

WHITE PAPER

**Optimizing Datacenter and Cloud Performance:
The TeamQuest Approach**

Sponsored by: TeamQuest Corporation

Tim Grieser
November 2011

INTRODUCTION

As the worldwide economy continues to bring many business challenges, IT organizations remain under heavy pressure to contain costs while providing competitive advantages by delivering effective business processes and applications to customers and internal users alike. At the same time, IT is challenged to manage ever more complex environments and infrastructures, including virtualization, Web, and cloud, and to satisfy service guarantees to end users across a wide variety of desktops, laptops, and smart handheld devices.

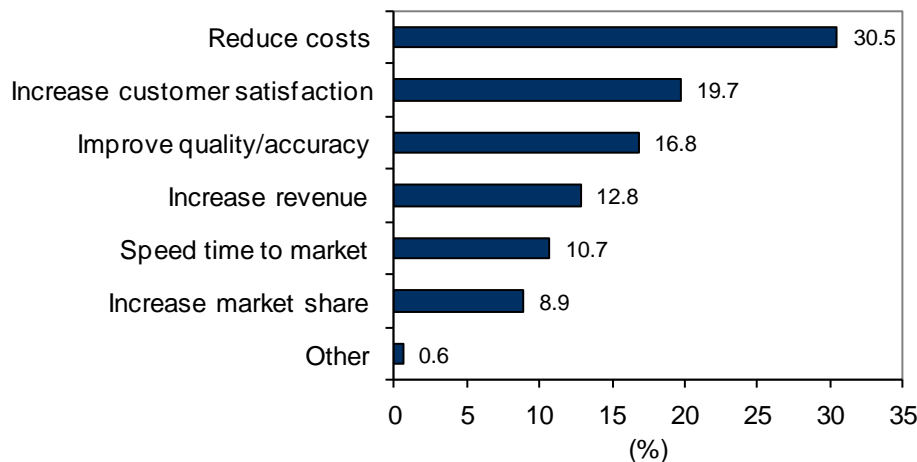
Key IT Business Goals

As business applications and processes become increasingly Web enabled and are directly accessed by end users, IT integration with business goals and priorities becomes ever more important. According to an IDC survey conducted to understand the relative importance of strategic business goals to IT organizations, cost containment continues to be of the highest importance even as IT focuses on other priorities important to business growth (see Figure 1).

FIGURE 1

Importance of Strategic Business Goals

Q. *Prioritize the following business goals as they relate to your organization by allocating 100 points to them. The more points you allocate, the more important the business goal.*



Source: IDC, 2010

According to the survey, other important priorities include increasing customer satisfaction, improving service quality, increasing revenue, increasing agility by speeding time to market, and increasing market share. Collectively, the key goals amount to improving IT efficiency and increasing IT integration with the business. These business goals are major drivers for IT organizations today.

PERFORMANCE AND CAPACITY MANAGEMENT OBJECTIVES

Performance management and capacity planning functions are needed to ensure that service objectives for key business applications and workloads are achieved in terms of performance, availability, growth requirements, and costs. Key functions include the following:

- ☒ Support key business processes by monitoring and managing the performance of critical IT services and applications so that required service levels — response times and workload volumes — can be met for service owners and end users
- ☒ Optimize costs; identify, size, and configure server hardware resources for provisioning and deploying applications and workloads; evaluate cost trade-offs
- ☒ Support IT business decisions needed for major infrastructure changes such as datacenter consolidations, moving to virtualized and dynamic environments or adopting cloud strategies
- ☒ Proactively prevent slowdowns or outages for business services and IT services, especially during peak periods; anticipate peak loads and plan for workload growth; prioritize workload requirements in terms of business importance
- ☒ Troubleshoot and resolve performance problems, diagnose and eliminate bottlenecks, meet service-level agreement (SLA) requirements
- ☒ Streamline and simplify IT processes for performance monitoring and capacity management to contain costs and reduce operational complexity
- ☒ Increase IT efficiency and make IT more productive, process more transactions with existing resources, reduce IT staff time needed for firefighting performance problems

The functions needed for performance management and capacity planning are typically enabled by a combination of IT staff expertise, software tools, standardized processes such as ITIL, and best practices.

The Business Role of Capacity Management

Capacity management is an ongoing process to ensure that IT hardware and software infrastructure components are correctly sized and priced as well as efficiently utilized to deliver IT-based services at required service levels. Capacity management consists of capacity planning and performance management functions. Traditionally, these functions have been conducted by senior IT professionals who

understand performance analysis principles and best practices — and who can apply their knowledge to key technical tasks such as performance measurement, workload characterization, workload trending and forecasting, bottleneck analysis, and "what if" scenarios for changes to hardware, software, and workloads.

Capacity management should always be conducted in a manner that supports business objectives and provides business benefits. In practice, capacity planners perform important business functions by applying their expertise to understanding and optimizing the trade-offs between delivering required performance to meet business service requirements and minimizing computer hardware capacity and usage costs.

Capacity planners should strive to closely integrate their analysis and recommendations with the business functions and workloads being supported by IT, not just in terms of hardware requirements. This means being able to analyze and report IT-based services in business terms such as transaction volumes and response times for key applications and workloads. This perspective allows performance and capacity decisions to be made on a business-centric basis and helps to better position and integrate capacity management experts with the owners and users of business services and production applications.

NEW ENVIRONMENTS BRING NEW CHALLENGES

Performance and capacity management processes were originally developed to optimize the cost, performance, and utilization of large servers, especially mainframes. Today, while large servers are still important, the requirements for performance and capacity management have greatly expanded and increased in complexity with the widespread adoption of complex multitier applications running in dynamic, virtualized and Web-based environments. Typically, these environments are based on x86 hardware architectures, often spanning large numbers of servers. Such environments bring a multitude of challenges for performance and capacity management:

- ☒ Provide more real-time insights into workloads, performance, and resource consumption in large scale-out environments spanning hundreds or even thousands of servers
- ☒ Monitor performance and availability metrics for key workloads on both physical and virtual servers, with the ability to alert if performance thresholds such as server utilization are exceeded
- ☒ Plan and anticipate the effects of physical to virtual server migrations and consolidations and analyze the utilization and performance impacts of deploying multiple virtual images onto a single host server
- ☒ Analyze both static and dynamic virtual environments and plan and analyze the effects of virtual image motion, virtual clustering, and automated workload migration

Limitations of Traditional Capacity Management

IT organizations using best practices for capacity management employ capacity planning to ensure that the IT infrastructure is optimized to meet business needs. This function is as important today as it has been in the past.

Traditionally, capacity planning required well-trained and seasoned experts with skills such as the ability to use complex mathematical models for decision making or the ability to gauge future capacity requirements based on years of experience and rules of thumb. Because of the time and expertise required for capacity planning, it was typically employed only for the largest (often the most expensive) and most critical servers.

Today, most IT organizations do not have such seasoned capacity planners available. Even with such experts, today's environments typically consist of large numbers of inexpensive servers that house potentially critical applications. In these cases, there is no time for exhaustive expert analysis of individual servers before needed capacity changes are made.

Capacity Management for Large Scale-Out Virtualized Environments: What's Needed?

Performance analysts and capacity planners need comprehensive and sophisticated software tools to deal with the scope and complexity of managing today's large scale-out virtualized environments. Performance and capacity management tools to support these environments should provide a number of capabilities, such as the following:

- ☒ Tools to simplify capacity management, reducing the amount of knowledge and expertise required
- ☒ Tools that streamline capacity management so that problems are resolved more quickly and decisions can be made more rapidly
- ☒ Tools that can scale to optimize environments with large numbers of servers
- ☒ Tools that support virtualized infrastructures at both the physical server level and the virtual server level across heterogeneous hypervisor environments
- ☒ The ability to understand how workloads will perform when coexisting with other workloads in virtualized environments
- ☒ The ability to support dynamic environments — including virtual machine motion — as workloads and virtual machines change and move across physical servers
- ☒ Providing a macro view of datacenter capacity and performance, with the ability to look inside virtual machines when necessary to determine the root cause of performance problems

- ☒ The ability to determine when virtual and physical servers will run out of capacity and how best to configure systems to avoid capacity problems
- ☒ Support for capacity management of diverse, multivendor environments with both physical and virtual infrastructures, across technology silos, from a "single pane of glass"

Overall, tools for managing complex scale-out virtualized environments should present a high-level graphical view from which performance analysts and capacity planners can drill down into successively more detailed views of infrastructure elements, performance metrics, and alerts and exceptional conditions.

THE CHALLENGES OF CLOUD COMPUTING

Cloud computing represents a new set of challenges for performance and capacity management professionals. Cloud infrastructures are typically based on shared, pooled, highly virtualized hardware and operating environments. Clouds require extensive management software for enabling such functions as resource allocation, self-service catalogs, automated provisioning, service-level management, usage-based metering and billing, and support for dynamic expansion and contraction of resources or "elasticity." As such, clouds represent technology evolution for the delivery of business services and workloads, but cloud infrastructures still require the same performance and capacity management functions as more conventional infrastructures.

Clouds Need Capacity Management

When clouds first evolved, there was a perception that capacity planning was not needed due to the highly pooled environments and support for elasticity. However, as cloud usage has grown, this myth has been dispelled, particularly as service providers must satisfy SLAs for functions delivered as cloud services and as cloud-based service slowdowns and outages have received public attention. Indeed, managing performance and capacity for cloud environments must deal with increased complexity and needs to address requirements from two perspectives: cloud service providers and cloud service consumers.

Cloud service providers must provide operational environments that meet performance and availability SLA requirements for consumers of their services. They must execute performance and capacity management functions for the cloud infrastructure, ensuring that sufficient hardware resources are available to meet SLAs under variable loading conditions and in environments supporting multitenant consumers. If elasticity is supported, the overall capacity of the primary cloud infrastructure plus the additional infrastructure available for elastic expansion must be sized and managed to meet SLA requirements. Elasticity works only if sufficient resources are available through elastic expansion to meet expanded service requirements.

Cloud service consumers need to understand the service levels they are actually receiving from cloud service providers. Consumers also need to know whether there is sufficient service capacity to support growth or peaks in their workloads (as well as contention from other workloads in multitenant situations) and still meet required

SLAs. Insufficient service provider capacity can lead to performance degradations, slowdowns, or even outages as resources become saturated and bottlenecks grow.

Capacity Managers and the Cloud

The increasing adoption of cloud infrastructures delivered to consumers in such forms as SaaS, PaaS, or IaaS presents heightened complexity and challenges for performance and capacity management professionals. Cloud infrastructures can be deployed as public clouds, private clouds, or hybrid clouds, further expanding the choices and decisions that must be made regarding where best to host applications and business services to meet service-level requirements and optimize costs.

Responsibility for capacity planning has traditionally been part of enterprise IT, with the in-house datacenter being the location for IT infrastructure and also being the subject of performance management and capacity planning functions. Now, with increasing business integration and the emergence of cloud alternatives, capacity managers need to take on an expanded role by asking the right questions and providing expertise and advice to a variety of constituents. Specifically, they must:

- ☒ Help corporate IT managers evaluate cost/performance trade-offs for alternative infrastructure deployment strategies — use conventional infrastructure versus cloud infrastructure
- ☒ Help business service and application owners evaluate which services should be hosted on cloud infrastructure versus conventional infrastructure (What mix of external, internal, or cloud will best meet business service needs?)
- ☒ Provide support for monitoring service levels on cloud deployments for providers and consumers to see if objectives are being met
- ☒ Help size and optimize cloud infrastructures to meet SLAs (Can peak loads be supported? Is capacity available for elasticity?)

In addition to these functions, capacity planners should be able to make detailed cost analyses of the various infrastructure alternatives, including the ability to see how costs grow with increased volumes and workload growth. For example, while consuming SaaS-based services may be initially "cheap" due to lack of capex investment by the consumer, SaaS usage costs may eventually grow and skyrocket past the cost of in-house deployment alternatives as volumes increase.

TEAMQUEST

TeamQuest is a 20-year-old "best of breed" independent software vendor providing software and professional services for IT service optimization (ITSO). TeamQuest supports performance and availability management, capacity planning, and service-level management primarily for large, complex datacenters running mission-critical applications. TeamQuest began life as a division of Unisys, before becoming an independent company. Today, TeamQuest serves a large number of enterprise-class customers internationally, with concentration in key vertical segments including finance, insurance, telecom, and service providers. TeamQuest provides a wide

range of support for IT service optimization, helping IT organizations consistently meet service levels while minimizing infrastructure costs and mitigating service delivery risks.

The TeamQuest Approach

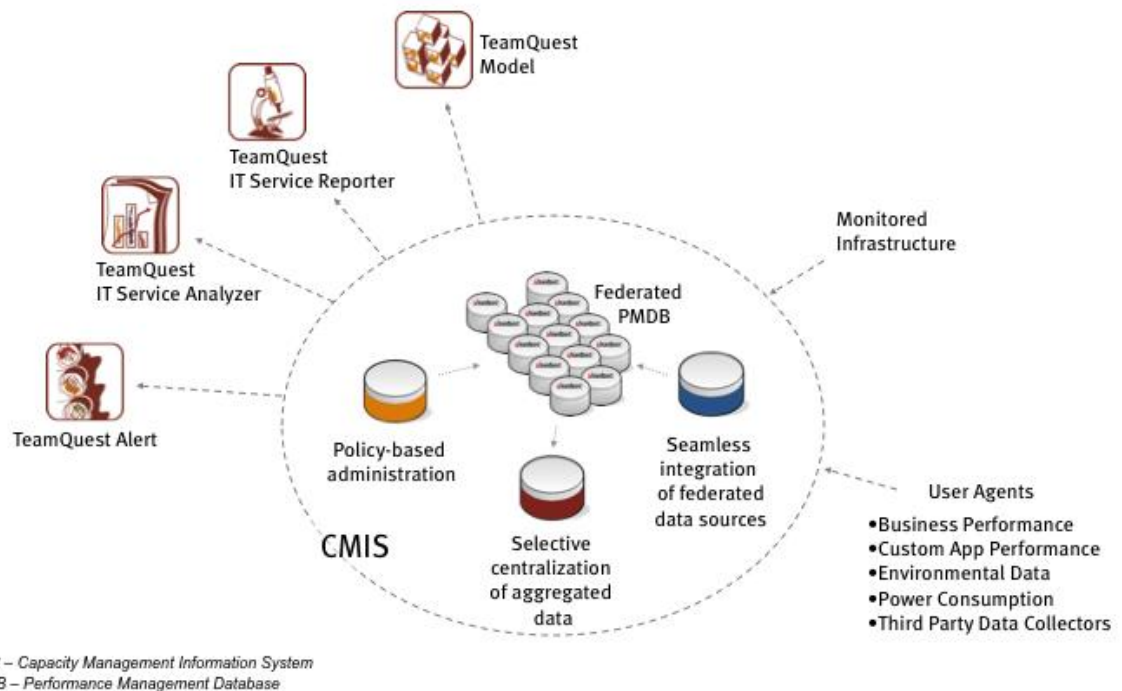
TeamQuest develops and markets software that supports IT service optimization by addressing a wide range of performance management and capacity planning functions for a variety of hardware, operating system, and application platforms. TeamQuest products address both physical and virtual environments, as well as cloud infrastructures. The TeamQuest approach is based on quantitative performance measurement and analysis that includes agent-based monitoring, performance reporting, statistical analysis, trending, and modeling for "what if" performance analysis and capacity planning.

TeamQuest Performance Software

TeamQuest supports performance management and capacity planning functions with product suites organized into four major groups, as illustrated in Figure 2.

FIGURE 2

TeamQuest Performance Software



Source: TeamQuest, 2011

The key TeamQuest product groups are as follows:

- ☒ **TeamQuest Alert** for real-time performance and status monitoring. TeamQuest Alert is a system performance console that provides performance monitoring and alarm event management together with "health status" displays for the managed servers.
- ☒ **TeamQuest IT Service Analyzer** for problem analysis. TeamQuest IT Service Advisor supports a Web interface to enable users to drill down and help diagnose service problems and perform root cause analysis — both for real-time events and for historical data.
- ☒ **TeamQuest IT Service Reporter** for customized reporting. A Web-based interface is used to create and customize reports for distribution to IT and business managers. Report dashboards can be used to communicate performance and availability information to service providers and service consumers.
- ☒ **TeamQuest Model** for "what if" capacity planning analysis. TeamQuest Model can be used to evaluate performance impacts for a wide variety of hardware capacity scenarios such as physical to virtual migration, workload and server consolidation (both physical and virtual), the impact of workload growth and peaks, server memory and configuration changes, and evaluating new hardware alternatives. TeamQuest Model incorporates a built-in analytic queuing network solver that provides powerful analysis capabilities that enable capacity planners to make more accurate infrastructure optimization decisions without requiring a Ph.D. or advanced math skills.

TeamQuest's performance software products access data from TeamQuest's Capacity Management Information System (CMIS), which collects, stores, manages, and administers performance data. The CMIS provides a scalable way to efficiently store detailed performance data. Each of the products retrieves data from the TeamQuest CMIS for performance reporting and analysis. The CMIS supports a federated performance management database (PMDB) to store performance information generated by monitoring agents and to provide the historical data for comprehensive reporting and analysis capabilities. Fine-grained detailed data can be kept close to the data source for the short term. Only summarized, aggregated data need be stored centrally for enterprisewide analysis. This approach provides selective centralization while preserving the details needed for solving real-time performance problems.

Key Performance Software Benefits

TeamQuest operates as a "best of breed" vendor with a sharp focus on performance management and capacity planning software for IT service optimization. As such, TeamQuest offers a range of benefits available to customers who implement its performance and capacity management processes and software. These benefits can be summarized as follows:

☒ **Performance management.**

- ☐ Ensure that key applications and business workloads achieve service objectives for throughput, transaction volumes, availability, and response times for end users.
- ☐ Monitor service performance compared with agreed-upon objectives for IT infrastructure and business processes.
- ☐ Report on infrastructure health and status.
- ☐ Generate alerts for pending slowdowns, bottlenecks, or outages.
- ☐ Provide drilldown diagnostics for root cause analysis.
- ☐ Avoid performance crises and eliminate the turmoil and overload on IT staff resulting from serious performance problems.
- ☐ Communicate service status over the Web.

☒ **Capacity planning.**

- ☐ Optimize hardware costs while achieving performance objectives.
- ☐ Evaluate hardware alternatives needed to support current and future workloads and workload growth.
- ☐ Reduce hardware costs through timely acquisition practices, based on ongoing performance management and capacity planning actions.
- ☐ Avoid overbuying new hardware.
- ☐ Analyze the cost versus the performance of major infrastructure changes such as workload and server consolidation, virtualization, dynamic movement of workloads, and adoption of public or private cloud infrastructures.

☒ **IT service optimization.**

- ☐ Optimize IT costs and achieve higher return on investment (ROI) for IT investments.
- ☐ Increase server utilization rates.
- ☐ Reduce the complexity of IT operations by providing a consistent methodology and supporting software tools for performance management and capacity planning.
- ☐ Manage proactively by detecting trends that can lead to performance problems and making adjustments before problems can grow to impact end users.
- ☐ Automate the detection and alerting of problem conditions.

- ☒ **Agility.** TeamQuest's tightly focused product set can result in shorter implementation times than those provided by more comprehensive products from enterprise system management software vendors. TeamQuest can react quickly to support changes in hardware offerings and operational environments.
- ☒ **Enterprise-class heterogeneous coverage and scalability.** TeamQuest's products cover a wide range of operating systems, including AIX, HP-UX, Solaris, Linux, and Windows, as well as virtual environments, databases, and Web servers.
- ☒ **Integration with other vendors.** TeamQuest integrates with key products from other management software vendors such as management consoles, frameworks, monitoring agents, and performance data repositories.
- ☒ **Independence.** Capacity planning decisions often require making trade-offs between offerings from competing hardware vendors. Customers can benefit from TeamQuest's position as an independent software vendor to help make vendor-neutral decisions regarding hardware alternatives.
- ☒ **Best-of-breed expertise.** As a 20-year-old best-of-breed vendor, TeamQuest provides deep knowledge and expertise around performance management and capacity planning processes and best practices.

CHALLENGES/OPPORTUNITIES

Perhaps the most important challenge facing TeamQuest is to gain greater identity and recognition in the IT management software space. This market is rapidly consolidating and is dominated by large enterprise vendors that sell a broad range of system and application management software. As noted in this analysis, TeamQuest competes as a dedicated "best of breed" software vendor with benefits to customers coming from a tight focus on the market requirements and product technology needed to support performance and capacity management. At the same time, TeamQuest must continue to leverage its ability to connect to the major management software frameworks in a way that effectively positions the company as the vendor of choice for performance management, regardless of the management platform that may be established at a customer site.

Another challenge, which is not unique to TeamQuest, is to further simplify the overall performance management activity through increasing automation and integration in the software. Application deployments are becoming more and more complex with multiple tiers that may include Web servers, application servers, and databases, often with multiple instances. Operating environments continue to expand and include virtualization, automation, dynamic motion, and cloud infrastructures. The challenge is to continually simplify performance management in environments with large numbers of managed systems, heterogeneous physical and virtual operating environments, and growing adoption of cloud. TeamQuest has made significant advances in product simplification through the use of Web-based interfaces and support for a comprehensive CMIS/PMDB product structure.

SUMMARY AND CONCLUSION

The benefits of performance management and capacity planning are well established and include improved service levels, better performance, and cost savings stemming from improved operational efficiencies, more productive business processes, and savings in capital expenditures for hardware. The ability of TeamQuest performance software suites to perform the associated management tasks based on a "best of breed" product set and the deep staff expertise developed during 20 years of company operations constitute core strengths and provide the basis for the company's competitive differentiation.

CASE STUDY: COOP GROUP

Coop is a large retail enterprise group headquartered in Basel, Switzerland. The company is organized as a nationwide cooperative society. According to company information, Coop operates over 1,900 stores, including supermarkets and megastores, that provide a wide variety of food, nonfood, and services products. Coop offers a mix of branded items, including its own brands, and flagship labels. It also operates specialist retailers and manufacturing companies. According to public financial reports, Coop employed approximately 54,000 individuals and had annual sales of over US\$22 billion in 2010.

Coop IT operations are housed in three datacenters situated in two locations in Switzerland. The IT infrastructure includes over 1,000 servers running a variety of operating environments, including Solaris, Linux, Windows, VMware, and IBM AIX. Key applications are based on SAP, including business intelligence and business warehousing, with data housed on Oracle and SQL databases. Another key application is WAMAS for warehouse operations and management. According to company officials, Coop is one of the largest SAP customers in Switzerland.

Coop needed a comprehensive software solution for performance and capacity management of its large, diverse IT infrastructure. The company needed to move to a more proactive process that focused on actions that deliver required service levels and prevent performance problems — not just react to problems as they occur. When evaluating potential software products, Coop had a number of key requirements for the software:

- Provide key functions, including monitoring, reporting, trending, and modeling
- Support all the operating environments and platforms in use
- Be able to collect data from a variety of sources, including any database
- Support both short-term and long-term performance management
- Have common Web-based access

TeamQuest was selected based on its ability to meet these requirements.

In practice, Coop uses the TeamQuest software to support a number of activities. Performance is tracked on an ongoing basis with daily, weekly, and monthly reporting across many different systems and applications. Service levels are tracked for individual systems — many with specialized performance requirements. For example, in terms of short-term performance, one key SAP process is run every 30 minutes and is monitored to see if the execution time stays within the required service objective. If exceptions or performance problems occur, Coop IT engineers can use monitor data and TeamQuest analysis to troubleshoot the problem. Another example is the ability to monitor a particular server and understand that it is very lightly utilized during the day but can run at high utilization at night because it is processing batches of transactions entered during the day.

TeamQuest software is used to manage performance for 200 Solaris systems and is successful in managing AIX LPARs. In terms of capacity sizing, TeamQuest Model was recently used to help size the acquisition of two large IBM Power systems.

According to Coop's manager of IT operations, the company is happy with the products and support provided by TeamQuest. Coop uses TeamQuest onsite consulting periodically for product customization and to assist with new releases and training for new features. He likes the ability to use scripting to help customize the TeamQuest performance management functions. Recent extensions to TeamQuest products are in line with Coop's ongoing needs.

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2011 IDC. Reproduction without written permission is completely forbidden.