Enterprise Capacity Management –
Assessment of Related Best Practices and Research Sponsored[1] by Nationwide Insurance

Project Sponsor: David Pike/Director of IT Architecture/Nationwide

Many Nationwide Contributors (including Chris Newland, John Northland, Jason McKitrick) have provided various inputs to this document

Enterprise Architecture Contributors: Aaron C. Roberts, Brett Gerke, Joe Bolinger, John Pan (Nationwide)

CETI Faculty: Jay Ramnathan, Rajiv Ramnath

5/30/2007

CETI
(CERCS[2] FOR ENTERPRISE TRANSFORMATION AND INNOVATION)
COMPUTER SCIENCE ENGINEERING, THE OHIO STATE UNIVERSITY
OSU-CISRC-5/07-TR38

This document provides 1) a state-of-the-art assessment of current methodologies, practices, 2) a proposed ACE – Adaptive Complex Enterprise - framework for future implementations, and 3) a short-term and long-term roadmap. The framework facilitates integrated and predictive enterprise-wide capacity management. Finally, the document shows that the very same processes that enable better capacity management also positively impact other enterprise-wide initiatives and can be cross-leveraged.
(This document is written for a wide audience and does not contain proprietary information.

---

[1] Nationwide has sponsored OSU CETI to research and synthesize best practices in capacity management and enabling architecture concepts that are applicable to a large-scale enterprise. The roadmap included here is developed in the context of many related and advanced on-going efforts at Nationwide Insurance. The CETI participants would like to acknowledge Nationwide's sponsorship of enterprise architecture projects and industry-university collaboration.

[2] CETI is a component of CERCs - center for experimental research in computing systems - is An NSF Industry University collaborative Multi-institutional research Initiative of Georgia Technology Institute and planned at the Ohio State University.

Table of Contents

0. Introduction

There are tactical and strategic imperatives to address the enterprise capacity management challenge.

Tactical reasons:

- Ensuring current and future quality of IT services
- Enabling shared IT services as a way of optimizing operational costs
- Business-IT alignment through activity-based costing and chargeback

Strategic reasons:

- Effective enterprise-wide operations through visibility and traceability - improving capacity management maturity through enabling practices and systems will also improve other enterprise functions (such as development of new services, customer service, business accountability and auditing). Traceability is more than trending of a particular set of variables within a function – it relates trends across functions in a value chain.
- Agility – innovation of new business services and through Service Oriented Architectures (SOA) needs traceability to underlying IT services and their quality to ensure robustness[i].

That is, the true value of capacity management is in addressing the issues and putting in place the infrastructure for enterprise-wide agility.

Project Context

Problem Statement

There is a gap in Business-aligned capacity planning, performance analysis and sizing/scalability techniques.

Project Charter

The project considers both practices and technologies within Nationwide and also the industry at large. In addition to reviewing existing frameworks and work products for managing infrastructure capacity, sizing, and scaling requirements, the project also identified emerging research.

Project Scope

The scope covered is as identified below:

Start with Nationwide current state and needs
Identify policies that relate SLAs to capacity management
Identify the Capacity Management requirements that need to be addressed
Research and provide industry best practice frameworks and identify gaps
Research and evaluate tools and technologies in capacity and performance management
Research CMDB as enabler for effective capacity management
Research architecture monitoring and relate to capacity management
Asses research status within the academic institutions
Provide roadmap and recommendations

Context within the enterprise

Typical relevant and important in-house efforts include.

- Capacity management team and advanced server-based methodologies for projecting capacity, in-house case studies showing the traceability of business service volume to capacity use
- Process for reviewing capacity requirements of solutions
- Business modeling
- Storage Capacity Management Process
- Application testing requirements
- Performance Engineering
- SOX audit

This document relates and integrates such efforts.

Deliverables and the Chapters

…of this document are:

Chapter 1: **Adaptive Complex Enterprise (ACE) Framework for Modeling, Monitoring, Analysis and Improvement**: to address challenges of implementing needed enterprise-wide communications.

Chapter 2: **Capacity Management Roadmap**: both for the short term and for the longer term, addressing 'what' it should be and 'how' it can be approached.

Chapter 3: **Existing Best Practices and Gap Analysis**: A summary of current practices and gaps related to capacity management.

The key findings are summarized below.

Key Findings, Recommendations and Log-term view of Capacity Management

Key Findings

As with other enterprise-wide initiatives, improving the maturity of capacity management is a difficult problem. Enterprises face organization and technology challenges in an attempt to better manage their *shared* IT Capacity. Differing business strategies and goals within different operational functions makes it difficult to:

1. connect past capacity usage with future capacity needs
2. provide visibility about available capacity, especially for on-demand and unpredictable events (business fluctuations, emergencies, unforeseen events etc.)
3. charge accurately for capacity use

Since IT capacity is expensive and also directly affects the quality of business services, existing best practices suggest that Capacity Management (Planning and Execution) aspects be included in the forefront of business and operations planning, and be informed from *business scenarios*. To accomplish this and establish a unified enterprise-wide plan, it is first useful to understand what Capacity Management is, and what it is not. Capacity management is described through related perspectives [2]:

- *Business service capacity:* referring to the demand a given business domain can satisfy as identified by its market segment or demographic. This is the potential of the business.
- *Service execution capacity:* supports business service capacity and referring to the demand a particular service can sustain. This is the current service as delivered to the customer. And,
- *IT resource capacity:* referring to the demand that a given resource can sustain in support of some level of Service Capacity. This is the application and underlying IT capacity.

The interaction of these perspectives is illustrated in Figure 1 and is defined as *Enterprise-wide Capacity Management*[ii]. The current state of practices and tools can be summarized as follows.

1. Existing IT best practices (ITIL, TOGAF, MOF etc.) related to capacity management focus more on improving the IT operational processes rather than measuring and relating IT assets (servers, applications etc.) used to meet business requirements.
2. Asset testing and monitoring tools that are useful for identifying capacity requirements are, by and large, specific to the asset. Thus, for ***shared*** and composite assets (e.g. a mainframe server on a network), it is very difficult to answer questions like: "How much of this server-asset capacity is available to handle emergency claims requests?" "How can 'claims' requests impact other business services, indirectly through the capacity of assets?" "How does seasonal variation impact other business domains?"
3. Emerging and existing monitoring and business intelligence tools for application-level monitoring have promise, but to be successful they must synthesize and relate IT usage to business-level abstractions for effective and predictive decision making. Currently there is significant overhead in the implementation these tools for true enterprise-wide *capacity management*.
4. ITIL suggests the use of a CDB (Capacity Database). However, this concept needs to be integrated with CMDB (Configuration Management Data Base) that maintains the installed IT items, their attributes (such as id, location, status), and associations
5. In summary, as illustrated in Figure 1 (refer to details starting page 16), it is difficult to precisely relate the IT resources used to deliver particular business service capacity. This disconnect in tools and methods has to be bridged with a methodology that provides the discipline for enterprise-wide planning.

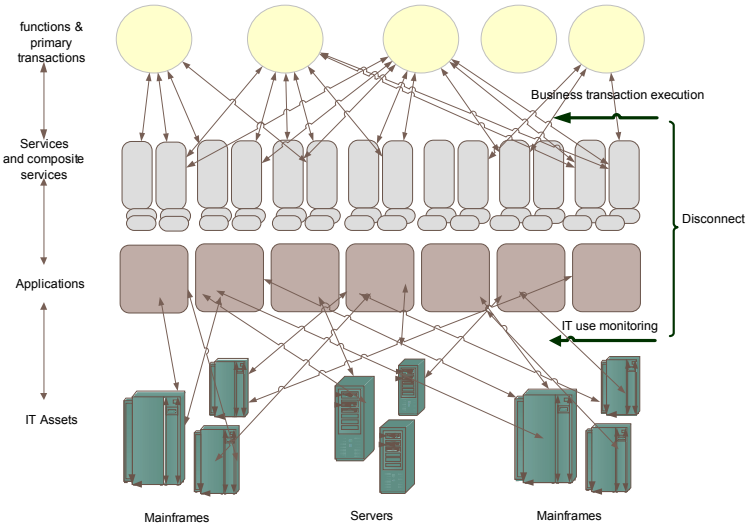

Figure 1: Many-to-many interactions between Business, Service and IT Capacity components and the Business-IT traceability disconnect.

Recommendations

Enable enterprise-wide capacity management through an *integrated* representation of the enterprise architecture that can be used by different Business and IT stakeholders to share their understanding of needs and to make decisions that affect multiple functions. This representation can also be used to provide traceability with existing practices/tools. We define *BIOS traceability* - *Business-IT use-Operations-Strategy* - as two-way as-is/to-be cause-and-effect change in metrics along the associations between Business services, *IT* resource use, Operations or process execution and Strategy. For each single *business service and request type*, this Business-IT or BIOS traceability will allow us to relate and analyze:

Business – $$ investments needed

Operations – # increase in the number of claim requests serviced

IT use – % increase in the use of specific assets (server, database etc.)

Strategy - % increase the presence in a market

With this traceability achieved at a *fine grain*, we can do mature enterprise-wide capacity planning. Understanding the point about the fine grain is important. - capacity management is done on a day-to-day basis because the quality of service on a daily basis is critical. In contrast, business planning however works on a *macro-grain* needed for the financial cycles.

Thus, at a high level the recommendations are:

1. *Roles and responsibilities:* Create an *enterprise-wide model* clearly identifying and surfacing the points of cause-and-effect metrics collection for traceability and capacity management. Identify roles and responsibilities for data synthesis and decision making. This is accomplished by ACE.
2. *Achieve process flexibility:* By creating a standardized enterprise-wide representation for traceability that allows diverse groups to make their IT use and business strategy explicit, the efforts of different groups within and across the various business domains can be published and shared for more effective decision-making an on-going basis.
3. *Data collection:* Provide a logical enterprise model layer with needed traceability relationships, capacity and runtime information to facilitate decision making at the right levels by the different BIOS stakeholders.
4. *Cross leverage:* Does the investment of time and resources in capacity management payoff? Capacity management is becoming important with a shift back to mainframes. However, the investment in the enterprise-wide frameworks can be further justified by having a positive impact not only on capacity management but also on many other related initiatives for shared services (e.g. service-oriented architectures, chargeback, SOX compliance). This will help justify the investment in more comprehensive enterprise-wide improvement programs and related organization, methods, and tools.
5. *Automation:* Research the automated monitoring of interactions between business users and IT assets in order to acquire SLA (service level agreements) requirements without human intervention. The objective of this proposed *socio-technical* research is to gather BIOS traceability information for predicting service requirements, but with minimum overhead impact to operations. This will also reduce the cost of information gathering.

Long-term Improvement Vision

Before getting into details, certain advantages of enterprise-wide frameworks must be noted. In

addition to having a positive impact on Capacity Management, by many other related activities also benefit. For example:

- Service catalog - a definition of the services, the transactions and resources to complete the services
- SLA to OLA mapping – the SLAs required for the service and the OLAs that apply to the underlying resources
- Capacity planning – defining, gathering, reporting, driving change from capacity metrics
- Chargeback - charging for services based on the true cost of the service
- Audit of IT controls based on the role of resources in service transactions
- Load testing and resource allocation – defining, performing, interpreting results, driving change in components from load tests
- Equipment Life cycle – budget, purchase, deployment, retirement (that has to be incorporated into capacity costs)
- Operational architecture – how to monitor, control, failover, services
- Service Implementation – taking business-side service requirements, integrating them with product development components to produce production services

It is possible to theoretically construct (this is not yet implemented and is a proposal) a self-monitoring architecture for capacity management as shown in the figure below. The architecture is composed of BIOS perspectives of individual organizations which, taken together, provide a unified model of the influences driving the capacity of an Enterprise. Although each perspective is conceptually independent, the real value of such an architecture is in eliciting the points of concern from one perspective which may be taken in account by the concerns of another. As illustrated in the Figure 2, the 'vertical' enterprise approach is conceptually motivated as follows:

- In the innermost level, denoted as *IT Transaction Monitoring*, the traditional fixed asset and infrastructure view of the Enterprise is captured. The primary concern of this viewpoint is that of maximizing the utilization of available resources with respect to the basic capacity elements, including disk, memory, and CPU time. Managing capacity from this perspective is rooted in the historical data of application usage which provides fundamental insight as to the capabilities of the infrastructure but lacks the predictive power that an agile Enterprise demands.

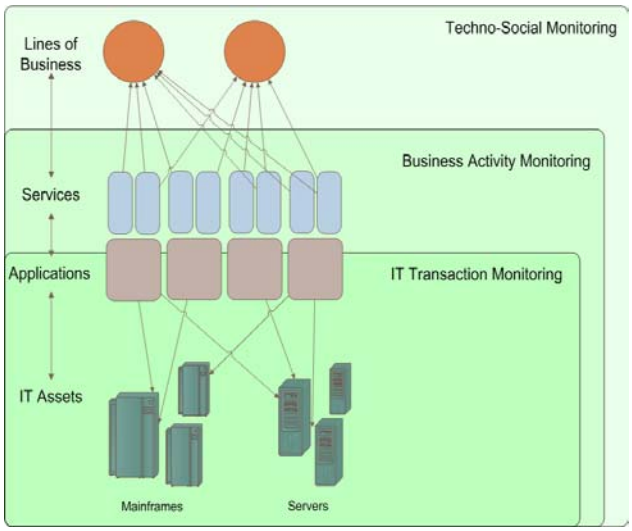

Figure 2: Long term improvement view.

- Moving to the *Business Activity Monitoring* perspective, the primary concern becomes adhering to service level agreements. This view captures the relationship between enabling services provided by IT and the value-adding services offered by the business. As with IT Transaction Monitoring, over time this perspective provides knowledge of how capacity may change as changes in business services occur, hence yielding additional predictive power beyond that of IT Transaction Monitoring. This increase in predictive potential is useful to Capacity Management because better estimation can be done before

the effects of a business activity change manifest themselves at the application layer. This also provides enough time for acquisition and change management cycles.

- The *Socio-Technical Monitoring* emphasizes understand the relationship between IT services and how the Lines of Businesses consume them [Argyris78, Kaplan]. The difference between this perspective and that of Business Activity Monitoring is very subtle but once again will yield additional predictive power, with respect to capacity, that the other viewpoints are not capable of alone. To clarify, recall that under Business Activity Monitoring the focus is on the relationship between IT services and business services. This is a relationship between IT and what IT's customers, the Lines of Businesses, do. While this is undoubtedly a useful relationship to understand in an Enterprise it leaves one with the, often extremely difficult-to-quantify, problem of determining when and how business services will change. This is a knowledge gap that is difficult for existing tools to capture and requires expensive explicit communication between IT and business staff, ultimately limiting its predictive power.

Under Techno-Social Monitoring the IT organization takes responsibility for recognizing when IT customers change their behavior before the change manifests in any formal communication or service level agreements. This is accomplished through a deep understanding of who the IT customer is and how they consume various IT services to satisfy any given business service, something which the existing methodologies we have surveyed do not address in great detail. The net result for Capacity Management is that capacity can be planned in accordance with a stated business plan and, concurrently with, the actual behavior of the Enterprise. Essentially this provides a "checks-and-balances" mechanism which ensures Capacity Management can be optimized from both the Business Activity and IT Transaction perspectives without unfairly biasing future capacity forecasts too heavily on the concerns from a single viewpoint.

The next three sections document the results of the project.

1. ACE Framework for Modeling and Analyzing

The Adaptive Complex Enterprise (ACE)[3] is a framework for modeling, monitoring, analyzing and improving enterprise functions and service-based interactions. The framework objectives are to:

- provide an integrated modeling and monitoring notation to create enterprise-wide models that are simple but complete for decision-making purposes
- model interactions between functions to surface the traceability disconnects that arise along interactions
- provide insight into what should be monitored and what information should be collected to analyze and address disconnects
- allow teams to represent an enterprise and use resulting models for assessments, for developing a roadmap, implementation, monitoring, and analysis leading to overall improvement
- allow teams to share knowledge yet work fairly independently

The next few sections show, for example, how this framework results in an enterprise-wide representation that allows us to identify and unify Business Service Capacity requirements that ITIL, TOGAF, and MOF have all recommend as being critical for good Capacity Management (refer to details of best practices in Chapter 3).

[3] This is based on previous Nationwide [John Northland] and CETI research [].

ACE Notation and Framework Application

Key Adaptive Complex Enterprise (ACE) framework [7] concepts and notations are introduced below:

- Complex enterprises can represent and relate their on-going Business-IT objectives by using the concept of a 'business domain'. Each business domain executes a primary transaction. And, there is an internal or *vertical BIOS value-chain* and *sense-respond* inter-relationship between Business and IT service functions as they execute the primary transaction and react to external events.
- Each domain handles a few external request types (e.g. commercial insurance, claims processing) and related transactions. There are important ways to characterize this external stimulus – routine (standard processing), non routine (non-standard processing), seasonal variation, and so on.
- A Line Of Business (or LOB) is also a business domain that is a collection of other related business domains. A LOB business domain might have a *horizontal* value chain of primary transactions: "sales →title search→claims processing". Each activity in this horizontal chain will have a vertical BIOS chain.
- For each request type, primary 'RED Service Transactions' are executed within the *service transaction execution* function of each domain using resources (human and IT). The execution provides its BIOS stakeholders with the performance views of the executions. A domain may support many secondary RED transactions types. A transaction is executed with its related request event. The transactions' strategy and business objectives are planned by other service functions.
- A RED service transaction and abstraction focuses on the *milestones* reached as a result of executing underlying business processes using resources, but *not on the details* of business processes or resources themselves. The focus of the transaction is thus on performance that results from interactions between providing services and consumed resources. When the transaction - called the 'RED (for requirements, execution, and delivery milestones)' – is executed, it provides metrics and BIOS traceability.
- Shared domains can exist also exist. An example is the IT domain.

An ACE model is also (when integrated with IT instrumentation) a dynamic model-driven business monitoring layer. Its objects represent on-going operational functions, the arrows represent on-going interactions, and instrumentation can populate the model for decision-making. Finally, model fragments can be developed by teams working independently and later connected.

Enterprise-wide Capacity Management using Cross-Domain Traceability

The application of the ACE framework to develop a enterprise-wide capacity management model is illustrated in Figure 4 and discussed below.

- *Modeling of Different types of Domains:* Key aspects of the business domains illustrated include functions that are aligned along the internal value-chain dimensions. Each domain also has a request-transaction-type specific business strategy supported by the information flow between these dimensions. The needed information is captured as attributes of the nodes in a model.

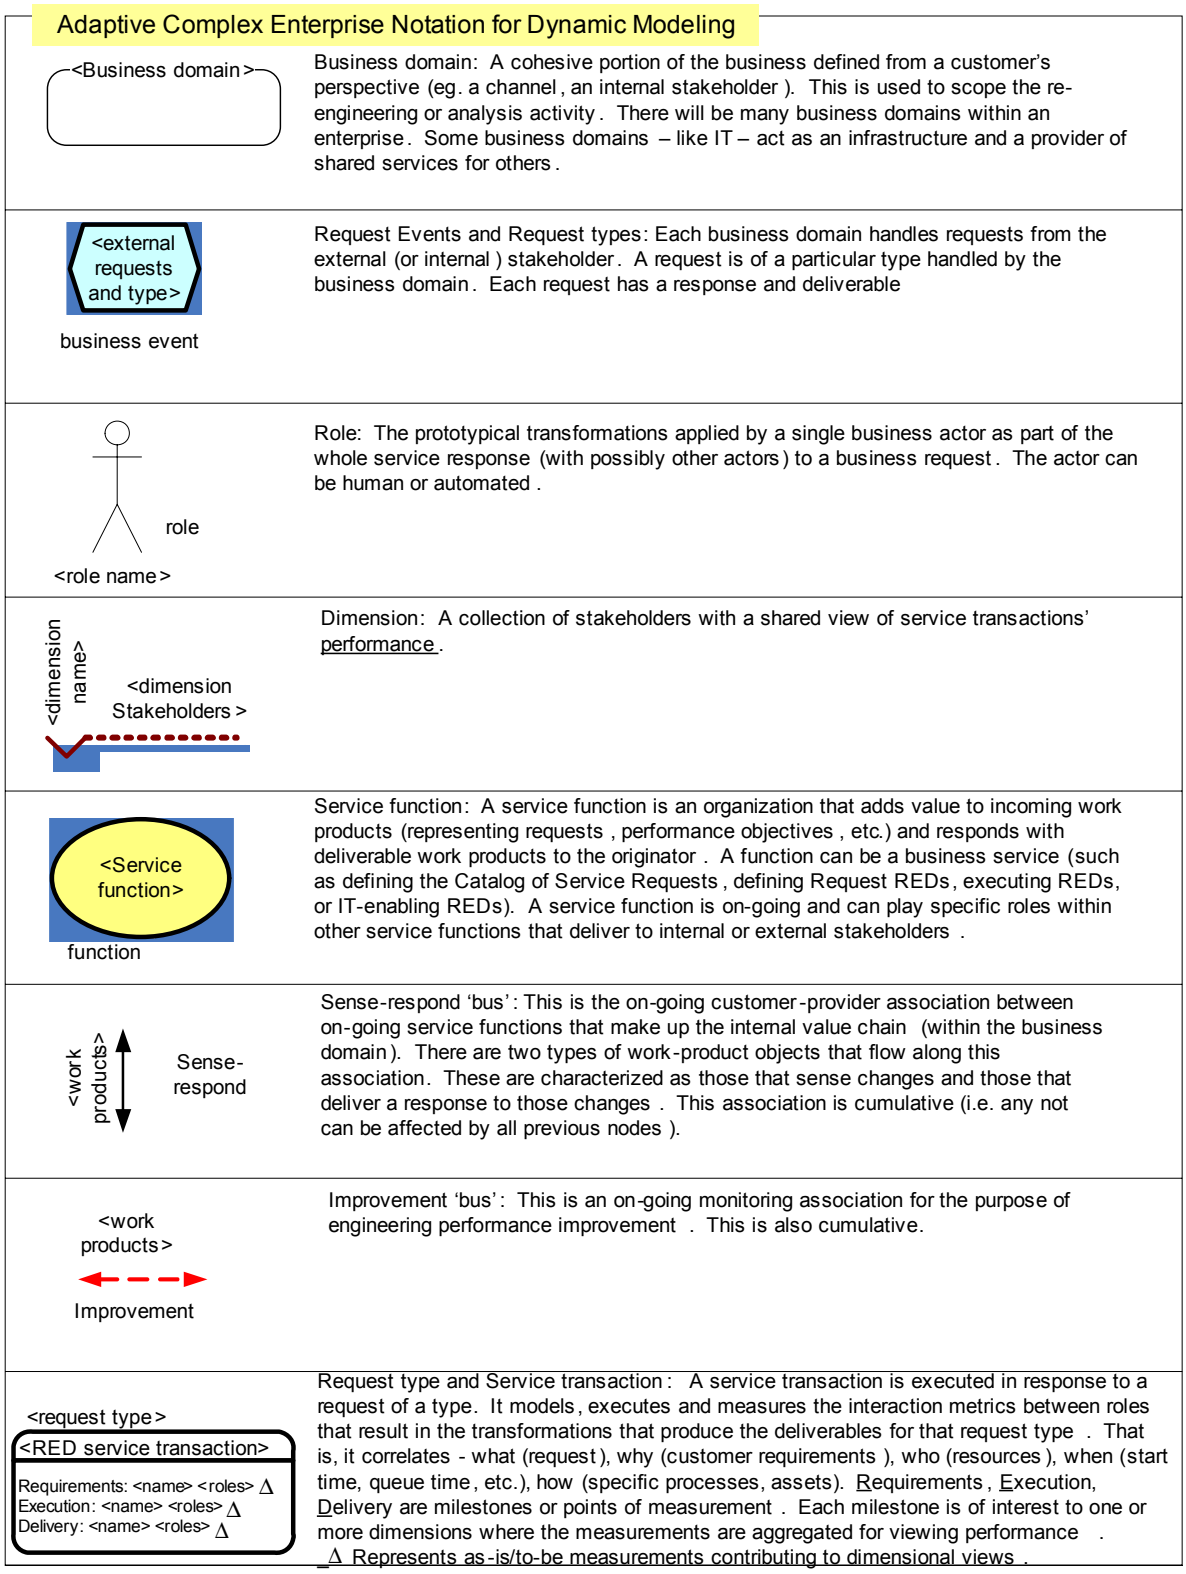

Figure 3: ACE Dynamic Modeling Notation

- *Certain domains are shared*. Specifically, IT supports the increase/decrease/retire/create anew the service execution capacity to respond to fluctuations in business service requests of the primary domains. Since the IT infrastructure is itself a business domain and is shared as illustrated in Figure 4, capacity management is no longer locally determined but is driven by different primary business domains that are external to the IT domain.
- *Business Improvement Domain:* (also illustrated) captures and consolidates the changes in functions and interactions between other domains for the purpose of improvement. In particular this domain captures the collective service requirements and business constraints on the IT services.

BIOS Traceability: Each vertical value chain typically consists of dimensions - BIOS stakeholder perspectives – with information that should be co-related for the creation of the Business Capacity Plan. This requires Traceability. We formally define traceability as the sense-and-response history of as-is/to-be (△) changes in model attributes (i.e. transaction metrics, derived performance indicators) along the model objects and associations (e.g. sense respond, improvement).

Note that the traceability is based on a two-way 'bus' association (represented as ⇔) where each node senses and posts responses for other functions. In the above example, the business work products are provided to strategy and the change of strategy is a response back to business. Thus, traceability is related to (and is a generalization of) the balanced scorecard [Kaplan]. There are also alternate views of traceability. One that is relevant for IT use is: **IT Resource △ ⇔Service execution △** (requests ⇔ execution metrics⇔ business process performance) **⇔Business services △** (business ⇔strategy)**'.**

Identifying and Addressing Traceability Disconnects: Traceability is achieved by using an ACE model first to identify interactions within and across business domains along which change must be communicated to the different dimensions and stakeholders. Attributes are next associated with the model elements to address different business scenarios that affect capacity and may take place within an LOB. For example, business domain△ in terms of increase/decrease/creation/elimination of a domain; or request type △ in terms of the number of service transactions within a domain.

The ACE model helps identify points where information is needed and disconnects (missing ⇔) between functions that need to provide the information for downstream use. Specifically based on the complete ACE model, we can now synthesize △ (as-is/to-be sense-respond metrics) information for function services and associations within each business domain and $\sum$ it across business domains to get overall impact to SLAs that IT is supposed to meet. However, as the model illustrates, in this case the transaction execution metrics that are used to develop SLAs for IT services are broken! Using the model these can be traced back to the lack of service strategy knowledge.

Examples of this related to capacity are illustrated by the 'X's' in Figure 4 identifying the source of missing traceability to information to be provided by the domains. Note also that the 'X' now pin-points the root causes of the disconnect that was identified at a high level in Figure 1 and the functions that are responsible for providing it. For example, such information helps with capacity to prepare for 'non-routine' requests triggered by disaster events. During such events, there is interference between business domains and increased needs for capacity that would not normally exist. In such cases, the traceability information needed along the business improvement domain and the BIOS value chain is:

Business constraints and investment ($) ⇔ IT resources ($\sum$ capacity used for routine and non-routine requests types) ⇔ Operational service execution capacity metrics (requests processed and quality) ⇔ Strategy (non routine, routine requests growth)

2. Roadmap

Based on the ACE framework we can now suggest the following tactical steps to achieve enterprise-wide traceability. This is followed by the more strategic research (no yet implemented) concepts.

Tactical steps

Step 1: Create Business-IT-Operation-Strategy (BIOS) traceability for each business domain

Develop ACE model with business domains representing the selected initial scope (e.g. LOB).

- At the highest level of LOBs, there is often related business domains organized into a 'horizontal' value chain. Starting with this subdivision of operational functions or business domains for the enterprise, each LOB is model is expanded with its vertical 'value-add' chains.
- For each ACE business domain create identify its functions and its transactions. Next attach the attributes for its BIOS traceability - identify projected routine/non-routine request volumes, related service transaction types and the value-add (Δ) to the customer and the LOB. Incorporate attributes for capacity management decision-making building upon on-going capacity management team's work and existing case studies.
- Identify where traceability is broken and what attributes are needed.
- Note: A transaction can be expanded into sub-transactions and at its lowest levels might eventually correspond to a data base transaction. But in most cases however, this is too much detail and would *not serve a useful purpose*.

Step 2: Assign Business Improvement Roles and Responsibilities to achieve to-be Traceability

Define the Business Improvement domain functions. Assign roles and responsibilities to the functions to gather information and improve traceability information flow (e.g. reports, metrics).

- Include responsibilities such as estimating the human/system cost of providing this information. Answer questions like - are we spending the right amount of effort in monitoring the right type of things to result in real benefit to business? For example, what is the cost of modifying an existing system to provide better information versus the real value gained?

Step 3: Implement CDB to provide BIOS Traceability

Implement CDB Capacity Management Database – a critical component (identified by best practices) to share information needed for Enterprise-wide Capacity Management.

- The enterprise must understand both externally-driven changes to LOBs and service utilization as well as internally-generated incidents related to asset capacity. This must also include aspects of disaster recovery and resilience.
- The MOF details what CIs/attributes and columns might exist in the CDB [see 3]. Based on these CIs/attributes, provide a CDB that is a ACE logical model for BIOS traceability as follows:

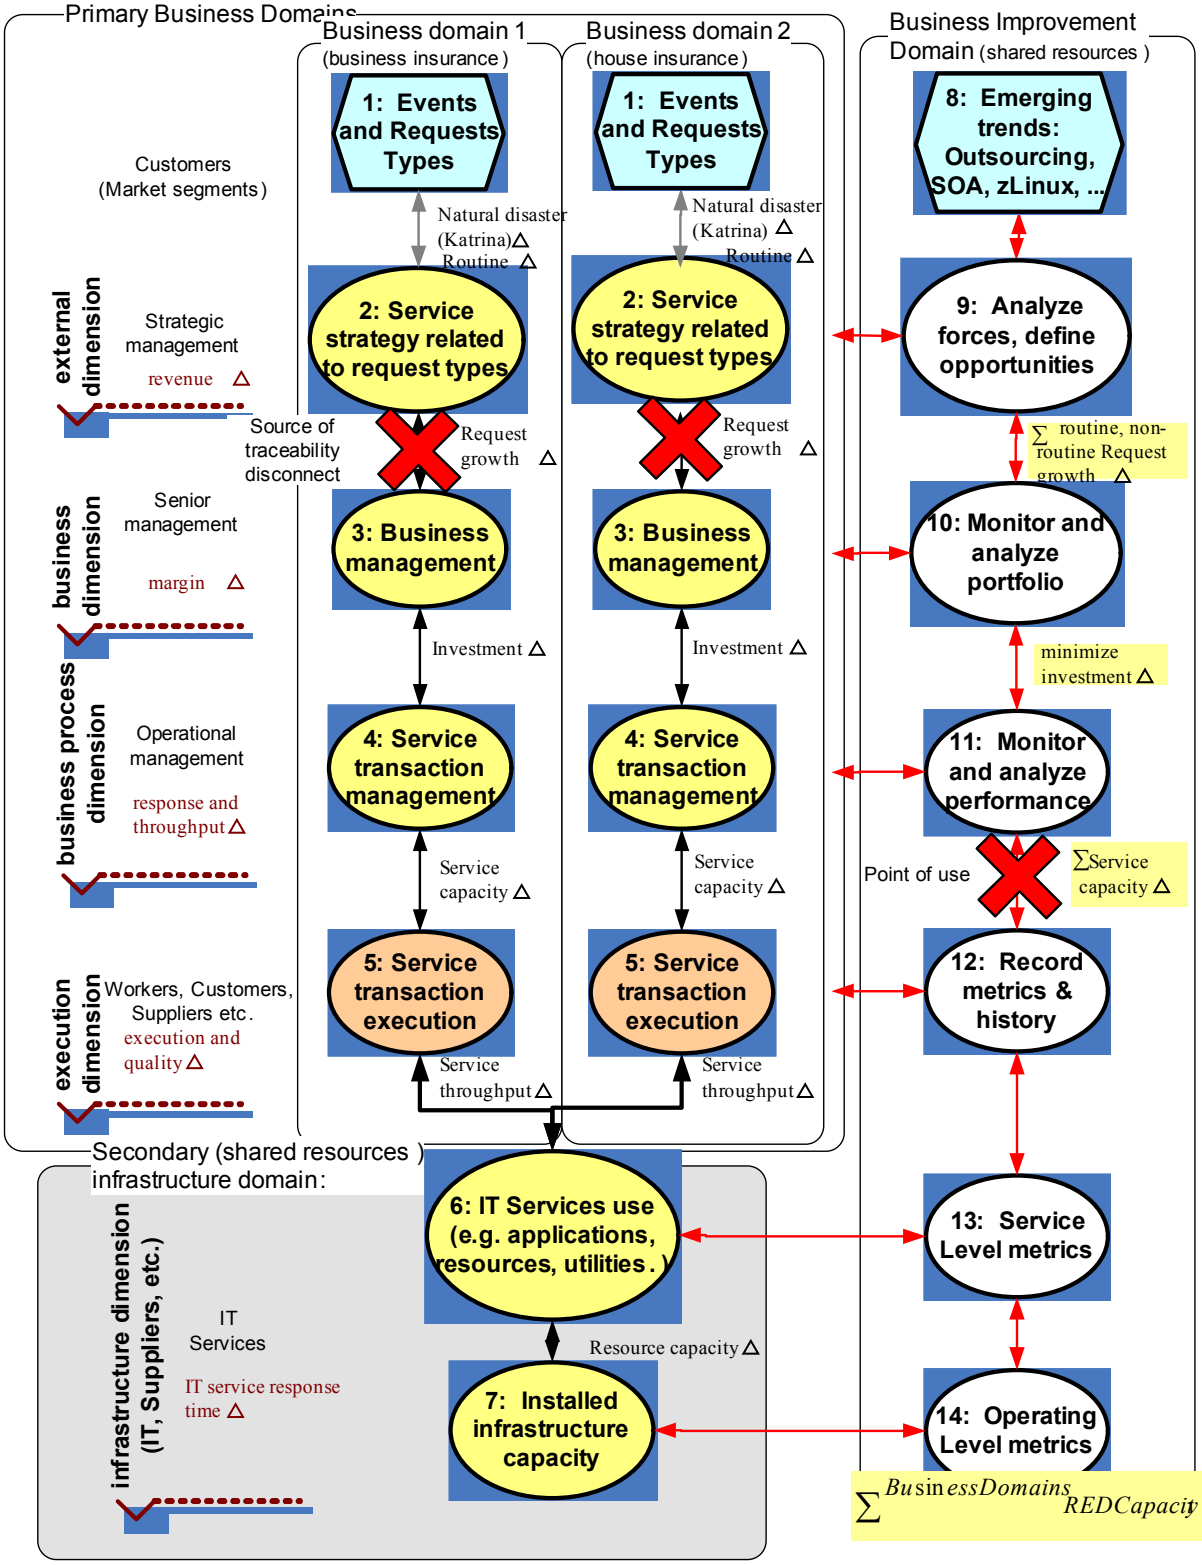

Figure 4: Enterprise representation with more than one business domains, shared domain with IT service functions, and improvement domain. The disconnect – lack of information needed to for improvement is more precisely defined.

1. *Predict future needs:* This is through BIOS traceability. Relate Configuration Items to Service Transactions they support, and those, in turn, to the Service Strategy related to the request types (i.e. relate functions as in the Business Domains). This enables prediction and resilience.
2. *Identify on-going operational needs*: Secondary IT processes such as incident, problem and change management handle any problems with the current IT infrastructure capacity. For example, an incident may affect asset configuration - a change in a network component and traffic may impact the capacity. Capacity availability may be affected for the time it takes to resolve day-to-day incidents.
3. Provide monitoring at a high level for rapid decision-making and improvement: This is at the abstraction of the ACE model. (This is covered in the next section.)

- Create CDB by associating *Assets* to *Service Transactions that use them* and then back to any *incident transactions* that manifest with those services or assets. That is, in addition to capturing configuration relationships between assets, services and the RED transactions that occur while delivering a specific service, include the capability to capture RED type transactions for Incident management as well as Incident resolution. Whereas the actual structure needed to model IT assets and the services they support is fairly simple, consisting of an Asset table, a Service table and a Configuration Item table, research has shown that all the association between these during enterprise operations, that is the *socio-technical* patterns, impact overall capacity availability.
- An example of inclusion of all the related aspects of operations, like Incident Management and SLA compliance as the starting pieces to map the socio-technical activities to the IT assets and services is in Figure 5 below. This sample, and simplified UML diagram, illustrates some of the common entities involved in the incident management process and shows their relationships to transactions. Some of these entities have been related with associations, denoted with dashed lines, to the corresponding RED transaction for incident management. These associations help establish the traceabilty between one transaction type and other but more importantly elicit the informational requirements of the CDB as a monitoring and improvement tool for capturing the value of the incident management process vertically throughout the enterprise. Specifically, the diagram illustrates that the CDB must be capable of representing the performance of the incident management process both in terms of a single incident, reported though metrics in terms of a *user* and *incident,* and the incident process, in terms of a *customer* and a *user.* In this manner the scope and detail of the CDB items can be justified with respect to the entities, the *customer* for example, that will use them to monitor the performance of the transaction in a meaningful way. Without such an explicit means of identifying items that should be present in a CDB it is difficult to ensure that the metrics are representative of any truly meaningful interaction between business dimensions, that they do not go into excessive detail, and that they can maintain traceability between various units in the business.

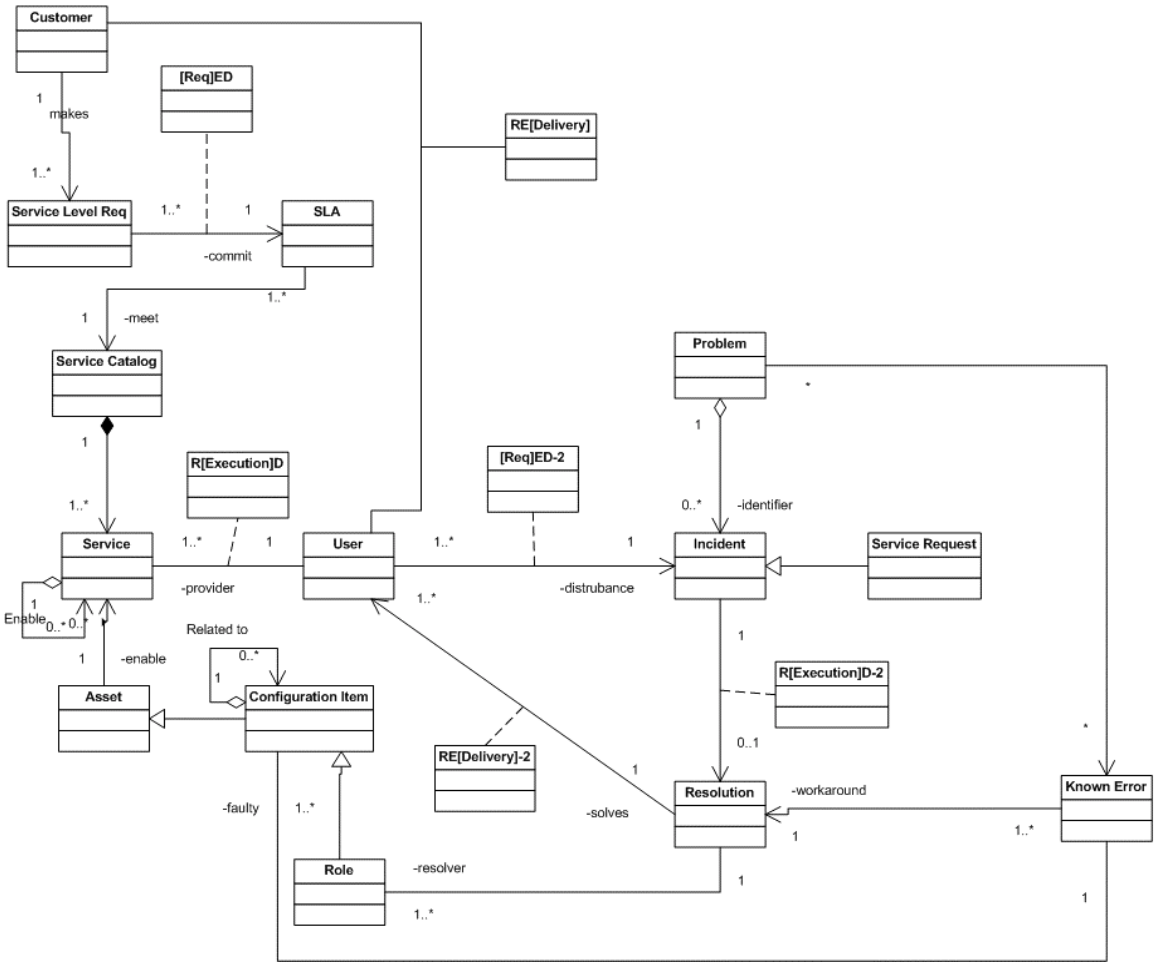

Figure 5: Conceptual Capacity Management View of the Configuration Management Data Base –

Legend: RED Transactions
[Request]ED, R[Execution]D, RE[Delivery] → A transaction corresponding to the IT department offering a complete service to a business entity, or customer. (Inter-dimensional: between IT process and business execution)
[Request]ED-2, R[Execution]D-2, RE[Delivery]-2 → A transaction corresponding to a single incident (Intra-dimensional: IT process only)

Strategic Step

Step 4: Monitoring and Socio-Technical Technical Traceability

Develop Socio-Technical monitoring so that much of capacity management can be accomplished as a side effect (i.e. with minimum impact to business operations) and with minimum of human labor.

1. Develop monitoring and reporting tools at the level of the ACE abstraction.
    - Monitoring externally driven change: Conduct research to determine ways in which socio-technical interactions can be mapped, especially for non-routine requests and considering aspects like seasonal variations. Develop ways in which these maps can be co-related with application-to-application traffic using underlying application monitoring tools to determine overall relationships between business activity trends and IT use.

- One possible approach is to be researched is to enhance sign-on components with user intent and roles corresponding to the business domain and the interaction/execution pattern captured. For example, a sales transaction causes more application services to be used than a claims transaction, or the user that is a channel may move/process hundreds or records, whereas internal sales will work with just a record at a time.
- Another indicator of socio-technical data that cannot be overlooked is information that can be summarized through security (log on) records. By providing analysis of the roles played by the individuals logging in along with the frequency and timings of log-on, it is possible to manage (and audit) IT resources by 'side effect'.
- Viewed differently, seek news ways of monitoring SLA and QoS requirements for different business services and relate then to operating level requirements from capacity and other user intelligence mining perspectives. Correlate the underlying Business-IT fluctuations more automatically.

Thus, once it is recognized that it is possible to use standardized transactions as a basis for IT capacity management and planning, we can use these transactions as logical markers or anchor points throughout the enterprise for analytics. Just one example being the overall asset utilization through macro levels of QoS analysis. Using the change events described earlier, (Increase/ Decrease/ New/Elimination of LOB/Business Domains, Transactions) it is now becomes possible to automate the information gathering of the role of IT capacity in relation to the possibility of these events.

3. Current Practices and Technologies and Analysis

Although technologies exist to provide on-demand data and insight to complex business activities, most enterprises continue to struggle with Capacity Management. There are several contributing factors:

- Difficulty in making the organizational changes for communications and flows of information needed to make more effective use of the technology and related current industry best practices.
- Lack to training in systems thinking that goes beyond program management.
- Industry practices and technologies that are either too generic or attempt to force fit business activities with technology rather than aligning technology to business activities.
- A lack of prescriptive practices and techniques for 'management by side effect'. That is, simple ways to gather needed data without making major changes in operational processes.

There exist several best practice guides, or frameworks, technologies and tools for which Enterprises may use to assist their Capacity Management and Planning. However, we find these frameworks tend to be descriptive and not prescriptive in nature. Although it makes them more universal, it also leads to confusion about what to take from these frameworks and how to apply it to an individual Enterprise. With this in mind, there is still great value in the current best practices and what they describe. In the following section we shall outline the pieces of these best practices that we found to be most critical to the success of Capacity Management and assemble them into our Framework.

Current Industry Practices

Information Technology Infrastructure Library (ITIL)

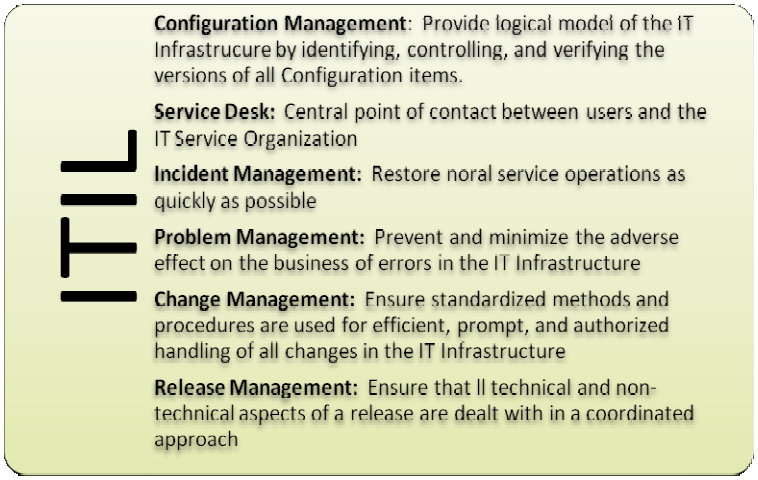

[ITIL Awareness – Greg Hines, Hines Consulting Group, Inc.]

The ITIL Foundations describe the IT Service Organizations that deliver the agreed upon services and maintain the infrastructure on which these services are delivered. What ITIL provides the enterprise is common functions and a high-level view of the organization, and also provide a common language for IT Service Management Processes [4]. With ITIL's focus on IT Service Delivery and Support of those IT Services, there are several aspects of the ITIL Foundations that are tapped for resources while others are not. The primary areas of the library that are used by organizations are detailed above.

In addition to these generalized components of Service Management, there are also specific areas located in ITIL that discuss the ways and means of managing capacity for an enterprise. Specifically, ITIL specifies that there are inputs, sub-processes and outputs to the capacity management process which feed off of each other in an effort to provide the necessary answers for the Enterprise to manage their IT utilization. Within each of these "phases" of Capacity Management, critical components and processes are detailed which, in the ITIL realm, create the situation in which good Capacity Management may be achieved. The specific details for each of the "phases" of Capacity Management for ITIL are detailed below.

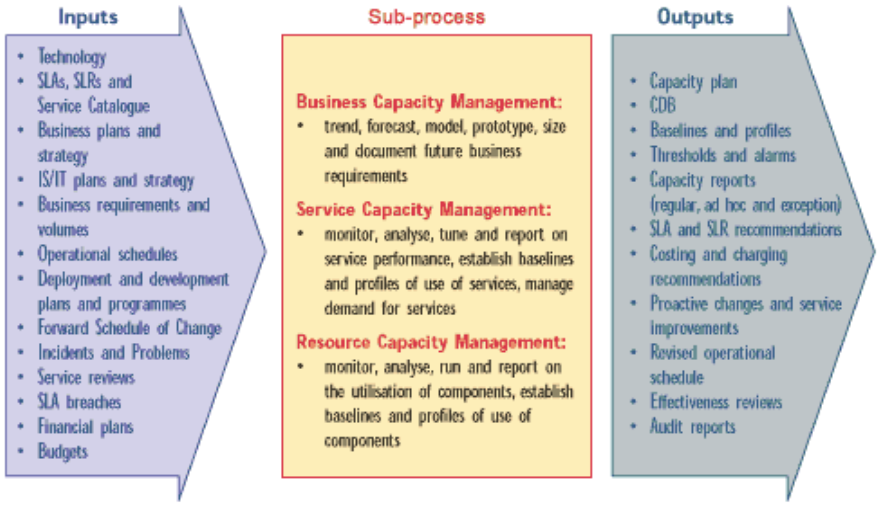

[ITIL Service Delivery – The Stationary Office of OGC 2001]

Upon close examination of each of these "phases", it is apparent what we mean by a descriptive and not prescriptive framework. For example, in the Sub-Process phase of management, listed for each of the areas of Capacity Management (Business, Service and Resource), items

like trend, forecast, model, monitor and analyze are listed. However, exactly what is supposed to be analyzed, monitored, and forecasted is not fully specified. Although the ITIL Foundations fall short in delivering a prescriptive framework, they do provide insight into some of the key components that can make Capacity Management successful in the Enterprise. The critical output components that ITIL brings to the stage are illustrated in the following figure. Again, we find that though entities like the CMDB are known in concept, there is much less consensus with respect to how they should be architected and maintained in practice.

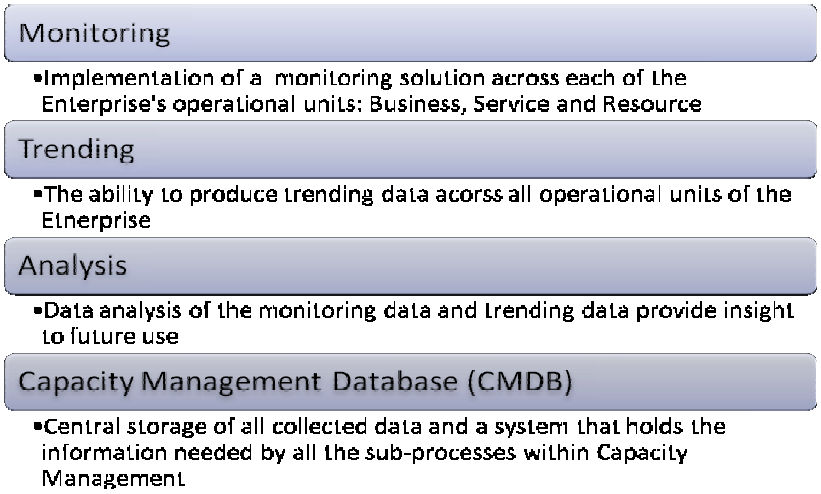

The Open Group Architecture Framework (TOGAF)

The Open Group Architecture Framework (TOGAF) is the combined views and solutions to enterprise architecture from industry and government representatives such as IBM, Microsoft, and the US Department of Defense. Specifically, TOGAF is the evolution of the US DoD's Technical Architecture Framework for Information Management (TAFIM) [5]. The end result of this evolution is the Architecture Development Method (ADM) detailed below with each of its respective steps. [Source for SAP/Gartner])

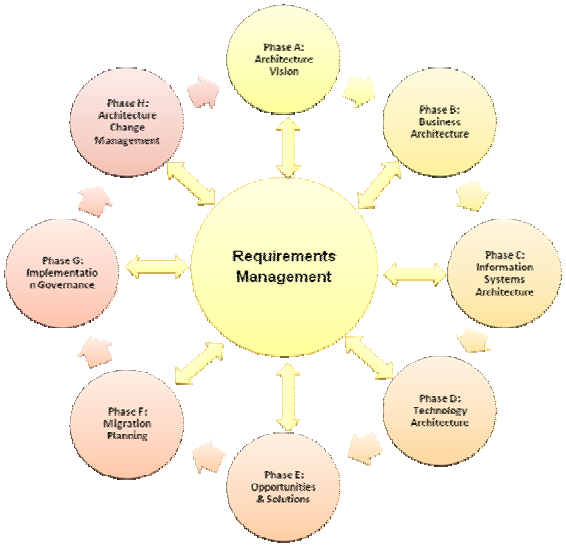

[TOGAF Model]

TOGAF is not designed to be exclusive to one modeling process or one type of notation, but is centered about the belief that architecture for an Enterprise needs to incorporate all relevant personnel and that everyone involved is committed to the success of the architecture process. What this provides the Enterprise is a high-level understanding to the architectural goals for the organization, a clear and defined "footprint" for the organization, and finally emphasizes the need to set up and monitor a process to confirm the purpose of the defined framework.

Additionally, what TOGAF does, not noted in other frameworks, is to clearly state that it is not designed to be implemented on its own. Alternately stated, TOGAF is designed to work with all existing frameworks and in fact, recommends doing so. They recognize the fact that there is no one unifying framework to solve the various difficulties the IT Service Enterprises face in day-to-day operations. The model for TOGAF, more simply put, is to define the architecture across the Enterprise and implement that architecture using the most appropriate framework for your specific situation. So, from TOGAF, we are again left with a descriptive framework that advises one to use whatever is best suited to get the job done and not a prescriptive path to achieve good Capacity Management. What we do find in TOGAF that supports our recommendations are the following critical components.

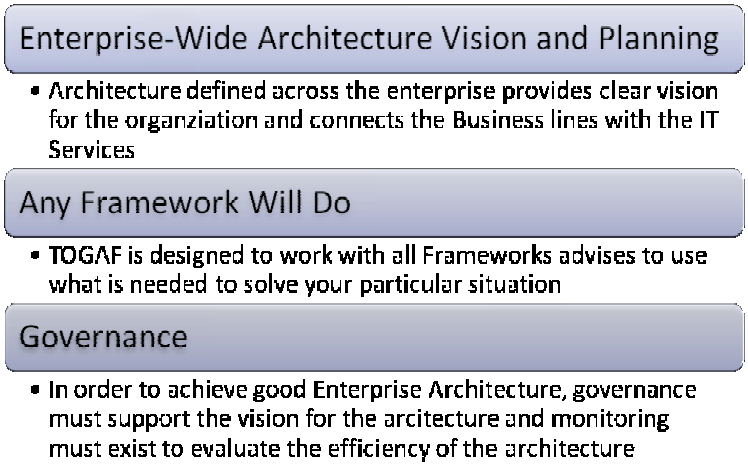

[TOGAF]

Microsoft Operations Framework (MOF)

Microsoft contributes their collective knowledge to the world of Capacity Management through their Operations Framework. The MOF is aligned with the ITIL Service Delivery book and represents an attempt at a prescriptive approach to implementing ITIL. Since the MOF is aligned to ITIL, there are some of the similar components that ITIL identifies as being necessary for good Capacity Management, such as Business, Service, and Resource Capacity Management. The Capacity Management process cycle is depicted below and is a combination of some of the items already identified as critical components of Capacity Management, namely Analysis, Monitoring, and Optimizing. Adding to our existing list, the MOF recommends Modeling and Change Initiation as essential pieces of the process.

What the MOF provides that ITIL does not, is a guide to the Quality of Service (QoS) metrics that need to be captured throughout the Capacity Management process in order to manage the quality and effectiveness of the overall process [6]. In addition to the QoS metrics that the MOF details in their Capacity Management guide, we are introduced to the importance of Service Level Agreements (SLAs)

for IT Services and how these play a critical role in Capacity Management.

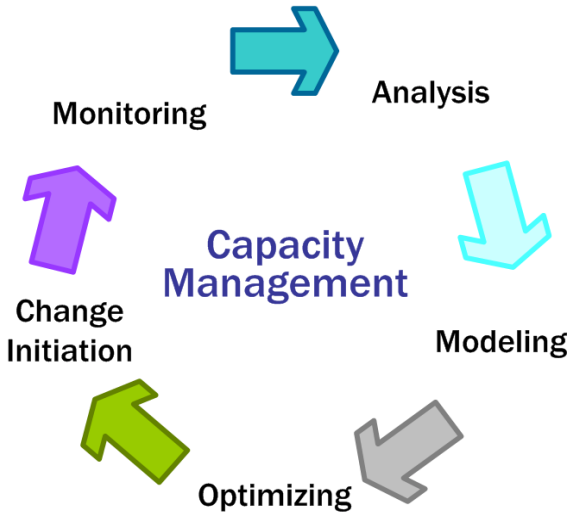

[MOF: Capacity Management Process Model]

Business forecasts:

- Timely production of workload forecasts.
- Accuracy of business trend forecasts.
- Incorporation of business plans into the capacity plan.

Technology:

- Monitoring of performance and throughput of all services and components.
- Deploying new technology in line with business requirements (time, cost, and functionality).
- Maintaining SLAs regardless of problems with support or performance of old technology.

Cost-effectiveness:

- Reducing or eliminating panic buying.
- Limiting over-capacity that cannot be justified in business terms.
- Accurately forecasting planned expenditure.

Alignment to business need:

- Reducing the number of incidents due to poor performance.
- Reducing the amount of lost business due to inadequate capacity.
- Implementing new services that match service level objectives.
- Acting on the recommendations made by capacity management.

[MOF: Capacity Management QoS Areas]

Capacity Management in the MOF guide works very closely to specify, and refine, requirements resulting from SLAs, which in turn drive the operating level agreements (OLAs) for a particular IT Service [6]. This, in essence, is the connection of Lines of Businesses (LOB) with SLAs to their supporting IT Services and assets (OLAs). In addition to making the connection between Lines of Business (SLAs) and the IT assets (OLAs), the MOF guide also attempts to articulate what the CDB would have in it. An example of a typical CDB table, as advised by the MOF, is listed below.

Field Name	Description
CapacityItemName	Name of item
CapacityItemType	Type—for example, service or resource
CapacityItemDescription	Free text description
FunctionOfMeasurement	Measurement functions—for example, kilobyte, milliseconds, or megahertz
FunctionCost	Standardized financial cost of each function
MaximumCapacity	Maximum possible capacity achievable by item
AlertThreshold	Current threshold at which an alert is first generated
ChildHierarchyItems	List of capacity items that this item depends on
ParentHierarchyItems	List of capacity items that depend on this item
TrendDataTimestamp	Timestamp when trend data is collected
TrendData	Data value captured

[MOF: Capacity Management Database example]

Where ITIL and TOGAF fall short in being prescriptive, the MOF Capacity Management Service Management Function guide attempts to more clearly articulate the path to Capacity Management through recommendations developed both in house and through standards bodies like ITIL. By aligning their strategy to ITIL, the MOF guide not only rides the wave of industry standards, but also feeds back to the standards body by contributing the in-house processes and procedures that they develop. The critical components that we can take away from the Microsoft Operations Framework are as follows:

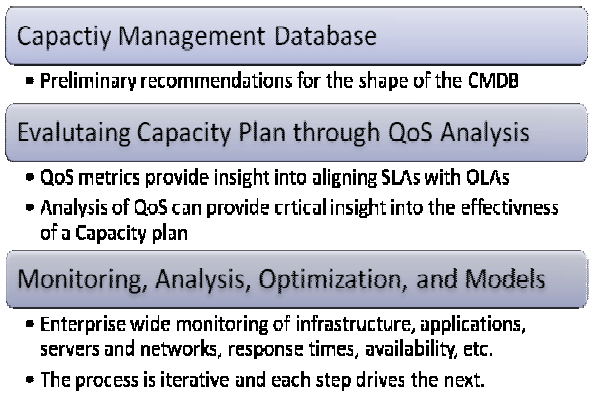

Current Tools and Technologies

Tools and technology constitute the other side of the Capacity Management processes, the pairing of tools and technologies with best practices provides a solution that Enterprises can implement. In the following section we will discuss some of the critical tools we found that could be used to facilitate good

Capacity Management. However, similar to the preceding analysis of best practices, these tools on their own do not avail the Enterprise to successful management of their IT Capacity. We submit that the complex nature of the IT Services interactions with IT Assets is not easily represented or captured by tools alone, nor do Enterprises deploy these solutions in the optimal way. Though the tools and technologies by themselves, provide some of the critical components previously identified for successful Capacity Management, they do not provide a comprehensive vertical and horizontal views of the organization that is needed.

Application-to-application Monitoring

Several tools on the market perform monitoring via standard logins and protocols. One example is Mercury SiteScope [http://www.mercury.com/us/products/business-availability-center/sitescope/]. Another example is InsightETE - a monitoring tool that captures packet data transmitted across the Enterprise Network. The monitoring tool extracts data from these packets that is then reported back to a central analysis database. From the analysis database, statistics and other data can be presented to the user both graphically and in a datasheet. It is an adaptable solution that can be modified and extended to process custom, or new, protocols per the needs of the Enterprise. The InsightETE solution is passive network analysis, which is to say that there are no agents installed on assets within the enterprise. This makes for a monitoring solution that is effective, easy to manage and easy to deploy. Additionally, a fascinating aspect of the InsightETE solution is its ability to track and monitor transactions across the IT Infrastructure. This provides the ability to view, on a transactional level, the daily operational signature of the organization and perform analysis and comparisons to changes in business activities. A typical InsightETE deployment looks like the following.

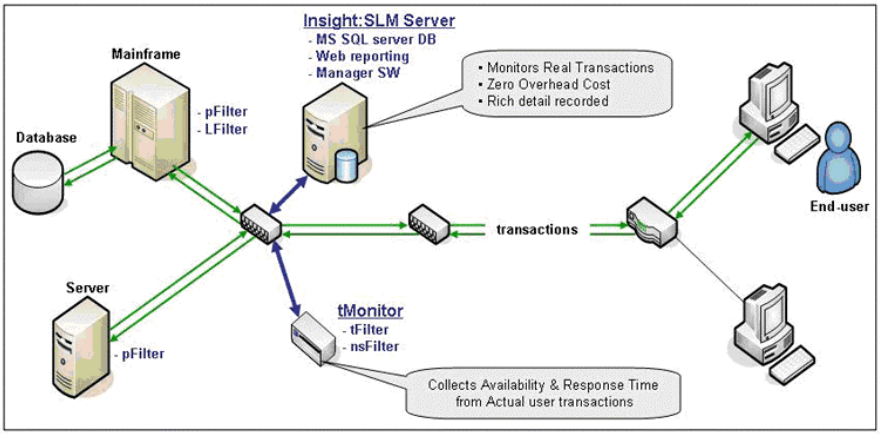

[InsighETE Solution Architecture]

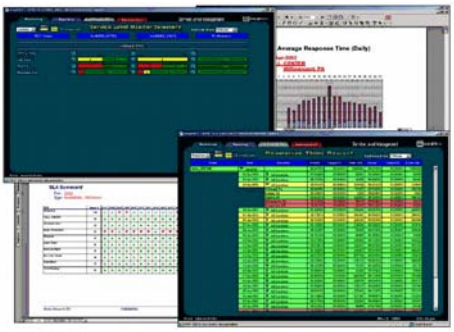

Similar to the BI applications previously mentioned, the InsightETE solution provides custom dashboard views of the data that it collects; a growing need and desire for the Enterprise commfunctiony. As a result of the capturing of transactions across the infrastructure, the Enterprise can perform complex analysis and comparison to determine changes in business activities that may affect available IT Capacity.

[InsightETE Dashboards]

InsightETE could be augmented by other agent based solutions which would provide the other pieces of the puzzle to true capacity.

Business Intelligence Suites

The market space for Business Intelligence suites is booming. Gartner describes these as Business Activity Monitoring platforms and feels that the market will continue to grow through 2010 [3]. BI applications are a combination of data reporting, data visualization, and data analysis tools deployed under one hood. There are several of these applications available from companies, such as Microsoft, Systar and IBM, each with their own twist on the same thing: Understanding Business Activities leads to successful management of the Enterprise. As with any product, there are differences in what and how data is delivered to the user, but fundamentally what they do is very similar. Some of the more interesting features to these applications are their ability to project trends, forecast performance, and most importantly, connect operational function both horizontally and vertically across the Enterprise. Advancements in data mining, data warehousing and client side rendering capabilities are allowing servers to perform more complex calculations and provide a richer viewing experience for the Enterprise.

TIVOLI

IBM's Tivoli product suite is a particularly good example of a mature set of tools designed to support businesses in this regard. These products have closely followed the development of best practices outlined in IT service management frameworks, such as ITIL, but have more recently expanded their scope to include service management in its entirety throughout the business. This focus underscores the importance of coordination throughout the vertical dimensions of the business when attempting to support those service management functions that have traditionally been seen as isolated within IT groups, including Capacity Management[cite - ftp://ftp.software.ibm.com/software/tivoli/analystreports/ITSM_Vendor_Report_Card_-_Turner_12-06.pdf].

Below are some examples of the types of views of data that these applications provide. User built dashboards and scorecards are provided, service delivery views tailored to your business objectives and infrastructure analysis views to connect monitoring data to the respective assets.

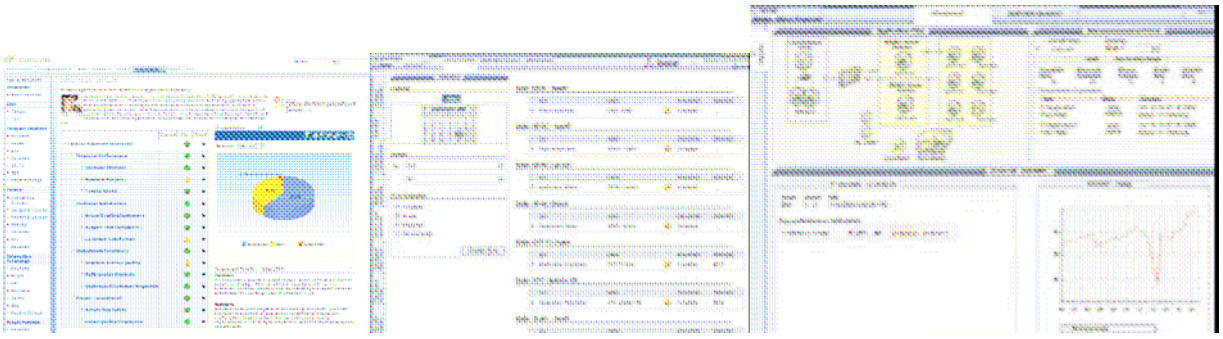

[Microsoft Office PermformancePoint Server 2007 and BusinessBridge SolutionVision by Systar]

Although these applications tend to be geared toward a specific type on Enterprise operation, for example sales or marketing, the newer architecture of these applications allows for extending the functionality of these solutions to work with any aspect of the Enterprise's operation, including Capacity

Management. In other cases, the out of the box solution simply needs to be attached to a data warehouse, have some views built from the data, and the Enterprise is off running. With the well thought architecture of the next generation of BI applications, the potential to use these technologies as solutions for Capacity Management is increasing.

Individual best practices and technologies in Capacity Management provide insight into potential solutions to these areas of capacity change, but fail to be integrated and prescriptive enough to implement consistently and successfully across all areas of operations and management. A prescriptive and holistic Capacity Management methodology must unify the three perspectives – business capacity, service capacity, and IT resource capacity - such that the capacity requirements of the business are clear and visible to all stakeholders concerned with various aspects of the enterprise. Currently, the best practices and research encompassing Capacity Management either fails to address all three distinct perspectives or is merely descriptive, and not prescriptive, in the explanation of methods for managing capacity within and throughout all three areas of concern. For example, Porter's Model identifies external competitive forces will drive change in the enterprise, but stops short of the steps for managing that change. Similarly, Capability Maturity Models (CMM) and the RED transaction model suggest that an examination of processes, combined with continuous improvement cycle of those transactions, can provide insight as to the assets, resources, and sub-transactions used to sustain a given type of business service, but falls short in an explanation of what to do with collected data or how to project capacity needs from the collected data.

As a result of the lack of prescriptive integrated approaches to Capacity Management, enterprises are left to fend for themselves and attempt to manage their capacity through a disparate collection of "best practices" and technology resulting in a poor and fractured understanding of their true capacity needs and capabilities and a general inability to respond, under a reasonable budget and timeframe, to business activity changes within their industry.

Analysis

Challenges in Aligning Differing Business Strategies to Service Level Goals

Lack of Consolidation Across Business Domains: Strategies for operational functions within an enterprise tend to be silo-oriented and don't span horizontally across other functions. Compartmentalization of planning can result in operational domains providing incomplete or conflicting service-level objectives to the IT services domain. An organization for cross-function business planning and the timely dissemination of relevant capacity-related knowledge becomes a necessity (see Improvement business domain in Figure 4). By instituting enterprise-wide business improvement, this organization can consolidate the strategies for separate operational domains to achieve a greater understanding of the enterprise's business activities and provide visibility to all parties. Specifically for unified Capacity Management the goal is to consolidate the service level objectives to become fluid in their daily operations and to gain the capability to adapt to key business activity changes. However there are no good communication tools that will surface and address the issues that relate business challenges to service-level and operating level views. The ACE dynamic modeling notation addresses this requirement to provide a conceptual architecture for requirements and implementation.

Lack of Socio-Technical Service Level Monitoring: Another problematic trend in Capacity Management lies in the technology deployed by enterprises to assist with their Capacity Management activities. The trend for these products is to focus on a fixed set of assets (or configuration items of hardware and software) as opposed to an evolving IT environment, and to focus on the utilization of these fixed assets. For example, a prototypical monitoring solution may examine a fixed set of IT assets,

as programmed and configured by a network administrator. However when a new asset is added to or removed from the environment, prompt intervention by a network administrator is required or the set of monitored assets may drift away from reality. A related problem with modern monitoring solutions is that although they may have the ability to monitor internal applications and services providing limited but valuable knowledge to capacity managers, they lack the ability to make correlations in application and service level trends with business trends easily visible to interested stakeholders. One solution might monitor applications by examining available memory, disk usage, processor spikes, request patterns, etc., but this still does not capture the full details of how and why a particular application or service is being used by customers. This type of information becomes especially important in the context of shared application services. Documenting this type of information could be extremely valuable but can also be prohibitively expensive. The central point is that the motivating reasons for which applications and services are used changes over time, and it is exactly this and not usage alone, that is critical to holistically understanding and managing capacity at all levels of concern. We call these motivational forces the *socio-technical* characteristics of the business and posit that understanding them through each of the perspectives of capacity management is the cornerstone of a prescriptive unified Capacity Management framework.

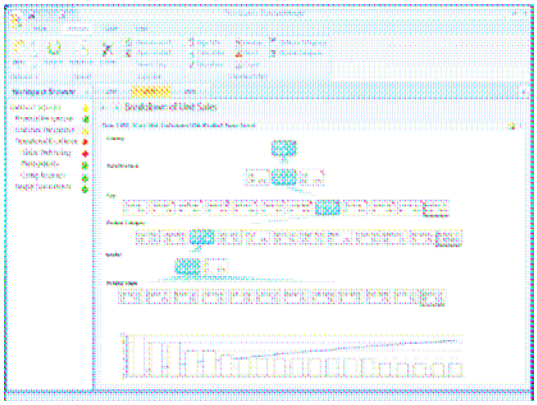

Lack of Connection between Past Micro-level Uses with Future Macro-level Prediction: A common practice of current Capacity Management, is the collection of micro-level transactional data in an attempting to align IT usage with on-going business activities. These data elements are collected to provide insight into the past capacity utilization for an organization but, on their own, do very little to explain the "what and why" factors that cause available capacity to change throughout the life of the Enterprise. The micro-level events are captured but often can't be easily or directly correlated to the high-level business activities and macro-level transactions. Typical pieces of data that are collected in this model are memory usage, disk usage, processor usage, bandwidth utilization, etc.

What is needed is a good understanding and view of the macro-level events that will provide insight into the socio-technical changes of the business activities within the Enterprise. This is not to say that increases in memory usage or increases in disk usage don't point to changes in the IT environment, but what is becoming more evident is that these items alone do not produce results that facilitate good Capacity Management and Planning. Therefore, we would like to suggest that another critical component to holistic Capacity Management is the understanding of and the capability to view real-time changes in the business activities within an enterprise and to understand what impact these socio-technical forces have on capacity requirements. For an organization to be successful at managing and planning their IT Capacity, enterprises must be able to align these activity trends down to the IT assets used to support them.

Lack of Application of On-demand Visibility and Data: Business Intelligence (BI) suites have been gaining in popularity for use by Enterprises in a multitude of areas. The use of these applications has primarily been focused on business performance. The use for business Capacity Management is often indirect. Many BI suites could be useful for the forecasting activities of a Capacity Management strategy, however, their focus on specific business activities, like sales or production, makes holistic Capacity Management throughout all the related perspectives an expensive task. Nonetheless, what BI

applications have done for Enterprises, and the business as a whole, is provide them with on-demand visibility of data about their performance. Dashboards, KPIs, and other techniques present in BI suites and are quickly becoming standard pieces of the CIO's arsenal for communicating, disseminating, and presenting the performance of their Enterprise.

Customized views of Enterprise data from within these suites allow practitioners to gain insight into their particular area of concern, and thus the ability to view the Enterprise from multiple vantage points is becoming easier. Horizontal and vertical views are quickly presented to the user as the BI applications provide drill down, across, and up to link together the various pieces of the Enterprise used to process any given transaction. One such example of this type of BI application is Microsoft Office Performance Point Server 2007, which provides views, like the one shown here, provides a visible representation of the connections between horizontal and vertical functions within the organization. Additionally, some BI products provide a set of analytics for forecasting, custom dashboards and all the other features that Enterprises need for connecting the business activities to their respective IT assets. The difficulty for most BI suites is the data they run on. BI suites cannot make reliable projections from inadequate data, and the best forecasting and data visualization tools can only forecast and render from the data that is provided.

A bridge is needed to enable Enterprise Capacity Management by connecting the monitoring tools which collect data and the BI suites which operate on the data. One such bridge is the ACE representation and conceptualization integrated with Configuration Management Database (CMDB) to provide a holistic Capacity Management strategy. A general architectural framework for Capacity Management can be constructed from monitoring solutions to collect data, a CMDB to relate it with IT assets, and ACE conceptualization-based BI implementations to facilitate analysis, projection and trending. The combination of these three components is what Gartner refers to as Business Activity Monitoring (BAM) platforms. Furthermore, what Gartner finds in their research is that "[these] platforms are adaptable. They are not locked into supporting any specific application or event source, but can be used to monitor almost any form of business activity [3]". With these expanses in the BI market, a technological solution is beginning to come forth from software companies and moreover, the implementation of these technologies is becoming easier for Enterprises to deploy and use to manage their business activities.

Prohibitive Cost of Enterprise-wide ACE Management Methods: In general the investment in an improvement domain for a single requirement might be expensive and appear to be prohibitive. However, by enabling many other improvement requirements (e.g. Sarbanes-Oxley compliance, help desk, new business services) an enterprise can more easily justify the improvement domain.

Bridging the Gap through the use of Practices and Technology

An analysis of the tools and technology seem to indicate that technological solutions exist, but are alone insufficient for most Enterprises to provide a sufficient Capacity Management strategy. The acceptance of best practices provides a common understanding to bootstrap Capacity Management initiatives but again for many Enterprises managing such a large number of assets, fully connecting business functions, and deploying monitoring in legacy environments extremely difficult to accomplish successfully in practice.

Our proposed roadmap in the next section is based on the integration of best practice, technology, and business process enhancements to facilitate optimal Capacity Management for any enterprise. However, unlike existing frameworks, we will sacrifice some generality in order to be prescriptive with respect to critical enablers of Capacity Management. Although incomplete in its current form, we hope

that this will provide the foundations for construction of a framework that will provide more than a shared understanding but also a shared plan of action.

4. Bibliography and References

[1] “Improving Business Functionality Through Model-based Service-Oriented Business Applications” - Ronald Schmelzer, ZapThink!, January 2007

[2] “An Introductory View of ITIL” – itSMF 2004

[3] “MarketScope for Business Activity Montioring Platforms, 3Q06” – Gartner Research, August 16, 2006

[4] “ITIL Awareness” – Greg Hines, Hines Consulting Group, Inc., July 2005
ITIL Service Delivery – The Stationary Office for OGC, 2001

[5] “ACE TOGAF” – Jay Ramanath, CETI, March 2007

[6] “Microsoft Operations Framework: Capacity Management Service Management Function” – Microsoft Corp, January 2005

InsigthETE – http://www.insightete.com
Business Activity Monitoring by Systar: BusinessBridge Service Vision - Whitepaper
Performance Point Server 2007 – Datasheet

[7] “Co-engineering the Adaptive Complex Enterprise for Service Delivery – an Integrated Interdisciplinary Approach for Performance” – ACE Theory, Jay Ramannathan and Rajiv Ramnath, CETI

[8] “Enterprise Architecture As Strategy: Creating a Foundation for Business Execution” - Jeanne W. Ross, Peter Weill, and David C. Robertson – Copyright 2006
[Argyris78]: Argyris, C., and Schön, D. Organizational learning: A theory of action perspective, Addison Wesley, 1978.

[9] Kaplan, R. S., Norton, D. P. The Balanced Scorecard: Translating Strategy into Action, Harvard Business School Press, Boston Mass., (1996).

[10] “Business Intelligence for IT” - Bill Cason – Architecture Magazine Vol 3 Issue 1 - February 2007

[11] “IT for Analytical Competition” - Thomas H. Davenport and Jeanne G. Harris - Architecture Magazine Vol 3 Issue 1 - February 2007

[12] “Gartner on CMDB” – Gartner RAS Core Research Note G00137125, by Ron J. Colville, March 13, 2006

[13] “Information Technology Infrastructure Library “ITIL” Review” – Presentation by Michael W. Gourley and Alan J. Steible

[14] “System z CPU Capacity Planning Information for SNI and Boundary Function Workload” -

[15] “Capacity and Performance Management: Best Practices White Paper” - Document ID: 20769 – Cisco Systems – October 4, 2005

[16] “CMDB or Configuration Database: Know the Difference” – November 21, 2006 - Gartner, Inc.

[17] “Optimizing ITIL Best Practices with Mercury BTO” - Mercury Systems

[18] “The Role of the Capacity Management Database in the ITIL Capacity Management Process Extending the CMDB with Capacity Management Data” – BMC Software Whitepaper – April 27, 2006

[19] “The Value Proposition for Capacity Planning” - A White Paper Prepared by Enterprise Management Associates – 2003

[20] “How to Do Capacity Planning” – Team Quest - White Paper TQ-WP23 Rev. B – 2003

[21] “Capacity Planning DISCIPLINE FOR DATA CENTER DECISIONS” – Team Quest - TQ-EB01 Rev. A – 2004

[22] “White Paper: The Principles of Capacity Management” – Sarquol Limited – March 16, 2006

[23] “Capacity Management for ASPs” - Unisys Corporation: Jeroen Bom, Joe Helm, Kelly Millsaps, Hilda Willems Microsoft Corporation: Kathryn Rupchock, Kent Sarff – Microsoft Tech Net – Microsoft Enterprise Services White Paper – September 1, 2000

[24] “Service Management Functions - Capacity Management” – Microsoft Tech Net – November 21, 2006

[25] “Capacity Planning” - By Louis de Klerk (Inobits Consulting Pty., Ltd.) and Jason Bender (MSNBC) - Microsoft Enterprise Services White Paper E-Commerce Technical Readiness – Microsoft Tech Net - April 1, 2000

[26] “Using ITIL Best Practices to Create a Capacity Management Process” – By Chris Malloy – As presented at NaSTEC 20 October 20 – 21, Chicago, Ill. – Document dated October 1, 2003

[27] “The Missing Link; Capacity Management and Business Requirements” - By Jan Vromant August 16, 2004

[28] “Building Effective Storage Capacity Management in Support of Information Lifecycle Management – Guidelines for a practical, ITIL process-based approach – By Joel Wenger

[29] “IT Services Management Service Brief – Capacity Management” – By Rick Leopoldi – May 25, 2002 – Copyright RL Consulting

[30] “AN OVERVIEW OF TRR CAPACITY MANAGEMENT PROCESS” – By Kenneth W. Kolence Kolence Associates, 3591 Louis Road, Palo Alto, CA 94303

[31] “A Capacity Management Service for Resource Pools” – By Jerry Rolia, Ludmila Cherkosova, Martin Arlitt of Hewlett-Packard Laboratories, and Arthur Andrzejak of Zuse Institute Berlin (ZIB)

[32] “Computer Capacity Planning: Strategy and Methodologies” – By I. Lynne Carper, Susan Harvey, and James C. Wetherbe – from Database, Summer 1983

---

[i] Various approaches to managing IT Capacity exist, each with their pros and cons. However, one thing stands true, regardless of the approach: Good Capacity Management is extremely difficult to achieve. Exacerbating this problem is the current trend, for the organization to become more Service-Oriented in their business model. Ronald Schmelzer writes in his January 2007 ZapThink! Article, *Improving Business Functionality Through Model-based Service-Oriented Business Applications*, “[This realization] signifies a fundamental shift from today's rigid and inflexible systems towards loosely-coupled approaches to application development and deployment. This new shift requires a redefinition of the concept of business functionality in a way that marries continually changing capabilities with new requirements for usability and flexibility. The aim is a new capacity for achieving competitive advantage, lower cost, and increased business value” [1]. With this industry shift comes the realization that industries need a better grasp on the overall capacity capability of their IT systems to support these types of changing business functions. To facilitate this heightened awareness within the Enterprise, Capacity Management and Planning should be brought to the forefront of business planning and recognized by the entire Enterprise as being a critical component to their success.

[ii] To illustrate the principle area of concern from each viewpoint, imagine the following analogy as drawn from a manufacturing perspective of yesterday: Product Line -> Manufacturing Process -> Tool (machine and/or human).