




Optimize Resources and Services to Meet Business Goals



Abstract: The introduction of dynamic operating environments through the application of Utility computing technologies poses special problems in assuring optimal allocation of resources to assure the consistent reliable delivery of business services. IT services directly link to the achievement of business goals. The complexity of the relationships, interactions, and interdependencies between resources and operations required to deliver services requires automated management tools. Dynamically changing resource allocations place unique demands on planning and management in order to balance current demands with the achievement of business performance goals. This paper provides a discussion about the unique challenges of managing the allocation of resources in a dynamic changing operating environment. BEZProphet™ from BEZ Systems provides an example of a solution designed to address these challenges.

Introduction

The current frenzy of enterprise focus on Business Service Management (BSM), Service Level Agreements (SLA), and IT-business alignment is meant to ensure that IT efforts directly contribute to enterprise business success. Enterprise management now fully recognizes the increasing dependency of business operations on its IT infrastructure and the potential financial impact for good and bad. Far from being an inconsequential 'commodity' service, IT and technology represent the foundation for security, business continuity, and services that connect the enterprise with its customers, business partners, suppliers, and employees.

For IT, this dependency has forced a re-thinking of their responsibilities and tasks. Once IT focused its efforts on simply maintaining what were relatively 'isolated' pieces of infrastructure. Today, IT's responsibility extends far beyond infrastructure maintenance to include efforts to support and assure the availability and performance of the business services and processes that drive top-line revenue gains, increase customer satisfaction, and lower operational costs. It is through the delivery of these services and processes that IT connects directly to and influences most strongly the success of business operations.

The distributed nature of these business services and processes frames the requirements for flexible integration, dynamic applications, and constant access to corporate data. The combination of radically changing operational tasks with new responsibilities means that to succeed IT must initiate and undertake fundamental change in the definition, prioritization, and processes that make up its day-to-day tasks. Let's examine how this affects the operational environment.

The Operational Profile of the New Datacenter

The task of data accessibility, protection, and delivery has changed, taking on new importance in terms of efficient performance as application creation and implementation changes. For instance, the increasing flexibility of application deployments (driven by J2EE) and integration architectures (driven by SOA and Web Services) means that enterprise databases and data warehouses are accessed by an increasing number of discrete business services and must support an increasing number of complex, integrated business processes. These business processes support delivery of the fundamental services that directly affect business success. Clearly, the delivery of data directly impacts the ability of business critical applications and processes to function effectively. IT services enable the reliable functioning and delivery of the business services defined by these applications and processes. The emerging and escalating speed, reliability, accessibility, and security concerns mean that data delivery itself has

become, and requires management as, a critical IT service. The reliable delivery of such services connects IT to business success. Yet such improvements represent only part of the challenge facing enterprise IT. Scarcity and restrictions on staffing, budget, and skills pose cost control problems in a competitive environment increasingly operating under global pressures.

In today's globally competing environment, market share and revenue gains must balance with lowered operational costs for enterprises to reach profitability goals. IT today experiences increasing pressure to align its activities with business goals. Pressures to increase enterprise resource utilization lead directly to the interest in and piloting of the many forms of utility computing. The potential for increasing the utilization of existing enterprise computing power, permanently lowering future capital expenses on hardware, and increasing flexibility in resource allocation to hot business projects is a promise simply too great to ignore. But such computing has challenges and benefits.

Dynamic Management of Data

Specifically of concern for data managers are the utility computing options appearing in complementary infrastructure. Such options include new dynamic storage servers, data networking options, and database management systems with increased levels of parallelism available that permit coordinating across multiple data-stores. These innovations drive the emerging interest in dynamically allocating data serving resources. Such innovation is necessary to keep pace with the on-going advances in dynamic server and network resource allocation. The speed and escalating complexity of today's IT operating environment inevitably drives efforts to standardization and automation of IT services. Operations managers and administrators, already losing ground in their attempts to track relationship dependencies in relatively stable environments, find it beyond human capability to do so in a dynamically adapting environment.

The combination of utility computing flexibility and service management is ultimately driving the world of data management away from singular databases and towards a broader concept of data services. Ptak, Noel & Associates believes that this shift requires the integration of real-time capacity management and auditing capabilities into traditional data problem resolution and prevention processes.¹ Once again, maintaining consistency and coherency among these data resources drives support for automated functions.

¹ Read "Best Practices for Service Level Management in Utility Computing Environments" Ptak, Noel & Associates, for a more discussion on the impact of utility computing on IT management.

Problem Resolution and Prevention in the new Datacenter

The challenge is not simply one of real-time capacity management, which does not really represent anything new. Today, administrators routinely perform bottleneck analysis to determine whether the cause of a service problem arises from capacity constraints, configuration problems, software bugs, or software design issues. For example, upon finding a capacity constraint, administrators can easily start procedures that add service capacity resolving the problem. The management processes involved are familiar and far too limited to address the challenges of utility computing.

Utility computing changes more than just the speed and ease with which resources must adjust. Frequently overlooked is the complexity of the underlying analysis required to determine how to adjust all the involved resources correctly while proactively planning for the dynamically changing needs of the business service. At issue, is more than a traditional capacity planning focus on server workloads and available capacity. Figure 1 speaks of capacities while utility computing requires thinking and assigning resources based on the needs of and impact upon integrated services not just the individual resource.

The allocation of resources and assignment of workloads in today's complex, integrated operating environment requires a comprehensive understanding of interactions, relationships, and impacts to balance the benefits to one service against the negative impact on another. IT services must align with business goals. Therefore, resource allocations must consider the full resource set assessing immediate and future impact on the ability to achieve business objectives. Decisions must be proactive and based on the total results not simply reacting to current conditions and short-range concerns.

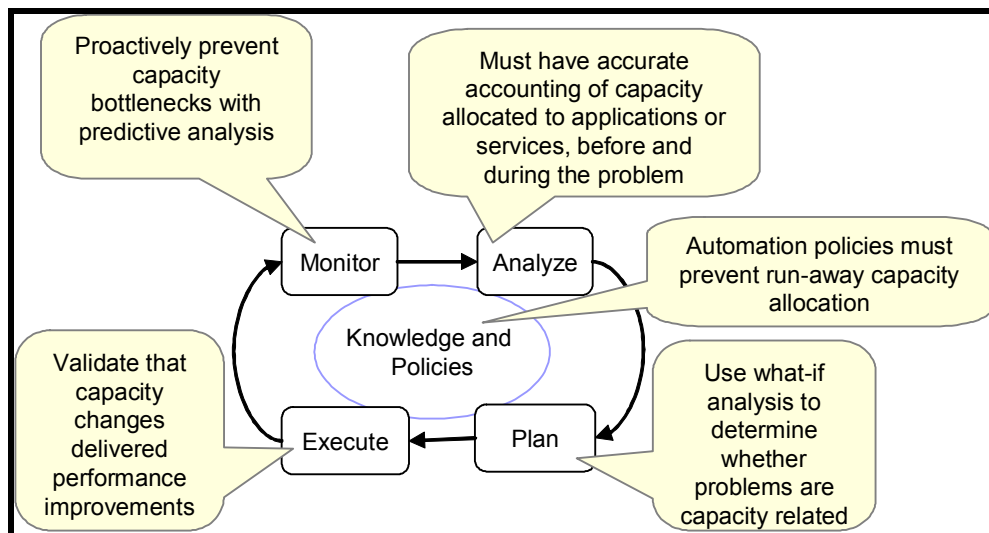


Figure 1: Data service management issues in a utility computing environment.

Ptak, Noel & Associates believes that the ability to easily allocate resources enables but does not guarantee more effective resource utilization. For example, too often enterprises have spent millions of dollars on additional storage capacity only to find they have attacked a symptom not a cause when the original data delivery problem reoccurs. Automated allocation policies can all too easily aggravate this scenario.

In the worse case scenario, a run-away capacity allocation system can potentially allocate all of the shared resources to a single data service disrupting other business services. However, the more likely scenario is that a single service will be consistently over-provisioned which is exactly the opposite effect enterprises desired when they implemented a utility computing system. Any experienced administrator knows that a good analysis tool will find the most cost effective way to resolve capacity related problems for a collection of resources. However, a more comprehensive solution design is necessary.

The problem requires a solution operating across the enterprise resource set looking beyond just what is taking place locally in, for example, a set of processors and their individual capacities. Workloads must be broken down into their line of business and application components. Such 'units of work' allow resource planners to understand the interplay and interactions among the data flows, workloads, and the available resource infrastructure. The result is not simply capacity planning but a solution able automatically to look at and help adjust resources that support data service delivery based on Line of Business. To monitor, measure, and predict needs and then to balance resource and workload allocations based on the business priorities that determine success.

The Challenge that Attends Easy Change

Utility computing dramatically simplifies resource allocation. The inevitable result is the introduction of more frequent change into the production datacenter environment. As any operations manager will attest, the frequent companions of change are unforeseen, and problematic consequences such as the disruption of another service.

Therefore, IT managers need an accurate map of which data and infrastructure assets relate to a particular service at a specific time before any serious problem resolution analysis can occur. After accurately identifying the service composition, it is necessary to arrange to monitor and measure service delivery. Monitoring and measuring service delivery provides a basis of agreement about performance. It sets a baseline for cooperation and measurement to report IT's contribution and support of business success. Looking at the issue from the perspective of a data service, the next step is for data collection that allows the analyst to understand:

- How the workloads are flowing through the data service infrastructure
- How the current workload patterns compare with baseline workloads
- How the response times are trending and where that time is being spent
- How will current performance change if various capacity variables are altered

Only when administrators have an accurate picture of their total service behavior and the ability to perform forward looking, comprehensive what-if capacity analysis can they truly resolve the underlying problem. New tools are necessary to simplify this proactive performance management analysis and extend it to the more complex utility computing deployments.

Finally, while problem resolution is an important part of delivering data services, the old adage of “an ounce of prevention” still holds true. Proactively predicting when capacity constraints will become an issue for future workloads can eliminate many availability and performance issues. Because utility computing makes data infrastructure capacity more flexible, the data service infrastructure will be changing more frequently. The increased number of configuration changes means that data managers will need to forecast future performance requirements more frequently to achieve uninterrupted service. The goal is to profile how capacity requirements change over time without requiring significantly more effort on the part of the data manager. Let's see what exists to address these challenges.

BEZ tackles the issues

BEZ understands that IT operations must align with business goals and that IT services form the basis of this connection. To achieve that end, BEZ provides automated analytical tools that simplify data service problem resolution and prevention in utility computing environments. With the introduction of BEZProphet™, BEZ provides what it calls Proactive Performance Management – the ability to look at data service delivery from the perspective of a line of business. BEZProphet™ measures the rate of delivery, then using acquired knowledge of performance, dependencies, and interaction - it predicts future service levels and potential service faults. Because of its insight into workloads and their performance footprints, BEZ can also perform and implement, if necessary, ‘capacity planning’ to avoid potential problems. BEZProphet™ can then monitor its effectiveness by measuring and comparing service delivery performance after a changing – comparing real to predicted performance.

BEZProphet™ achieves all of this by:

- Automatically discovering the data service infrastructure, including mapping the physical storage system's logical volumes to their database files, thus providing an accurate accounting of capacity assets.

- Creating a suggested list of critical availability and performance parameters for monitoring.
- Collecting the metrics through an agentless architecture, which simplifies deployment.
- Automatically creating a profile of the data service environment as well as establishing baselines for performance and capacity thresholds.
- “As is” analysis that forecasts capacity constraints based on current workload growth patterns.
- “What if” analysis allowing administrators to test potential infrastructure changes for troubleshooting a current problem or planning upgrades of hardware and software.

BEZ does not limit its solution to leveraging the newest utility computing technologies. Indeed, many of their existing clients have adopted a proactive performance management approach built on existing traditional systems hardware. For example, BEZ customers, such as JCPenney and Lowe’s, focus on identifying workload and capacity requirements trends on a daily basis for their data warehousing systems. BEZPlus’ Statistical Process Control (SPC) capability helps them to distinguish between a one-time performance exception and an emerging trend that will affect allocation needs on a more permanent and long-term basis.

Conclusion

The combination of utility computing flexibility and dynamic service management is ultimately driving the world of data management away from singular, distributed databases towards a broader concept of data services. Data services are defined, determined, and managed with flexibility and dynamism comparable to server, network, and storage infrastructure. These data services exist as part of the larger set of IT services that connect IT to the business and allow IT to align its operations with business goals.

For too long IT operated outside of and independent of business metrics applied to traditional business functions such as accounting. Because of this, IT’s unique and vital contribution to business success went unnoticed. In today’s globally competitive, dynamic environment, enterprises recognize the contribution of IT. The result is the increasing interest in ways to measure IT services and their alignment with business goals. IT operations managers and enterprise Line of Business managers justify acquiring solutions that measure and report on IT services in this way. BEZ’s BEZProphet™ provides such Business Intelligence.

The availability and performance of these data services underpin the availability and performance of business services and processes. Most enterprises already have problem resolution and prevention processes operating in place for their datacenters. As these enterprises move to incorporate service management and utility computing into their datacenter environments, the weaknesses of existing capacity analysis capabilities will become apparent. Enterprises will demand automated solutions that are able to proactively manage service delivery performance. Equally important, will be the ability to manage resources with a view to consistently meeting future, as well as current, operational metrics. This means solutions that operate according to policies that assure consistent, uninterrupted services. BEZ is delivering in BEZProphet™ just such a solution.



This paper was sponsored by:

BEZ Systems
345A Summer Street
Boston, MA 02210
Phone: 617.532.8800
www.bez.com

No part of this document can be used in any medium without the written permission of Ptak, Noel & Associates."

All trademarks are the property of their respective owners.

While every care has been taken during the preparation of this document to ensure accurate information, the publishers cannot accept responsibility for any errors or omissions.

About Ptak, Noel & Associates

With a belief that business success and IT success are inseparable, Ptak, Noel & Associates works with clients to identify, understand and respond to the implications of today's trends and innovations on the future of IT Operations.

www.ptaknoelassociates.com

About the Authors

Richard Ptak has over 30 years experience in systems product management working closely with Fortune 50 companies in developing product direction and strategies at a global level. Previously Ptak held positions as senior vice president at Hurwitz Group and D.H. Brown Associates. Earlier in his career he held engineering and marketing management positions with Western Electric's Electronic Switch Manufacturing Division and Digital Equipment Corporation. He is author of "Manager's Guide to Distributed Environments," (John Wiley & Sons, 1998). In addition, Ptak was technical editor of "Cisco Internet Architecture Essentials Study Guide: Cisco Internet Solutions Specialist" by Mathew Recore, Jeremy Laurenson, and Scott Herrmann (Cisco Press, 2002). Ptak holds a master's in business administration from the University of Chicago and a master of science in engineering from Kansas State University.

A recognized expert in infrastructure management, **Jasmine Noel** has 10 years experience as both analyst and researcher. Noel served previously as director of systems and applications management at Hurwitz Group, where she formulated and managed the company's research agenda. She was also a senior analyst at D.H. Brown Associates. Noel is regularly quoted in publications such as CIO Magazine, eWeek, InformationWeek, InfoWorld, and NetworkWorld. She also has contributed articles to several leading publications on various IT management topics. Noel holds a bachelor of science from the Massachusetts Institute of Technology and a master of science from the University of Southern California.

October 2005