UNTIL)
    RUN OBJECT/ACHMEA/DASHBOARD/READ/CPUPERDAY (PARTITION, DATE-FROM, DATE-UNTIL);
    RUN OBJECT/ACHMEA/DASHBOARD/READ/CPUPERHOUR (PARTITION, DATE-FROM, DATE-UNTIL);
    RUN OBJECT/ACHMEA/DASHBOARD/READ/CPUPERMERK (PARTITION, DATE-FROM, DATE-UNTIL);
    RUN OBJECT/ACHMEA/DASHBOARD/READ/CPUPERMERKPERTYPE (PARTITION, MERKNAAM, DATE-FROM, DATE-UNTIL);
    RUN OBJECT/ACHMEA/DASHBOARD/READ/CPUTOP20 (PARTITION, DATE-FROM, DATE-UNTIL);
```

3. Develop (Algol) programs that query the Dashboard database and use the data to create raw reports. All programs accept a parameter that specifies the reporting month:

```
BEGIN JOB ACHMEA/DASHBOARD/REPORT (INTEGER MAAND);
%
RUN OBJECT/ACHMEA/DASHBOARD/REPORT/1; VALUE = MAAND;
RUN OBJECT/ACHMEA/DASHBOARD/REPORT/2; VALUE = MAAND;
RUN OBJECT/ACHMEA/DASHBOARD/REPORT/3; VALUE = MAAND;
RUN OBJECT/ACHMEA/DASHBOARD/REPORT/4V2; VALUE = MAAND;
RUN OBJECT/ACHMEA/DASHBOARD/REPORT/5B ("<PARTITION 1>"); VALUE = MAAND;
RUN OBJECT/ACHMEA/DASHBOARD/REPORT/5F ("<PARTITION 1>"); VALUE = MAAND;
RUN OBJECT/ACHMEA/DASHBOARD/REPORT/5B ("<PARTITION 2>"); VALUE = MAAND;
RUN OBJECT/ACHMEA/DASHBOARD/REPORT/5F ("<PARTITION 2>"); VALUE = MAAND;
RUN OBJECT/ACHMEA/DASHBOARD/REPORT/5B ("<PARTITION 3>"); VALUE = MAAND;
RUN OBJECT/ACHMEA/DASHBOARD/REPORT/5F ("<PARTITION 3>"); VALUE = MAAND;
%
END JOB
```

4. Develop an Excel Macro that reads raw monthly reports (phase 3) and automatically creates the detailed standard monthly reports.

   Because this Macro needs some input parameters, a Macro Menu has been developed to supply the required parameters.

Figure 4 - Macro input screen

The following parameters have to be supplied:

1. Share-name and MCP directory point to the MCP location where the phase 3 reports can be found
2. Month (yyyymm and text) indicate the month for which reports have to be generated
3. The Windows directory is the name of the directory where the detailed monthly reports have to be stored.
4. Partition-1 thru Partition-n are the names of the MCP partitions at the customer site for which reports have to be generated.

The result

Running the Excel Macro described in phase 4 will deliver an Excel file with a number of tabs, filled with tables and graphs, preceded by a menu-tab, which can be used to jump to one of the standard detailed monthly reports:

Capacity Usage Reports <customername>		
Level	*Partition*	*Report*
I	Total	Metering balance versus prognosis
I	Total	Metering usage & -prognosis per month
I	Total	Metering usage per day
I	Total	Summary Metering usage
I	Total	Processor time per month
II	Per partition	Metering usage per day
II	Per partition	Metering usage per month
II	Per partition	Processor time per month
II	Per partition	Prognosis Metering usage per month
II	<partition-4>	Metering usage & -prognosis per month
II	<partition-1>	Metering usage & -prognosis per month
II	<partition-2>	Metering usage & -prognosis per month
II	<partition-3>	Metering usage & -prognosis per month
III	<partition-1>	Processor time per workload per day
III	<partition-1>	Table - Processor time per workload per day
III	<partition-2>	Processor time per workload per day
III	<partition-2>	Table - Processor time per workload per day
III	<partition-3>	Processor time per workload per day
III	<partition-3>	Table - Processor time per workload per day
IV	<partition-1>	Processor time per type per workload
IV	<partition-1>	Table - Processor time per type per workload
IV	<partition-2>	Processor time per type per workload
IV	<partition-2>	Table - Processor time per type per workload
IV	<partition-3>	Processor time per type per workload
IV	<partition-3>	Table - Processor time per type per workload
Vb	<partition-1>	Summary top-N reports - per day
Vb	<partition-1>	Summary top-N reports - total
Vb	<partition-1>	Top-N reports (I/O information)
Vb	<partition-1>	Top-N reports (Processor time)
Vb	<partition-2>	Summary top-N reports - per day
Vb	<partition-2>	Summary top-N reports - total
Vb	<partition-2>	Top-N reports (I/O information)
Vb	<partition-2>	Top-N reports (Processor time)
Vf	<partition-1>	Summary top-N reports - per day
Vf	<partition-1>	Summary top-N reports - total
Vf	<partition-2>	Summary top-N reports - per day
Vf	<partition-2>	Summary top-N reports - total
Vf	<partition-3>	Summary top-N reports - per day
Vf	<partition-3>	Summary top-N reports - total
X	Total	Total per month
X	Total	Total Summary per day

Figure 5 - Menu tab

The first column gives the level of detail of the report, the second one specifies the Partition name and the last column describes the specific report. Clicking on one of the report names (for example: Processor time per workload per day) will cause a jump to the tab with the specified report, which is a graph or a table. On this tab a Return-button is available to jump back to the Menu table.

Level I & II – Historical usage and prognosis

This graph gives the total usage during each month of the last one and a half year plus the original prognosis for the specific month

The same kind of graph is available for the total used processor time per month during the last one and a half year and for the Metering usage (and prognosis) for each of the Partitions.

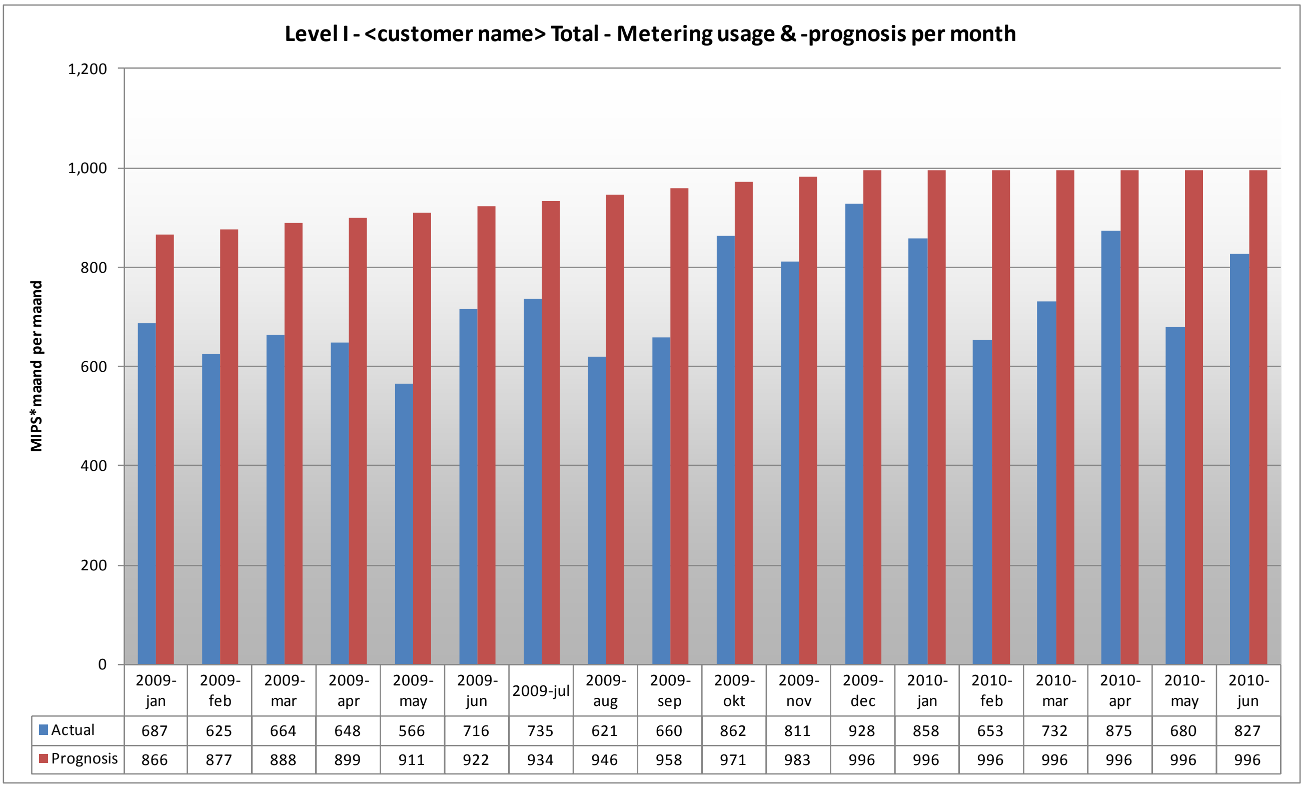

	2009-jan	2009-feb	2009-mar	2009-apr	2009-may	2009-jun	2009-jul	2009-aug	2009-sep	2009-okt	2009-nov	2009-dec	2010-jan	2010-feb	2010-mar	2010-apr	2010-may	2010-jun
Actual	687	625	664	648	566	716	735	621	660	862	811	928	858	653	732	875	680	827
Prognosis	866	877	888	899	911	922	934	946	958	971	983	996	996	996	996	996	996	996

Figure 6 - Level I - Historical Usage and prognosis

Level I – Metering balance and prognosis

The changing value of the Metering balance during the last one and a half year and the original prognosis is presented in a line diagram:

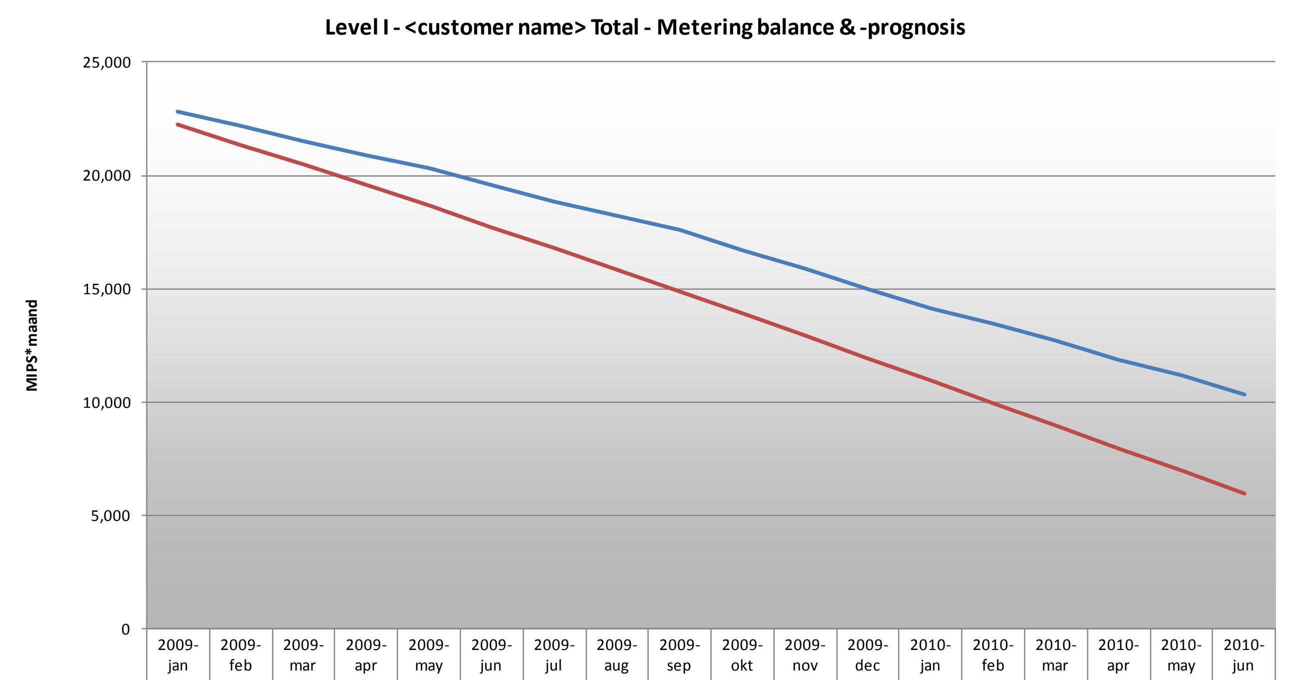

	2009-jan	2009-feb	2009-mar	2009-apr	2009-may	2009-jun	2009-jul	2009-aug	2009-sep	2009-okt	2009-nov	2009-dec	2010-jan	2010-feb	2010-mar	2010-apr	2010-may	2010-jun
Actual	22,808	22,183	21,519	20,871	20,305	19,590	18,855	18,234	17,574	16,712	15,901	14,973	14,115	13,461	12,729	11,854	11,175	10,348
Prognosis	22,235	21,358	20,470	19,571	18,660	17,738	16,804	15,857	14,899	13,928	12,945	11,949	10,953	9,957	8,962	7,966	6,970	5,974
Difference	573	825	1,049	1,300	1,645	1,852	2,051	2,377	2,675	2,784	2,956	3,023	3,161	3,504	3,767	3,888	4,205	4,374

Figure 7 - Level I - Metering balance

A table with the actual figures (both prognosis and actual plus the difference between the two) per month is presented below the graph.

Level I & II – Metering usage per dag

Figures for the total Metering usage per day is available in a bar graph. The same kind of information is available for each of the partitions.

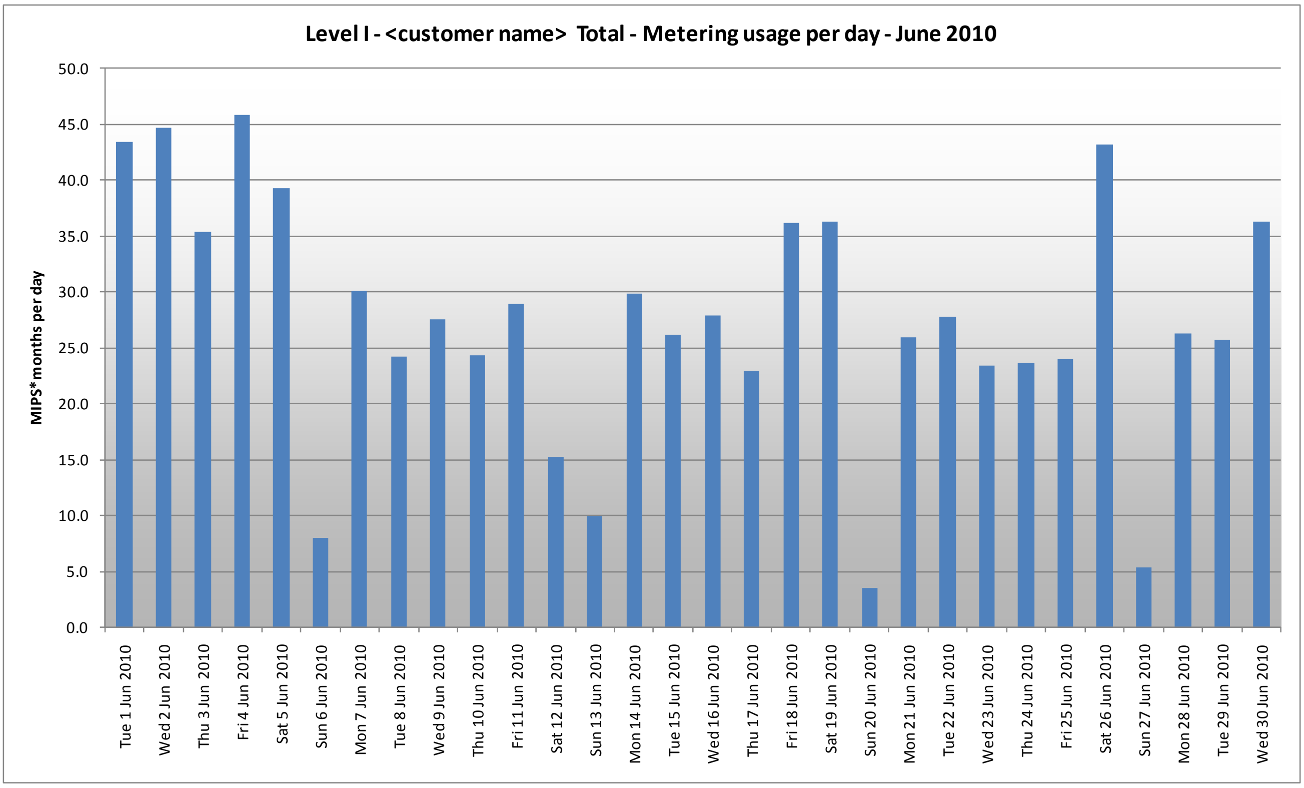

Figure 8 - Metering usage per day

Level III – CPU-usage per workload per day (per partition)

For each of the partitions the processor usage (expressed as processor time) is divided in processor usage per workload (or brand or user code or …):

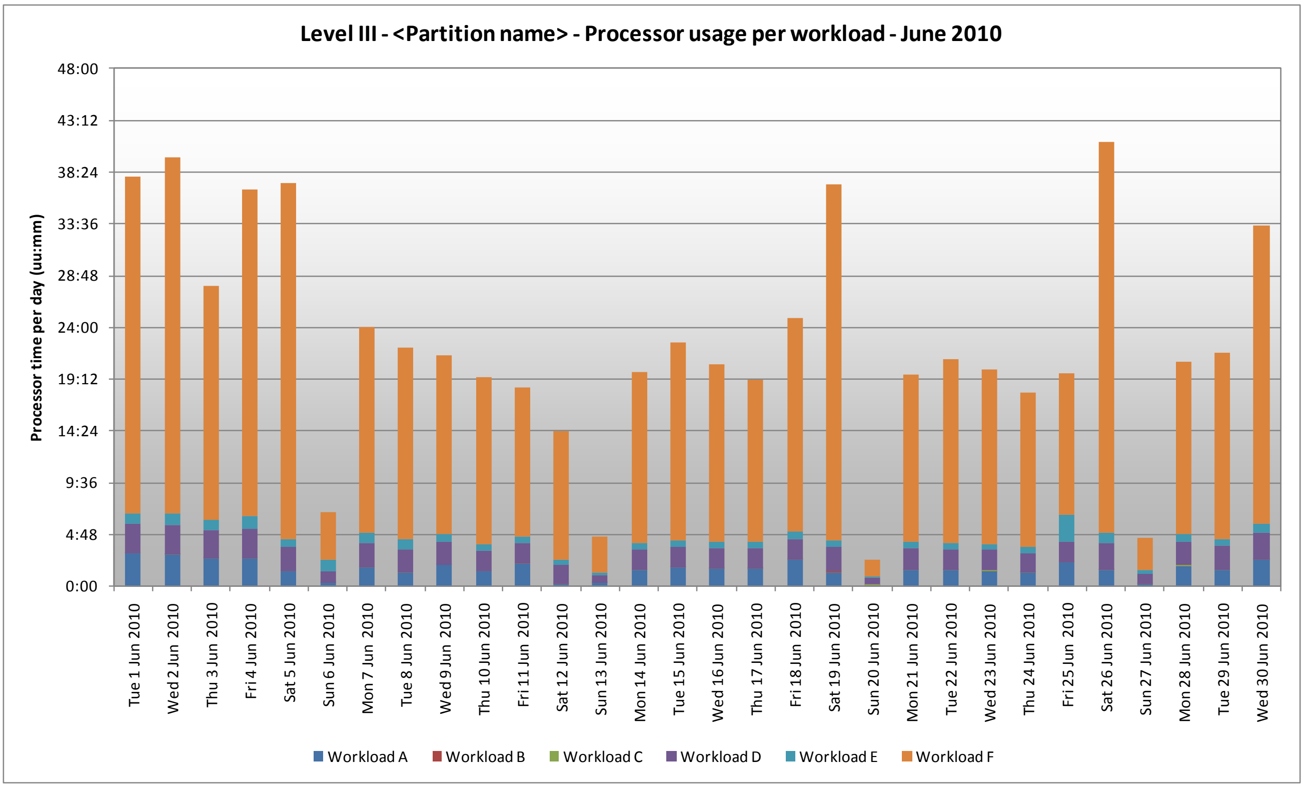

Figure 9 - Processor usage per workload

This information is also provided in a table:

Table 1 – Processor Usage per workload

Date	Workload A	Workload B	Workload C	Workload D	Workload E	Workload F
Tue 1 Jun 2010	3:01	0:00	0:00	2:47	0:58	31:14
Wed 2 Jun 2010	2:53	0:00	0:00	2:48	1:00	33:04
Thu 3 Jun 2010	2:33	0:00	0:00	2:39	0:56	21:41
Fri 4 Jun 2010	2:34	0:00	0:00	2:44	1:12	30:17
Sat 5 Jun 2010	1:26	0:00	0:00	2:10	0:46	33:00
Sun 6 Jun 2010	0:16	0:00	0:00	1:03	1:08	4:25
…………	…………	…………	…………	…………	…………	…………
Sat 26 Jun 2010	1:32	0:00	0:00	2:23	1:01	36:10
Sun 27 Jun 2010	0:08	0:00	0:01	1:00	0:18	3:02
Mon 28 Jun 2010	1:54	0:00	0:00	2:12	0:40	15:57
Tue 29 Jun 2010	1:32	0:00	0:00	2:09	0:36	17:17
Wed 30 Jun 2010	2:27	0:00	0:00	2:28	0:47	27:39
Totaal	**47:43**	**0:03**	**0:18**	**59:09**	**22:27**	**548:18**

Level IV – CPU-usage per type per workload

For each of the partitions, a monthly summary is available with processor time per workload type per workload, for example: Processor usage for all batch reports running under a specific user code:

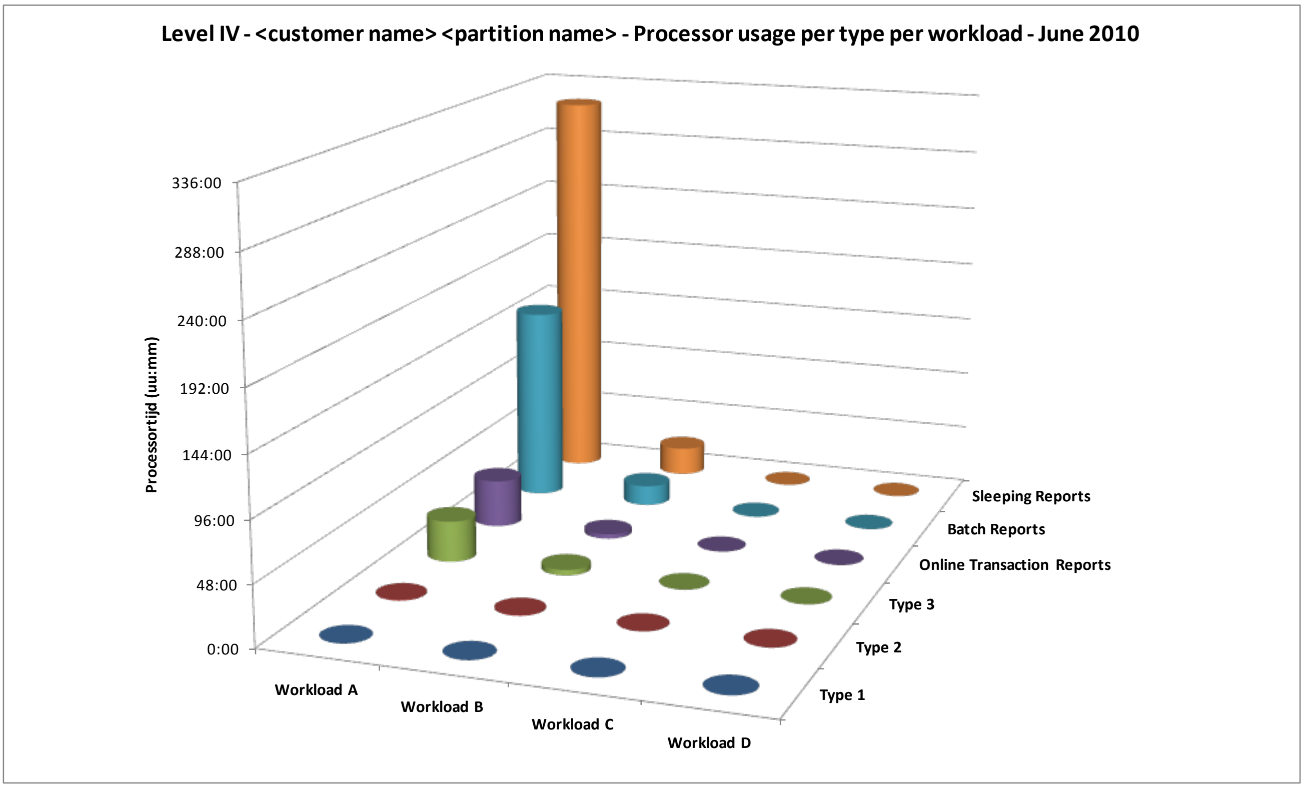

Figure 10 - Processor usage per type per workload

And again, this information is also available in a straight table:

Table 2 - Processor usage per workload

	Type 1	*Type 2*	*Type 3*	*Online Transaction Reports*	*Batch Reports*	*Sleeping Reports*	*Total*
Workload A	0:00	0:36	32:51	38:09	156:34	320:06	548:18
Workload B	0:00	0:32	4:27	3:34	16:21	22:47	47:43
Workload C	0:00	0:00	0:02	0:01	0:14	0:00	0:18
Workload D	0:00	0:00	0:01	0:00	0:00	0:01	0:03
Total	**0:00**	**1:08**	**37:22**	**41:46**	**173:11**	**342:55**	**596:24**

Level V – Top N-reports: Total Processor time per report

For each of a number of named reports (and separately for each of the actual top 20 reports), the total processor usage (time) is presented in bar graph format:

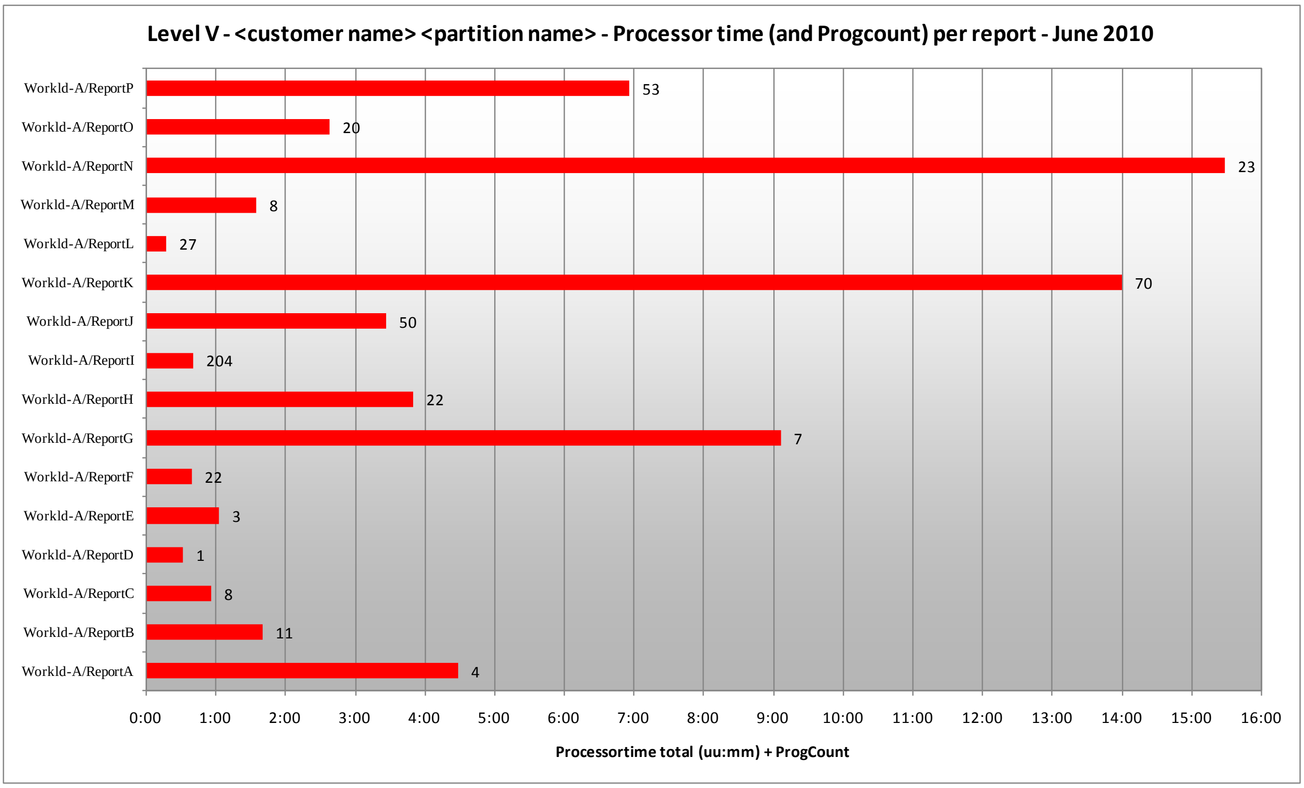

Figure 11 – Monthly processor time per (named) report

Besides the total processor time for the total month, for each report the number of times the report was executed (Progcount) is displayed in the graph.

Level V – Top-N - I/O-information per report

For each of the named reports (and separately for the actual top-20 reports), this graph presents the total number of database read-I/O’s, database write-I/O’s, non-database read-I/O’s and non-database write-I/O’s executed during that month:

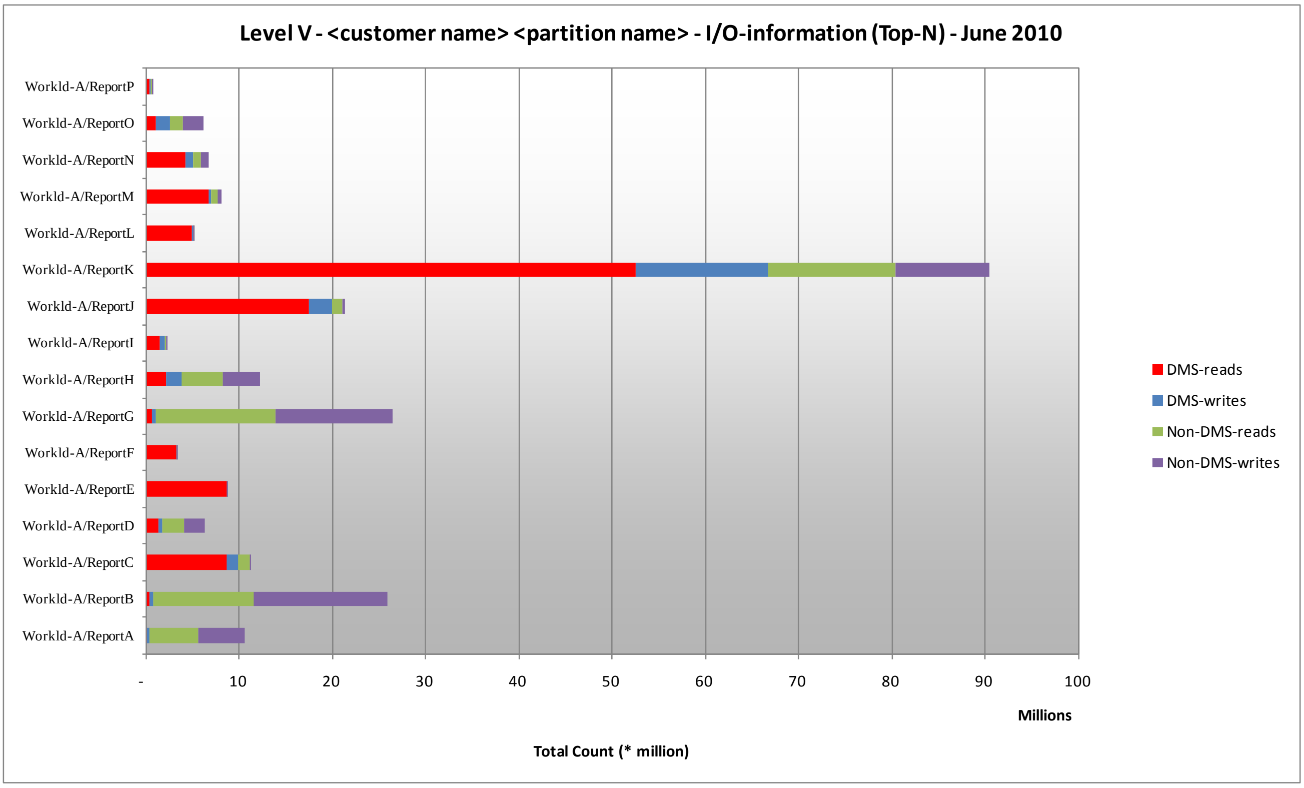

Figure 12 - Total monthly I/O count per (named) report

Again, this information is available in a table:

Table 3 - Total monthly I/O count per (named) report

Program name	Count	CPU (uu:mm)	I/O-count	I/O-reads	DMS-count	DMS-reads	DMS-writes	Non-DMS-reads	Non-DMS-writes
Workld-A/ReportA	4	4:28	10,621,824	5,384,586	319,964	91,951	228,013	5,292,635	5,009,225
Workld-A/ReportB	11	1:39	25,934,585	10,986,969	799,741	324,238	475,503	10,662,731	14,472,113
Workld-A/ReportC	8	0:56	11,303,284	9,946,964	9,849,372	8,671,077	1,178,295	1,275,887	178,025
Workld-A/ReportD	1	0:31	6,298,256	3,687,216	1,671,369	1,268,655	402,714	2,418,561	2,208,326
Workld-A/ReportE	3	1:03	8,660,057	8,604,706	8,611,164	8,556,385	54,779	48,321	572
Workld-A/ReportF	22	0:38	3,302,102	3,280,531	3,271,661	3,250,304	21,357	30,227	214
Workld-A/ReportG	7	9:06	26,390,176	13,546,873	981,840	599,662	382,178	12,947,211	12,461,125
Workld-A/ReportH	22	3:49	12,158,946	6,643,241	3,770,634	2,152,950	1,617,684	4,490,291	3,898,021
Workld-A/ReportI	204	0:40	2,141,961	1,547,584	2,039,195	1,473,261	565,934	74,323	28,443
Workld-A/ReportJ	50	3:26	21,315,095	18,527,343	19,968,352	17,455,910	2,512,442	1,071,433	275,310
Workld-A/ReportK	70	13:59	90,456,669	66,164,601	66,696,246	52,443,260	14,252,986	13,721,341	10,039,082
Workld-A/ReportL	27	0:16	5,049,525	5,000,768	4,972,026	4,925,005	47,021	75,763	1,736
Workld-A/ReportM	8	1:35	8,115,014	7,370,683	6,969,571	6,711,015	258,556	659,668	485,775
Workld-A/ReportN	23	15:28	6,703,629	5,092,419	5,065,642	4,252,037	813,605	840,382	797,605
Workld-A/ReportO	20	2:37	6,176,336	2,499,420	2,504,502	1,041,406	1,463,096	1,458,014	2,213,820
Workld-A/ReportP	53	6:55	714,492	495,643	473,814	334,783	139,031	160,860	79,818

About the authors:

John Meisenbacher **(john_meisenbacher@ieee.org)** is a Systems Engineer for the Server Capacity and Performance Management team at a major financial institution. He has more than 25 years experience with engineering system ranging from real-time embedded systems to large scale information systems, specializing in system modeling techniques to identify system and performance issues. Mr. Meisenbacher holds a M.S. in Systems Engineering from Johns Hopkins.

Karl Steger **(karl_steger@sbcglobal.net)** is VP/Technology Manager for the Server Capacity and Performance Management team at a major financial institution. He has more than 20 years' experience as a proven team leader, specializing in the identification of system performance and IT process issues. Mr. Steger holds a B. S. in computer science from the University of Missouri – St. Louis, and holds an IT Service Management Practitioner certificate in capacity management (ITIL v2).

Mr. Akinnuwesi Boluwaji Ade **(akinboluade@yahoo.com)** Mr. Akinnuwesi Boluwaji Ade has both B.Sc and Master of Technology (M.Tech) degrees in Computer Science. He is presently on his PhD degree program. He is a Lecturer in the Department of Information Technology, College of Information and Communications Technology, Bells University of Technology in Ota, Ogun State, Nigeria. He is a member of Nigeria Computer Society (NCS), Computer Professionals Registration Council of Nigeria and ACM. He is a student member of IEEE. He has publications in journals and conference proceedings. His research areas are Software Engineering, Software Performance Engineering, Expert System, Fuzzy Logic and Artificial Neural Networking. He speaks both English and Yoruba languages

Mr. Ezike J.O.J **(josephezike@yahoo.com)** Mr. Ezike has his Bachelor of Technology degree in Computer Engineering. His Master degree is in Computer Science. He is presently a PhD student. He is a lecturer in the Department of Computer Science in Bells University of Technology, Ota, Ogun State, Nigeria. His areas of research are e-learning systems, computer networking, distributed system, software engineering. He speaks English, Yoruba and Ibo languages.

Miss Shakirat Oluwatosin Raji **(tosinraj@yahoo.com)** Miss Shakirat Oluwatosin Raji has her Bachelor's degree in Computer Science; she has her Masters degree in Bioinformatics. She is currently a lecturer in Department of Information Technology in Bells University of Technology, Ota in Ogun State of Nigeria. Her research interest areas: Management Information Systems, Bioinformatics, Database System and Software Engineering. She speaks English and Yoruba Languages.

Dr Abhijit Ranjekar **(abhijit.ranjekar@gmail.com)** has a doctorate in the field of Operator Theory, Mathematics; was in the teaching profession and later joined the software R&D streams in order to apply the mathematical techniques. Dr. Ranjekar has worked in various areas from algorithms for planning and scheduling software (ERP-MRP) to performance modeling over the past ten years.

Swati Dorge **(dorge.swati@gmail.com)** holds a Masters in Mathematics and has been working in IT post teaching assignments for last four years.

Both authors have been involved in performance modeling using various tools – available in the market as well as self-developed. They have created many utilities to analyze performance and modeling techniques which have been applied across variety of applications ranging from portals to products to custom applications to COTS components. The techniques invented play critical roles and provide significant inputs to Performance Engineering practices.

Wim te Lintum **(wim.telintum@nl.unisys.com)** is a Systems Architect working for Unisys in The Netherlands. He has more than 35 years experience in supporting customers of ClearPath MCP systems, especially in the area of business continuity, availability, capacity management, database administration, security, storage solutions, automating system management tasks, unattended operations and ClearPath integration solutions.

2010 Board of Directors and CMG Staff

President
Donna S. Folkerts
ACS, Inc.
Bus: 630-468-2262
cmgpres@cmg.org

Vice President
Frank M. Bereznay
IBM
Bus: 626-564-7530
Fax: 626-564-3595
cmgvp@cmg.org

Treasurer
Charles Savage
Bank of America
Bus: 404-607-3060
Fax: 706-387-0395
treasurer@cmg.org

Secretary
Shana J. Bereznay
ACS
Bus: 714-630-0309
secretary@cmg.org

Director
Glenn R. Anderson
IBM
Bus: 312-635-1319
Fax: 312-635-1319
grander@us.ibm.com

Director
Ellen M. Friedman
Bank of Tokyo-Mitsubishi-UFJ
Bus: 201-413-8839
ellen@srmassoc.com

Director
Adam Grummitt
Bus: +44 (0) 1823 259231
Fax: +44 (0) 1823 334502
adam@grummitt.com

Director
Denise P. Kalm
CA. Inc.
Bus: 925-946-1384
denise.kalm@ca.com

Director
Susan Schreitmueller
IBM
Bus: 336-508-7728
Fax: 336-584-7499
rovrallovr@us.ibm.com

Director
Dave Thorn
Bus: 856-216-7550
dthorn55@gmail.com

General Chair
William Jouris
Jouris Consulting
Bus: 925-855-9512
cmgpc@cmg.org

Office
CMG Headquarters
Computer Measurement Group, Inc.
Bus: 856-401-1700
Fax: 856-401-1708
cmghq@cmg.org

Office Manager
Barbara Flemming
Computer Measurement Group, Inc.
Bus: 856-401-1700
Fax: 856-401-1708
barbara@cmg.org

Program Coordinator
David Troxel
Computer Measurment Group, Inc.
Bus: 856-401-1700
Fax: 856-401-1708
david@cmg.org

Conference Coordinator I
Michelle C. Cervantes
Computer Measurement Group, Inc.
Bus: 856-401-1700
Fax: 856-401-1708
michelle@cmg.org

Administrative Assistant
Kathleen M. Kinnarney
Computer Measurement Group, Inc.
Bus: 856-401-1700
Fax: 856-401-1708
kathy@cmg.org

Corporate Bookkeeper
Linda G. Stermer
LGS Bookkeeping Services
Bus: 805-238-2410
Fax: 805-238-2458
lgstermer@aol.com

Food & Beverage
Hugh Hunt
Hunt Conference Group, Inc.
Bus: 817-410-4660
Fax: 817-410-4661
hugh@huntconferencegroup.com

The Computer Measurement Group, Inc.

151 Fries Mill Road
Executive Campus, Suite 104
Turnersville, NJ 08012
www.cmg.org