

WHITE PAPER | OCTOBER 2014

Evolved Capacity Management

Why it's so critical and the six key steps to success.

Infrastructure and performance management.



Table of Contents

Executive Summary	3
Section 1: The Challenge: Tactical, Siloed Approaches	4
Section 2: The Opportunity: Establish the Six Critical Steps	5
Section 3: The Benefits of Business Centric Capacity Management	9
Section 4: Conclusions	10
Section 5: The CA Technologies Advantage	10

Executive Summary

Effective capacity management has become a critical differentiator for IT organizations. Those that can't effectively evolve their capacity management practice will continue to struggle with complexity and negligible insights into capacity sizing and the impact of changing demand and resulting service/application performance. However, those that gain advanced capacity management capabilities will be able to more effectively right-size investments, support key IT projects, and align resources with business objectives. This paper offers a practical look at capacity management, outlining the six key steps IT organizations can take to realize capacity management that delivers maximum value.

Introduction: Capacity Management, a Strategic and Operational Imperative

For most IT executives, there is a fundamental, common reality: The demands are high while budget and manpower are flat or declining. Within this context, capacity management—the process of aligning IT resources with current and emerging demands—is increasingly growing in importance. Capacity management is a key means with which IT teams can address their core operational objectives:

- Maximizing resource utilization in order to reduce investments and costs.
- Addressing the increasingly pressing demand to support business agility.
- Performing strategic IT infrastructure planning that guarantees capacity will be available when needed while providing acceptable service/application performance levels.
- Enhancing budgeting accuracy and intelligence in order to more effectively manage expenses and new investments.
- Strengthening service level agreement (SLA) definition and compliance to meet or exceed availability, performance and response times requirements.
- Improving visibility and transparency in terms of how business users are consuming IT resources.

Further, the extent to which an IT organization can support key strategic initiatives is increasingly contingent upon effective capacity management. Today, capacity management is essential in supporting these initiatives:

- **Cost reductions.** Effective capacity management enables executives to identify redundant or under-utilized systems, providing insights for eliminating hardware—and so realizing all the associated reductions in software license expenses, administration costs, power and cooling requirements, and more.
- **Successful application deployments and service delivery.** To effectively support the demands of new services and applications, organizations need to effectively build and track the capacity required.

- **Migrations.** Whether an organization is migrating to the cloud or a new infrastructure platform, capacity management is vital in preserving adequate service levels throughout the transition and beyond.
- **Data center and IT consolidation.** Sound capacity management is essential in identifying areas for consolidation, and for allocating existing workloads successfully.
- **Virtualization or hypervisor rationalization.** As organizations grow increasingly reliant on virtualization, they are prone to battle virtual machine (VM) sprawl and increasing license costs. Capacity management can help organizations rationalize the use of expensive hypervisor platforms and combat unnecessary VM proliferation.
- **Supporting business growth.** Whether an organization is expanding due to mergers or acquisitions, or an expanding customer base, capacity management is vital in enabling intelligent scaling of the resources required to support this growth.

In short, where systems monitoring may represent the senses of the IT organization, the prescriptive analytics derived from capacity management is vital to achieve a heightened insight for the department, a fundamental facility that can either strengthen or weaken virtually any endeavor the organization is tasked with.

Section 1

The Challenge: Tactical, Siloed Approaches

While capacity management is a critical endeavor, many organizations are being poorly served by their current capabilities. In large part, that's because historically many organizations employed capacity management in a more tactical fashion, with a specific team looking to address a specific set of infrastructure elements, rather than taking a strategic, enterprise-wide perspective.

Today, many IT groups are manually collecting performance data from siloed tools and systems. A virtualization administrator will gather metrics from their virtualization vendor's operations platform, a hardware administrator will leverage a hardware-monitoring platform, and so on. In addition, in these environments, capacity management efforts tend to be done on an ad hoc, inconsistent basis. As a result, if a CIO is looking to make plans or report on capacity, he or she will have multiple teams delivering many different reports. This reality poses several critical shortcomings:

- **It's inefficient.** Staff time is focused on labor-intensive, time-consuming reporting efforts, and analyzing data is cumbersome.
- **It's ineffective.** The siloed views available through these approaches provide very limited insights. These tactics don't enable managers to track or manage service performance, or gauge the capacity demands of business services or the data center as a whole. Further, IT can't efficiently map services and service levels to all underlying components.
- **It undermines IT value and reputation.** These tactical approaches significantly limit the ability of IT to support business objectives, and they don't instill confidence in business executives.

Ultimately, without comprehensive, effective capacity management, IT organizations are flying blind, which means IT teams have to resort to being reactive rather than proactive. Not only does this make it difficult to manage current infrastructure and capacity demands, but also it significantly hinders the organization's ability to support emerging requirements and initiatives.

Section 2

The Opportunity: Establish the Six Critical Steps

To deliver against their current charters, and better support evolving technologies and business demands, IT organizations need to adopt more sophisticated and comprehensive capacity management approaches. To achieve operational and strategic objectives, it's vital that these organizations move past tactical, ad hoc approaches and employ unified, prescriptive and cohesive capacity management capabilities across the enterprise.

As they embark on this journey toward capacity management sophistication, most organizations will be best served by taking a phased approach. Value can be measured by the first project alone without having to put the entire data center infrastructure under management. Starting off with understanding the utilization and capacity levels of your ecommerce application or internal mail system can provide insight into approaching bottlenecks or cost reduction opportunities never discovered before.

By adopting incremental steps, organizations will position themselves for both near-term benefits and long-term success. The following sections offer an overview of the key steps organizations should look to take in advancing their capacity management objectives.

Step 1: Establish a unified view of component capacity management data

To move past the siloed approaches of the past, it's vital that IT organizations establish a central archive and management interface for all components and elements across the enterprise. Gathering complete data sets is critical. If gaps exist, they can serve to mask great spikes in demand, and therefore lead to greater risk in your decision-making processes.

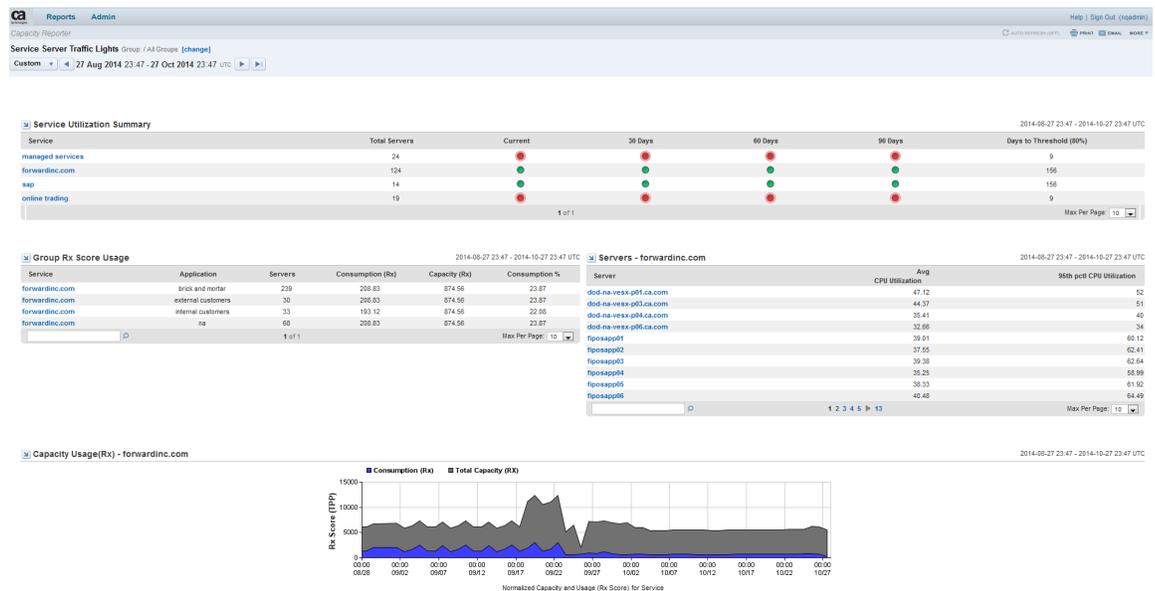
This archive should provide an aggregation and normalization of data from all components and resources, including hardware and software platforms, physical and virtual resources and on-premises and cloud-based systems. In addition, this data should be aggregated across distributed regions and resources, and provide detailed metrics on specific components, such as CPU, memory, storage and network utilization.

Once organizations begin to leverage a centralized archive, they can start to get insights, not just on the status of a specific component, but of the entire data center. This central archive is a critical building block, one that enables more efficient and effective management, planning and reporting on capacity across the enterprise. With this archive, it becomes much easier for IT organizations to establish baselines and identify exceptions, as well as risk and cost savings with the insight into over- and under-utilized resources.

Figure A.

Enterprise view
of capacity

An aggregate view
of data center
capacity across
physical and virtual
environments as
well as on-premises
and cloud-based
systems.



Step 2: Establish application/service capacity management capabilities

Once organizations have established a central, enterprise-wide archive of component metrics, the next step is to tap into configuration management databases (CMDBs) or equivalent systems and combine configuration, dependency and relationship data for workloads, applications and complete business services. This valuable information can then be used to map the relationship of business services to their associated infrastructure components. By understanding business service workloads and how they correlate to the usage of specific resources, IT executives can move from component capacity management to understanding and evaluating capacity at a enterprise-wide business service level. As a result, executives can understand, for example, how many more user transactions can occur before breaching capacity thresholds. They can also identify which dependent server component will create bottlenecks in application response due to high utilization levels and put the appropriate remediation efforts in place to reduce these risks.

With this service level visibility, management can much more effectively identify and analyze higher-level trends. As opposed to just looking at a specific executable in a vacuum, sound effective capacity management platforms enable easy correlation, for example, to see that a specific Java-based executable is associated with order entry, and so underpins a business service that directly contributes to the organization's revenues and profits. Consequently, the service levels of the executable, and the resources employed to support it, can be closely managed.

With these capabilities, managers can model workload growth, so they can start to identify trends that may introduce bottlenecks. They'll be able to see for example, how in one department, a new hypervisor is effectively collecting dust, while in another area, users are complaining about a service being slow because the underlying server is continually running at 85% of capacity.

Step 3: Leverage scenario planning capabilities

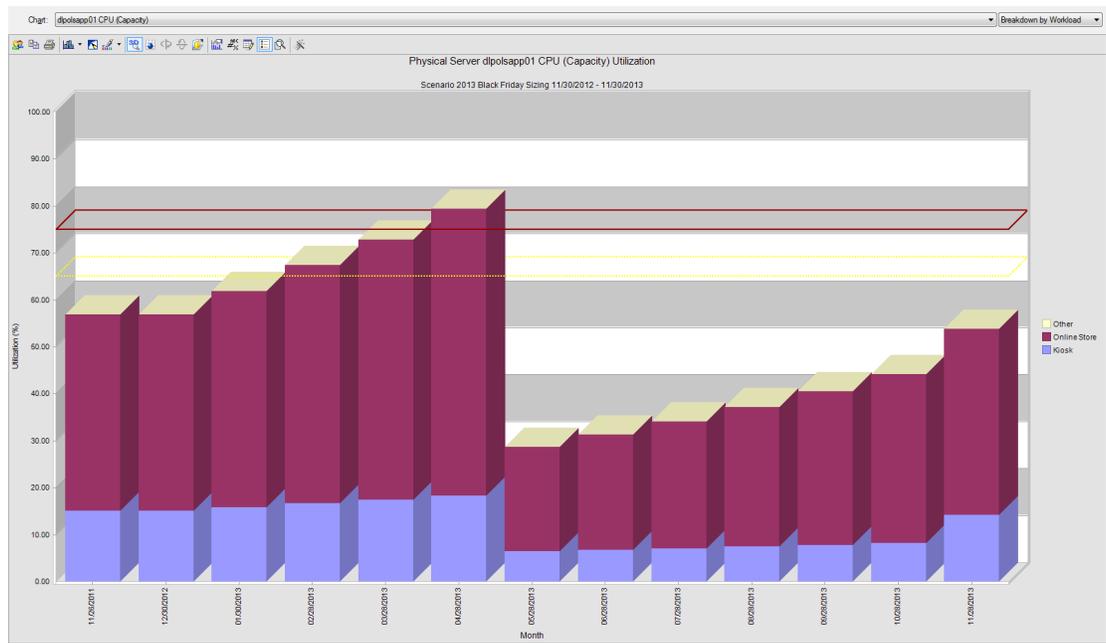
As organizations amass comprehensive, enterprise-wide capacity management data that spans significant lengths of time, they can begin to much more accurately identify patterns of behavior and trends, and use these insights to make more accurate predictions in terms of future capacity requirements. These ‘what-if’ modeling capabilities enable managers to not only understand the impact of hardware upgrades, virtualization efforts or cloud migration initiatives but also assess the impact of planned growth or changes in service demand, identify trends, discover bottleneck components and evaluate remediation efforts. All in an effort to minimize the disruption to the customer experience. Application performance management data can also be leveraged, which can help improve automation of workload definitions and profiles; providing input to modeling capabilities with real transactional data to right-size infrastructure environments for workload peaks taking into account throughput and response times while reducing risk of capacity related performance issues.

For example, at a global consumer packaged goods organization, IT executives implemented a sophisticated capacity management service to more intelligently manage their capacity plans. These executives were responsible for managing a large SAP implementation, which was used to support a range of business-critical processes, including order tracking and fulfillment. If the infrastructure wasn’t sized appropriately, bottlenecks may have arisen that could have had a significant bottom-line impact. Leveraging their capacity management platform and SAP application data, they ran hundreds of what-if scenarios, for example, assessing what would happen if they moved a front-end Web tier to a data center in Japan, or if they migrated 80 physical servers to a virtualization platform. Through these accurate models, users could better predict how business initiatives would impact capacity, and how to effectively support initiatives so they would have optimal chances of success.

Figure B.

Prescriptive modeling capabilities

Workload increase simulation and the appropriate upgrade effort to remediate possible service outage.



Step 4: Leverage business data

IT organizations should continue to work to establish the visibility needed to do capacity management of business services. To do so, they need to start leveraging not only IT data, but also business data, such as sales forecasts, hiring plans, and more, to factor into capacity management.

By comparing changing workloads against changes in the business environment, managers can bring entirely new levels of intelligence to bear in understanding evolving capacity demands. IT managers can more accurately forecast, for example, how a 20% increase in new customer wins will affect demand on the organization's order tracking system, or how, if the number of Web portal users exceeds 10,000 per hour, risk to the user experience increases as performance levels increase to approaching thresholds.

With a sophisticated capacity management practice, executives can start to view the impact of demand spikes within the context of business and technology influences, and better position their organization to anticipate and address their changing capacity demands.

Step 5: Leverage data from across the technology market

Over time, as organizations enhance their capacity management capabilities, they'll be able to aggregate vast amounts of data that will be invaluable in understanding the performance and capacity of the existing infrastructure and services. However, how do you make predictions if, for example, a new set of storage systems or a new cloud service is about to be employed? To evolve capacity management to the next level of sophistication, organizations need to be able to harness not only data center component data, but also industry metrics on alternative technologies and cloud services.

Consequently, advanced IT organizations need a capacity management platform that features a component model library, a repository defining the capacity properties of almost every permutation of hardware and hypervisor platforms, operating systems, and even mainframe and major cloud services available on the market today.

Armed with this kind of information, IT teams can begin to do detailed what-if scenario analysis for how both current and new technologies will accommodate emerging service requirements, so organizations can more effectively plan technology migration options, and more efficiently manage their execution. For example, if an IT organization is being tasked with wringing more value out of technology investments, executives may want to explore migrating to a new virtualization platform or IaaS vendor that purports to offer significant cost savings. Based on industry metrics of these virtualization and cloud alternatives, the IT team can build detailed models, including specific performance metrics and optimally sized VM template configurations, in order to do a detailed cost comparison, identify the most cost effective alternative and ultimately provide ample financial justification to senior leadership in proposing the move.

Step 6: Implement continuous optimization and improvement

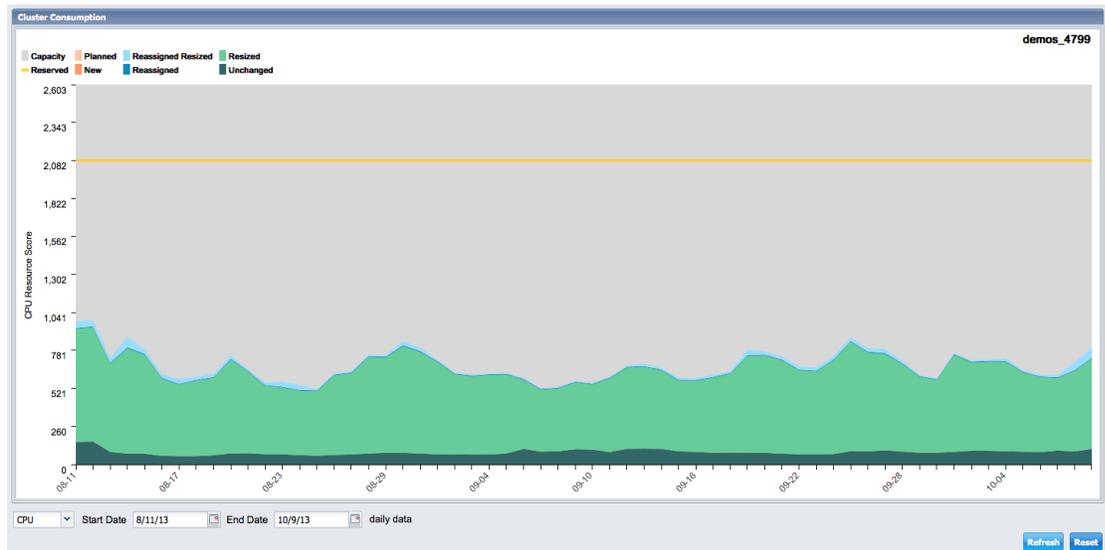
Capacity management is an activity that’s never truly finished. Even after a successful initiative has been completed, the ever-changing demands and technology landscapes in IT are such that organizations need to continuously manage capacity and performance in order to optimize and continue to properly run applications and IT based business services.

It’s critical to deploy ongoing processes not just for enhancing infrastructure capacity and application delivery, but to institute ongoing processes for optimizing modeling and forecasting efforts. For example, once a new deployment has been rolled out into production for some time, it’s important to compare the performance of this new implementation with the performance levels predicted by the models, and determine where the predictions were on the mark, where they weren’t and why. The insights gathered can help managers make appropriate adjustments in the calibration of predictive models, and help foster more intelligent analysis and predictions across the organization’s spectrum of deployments and initiatives.

Figure C.

Continual optimization of production

A right-sizing effort for a production virtualized cluster in order to ensure IT is continuing to meet business objectives and budget goals.



Section 3:

The Benefits of Business – Centric Capacity Management

After completing the steps above, organizations will be able to harness sophisticated capacity management, and enjoy a range of significant benefits:

- **Reduce costs and more fully leverage existing investments.** Effective capacity management can set the foundation for continual right sizing and resource optimization. By enabling the reduction in under-utilized resources, capacity management helps organizations eliminate unnecessary expenses, and gain maximum value from existing investments.
- **Optimize infrastructure performance.** Effective capacity management enables decision makers and administrators to ensure that critical business services receive the infrastructure capacity and performance required, even as demands change. By gaining the ability to better predict future capacity needs and resolve potential bottlenecks, they'll be able to avoid outages as well as having to react to implement rushed procurements and changes.
- **Support innovation, without compromising service levels.** Capacity management can be instrumental in helping organizations successfully harness new innovations, support new strategic initiatives and execute platform and service migrations. At the same time, capacity management can help ensure that, before, during, and after this evolution, the business doesn't have to be exposed to any compromises in performance or availability.
- **Boost IT's competitive position.** With effective capacity management, the IT organization can more effectively align investments with business plans, reduce costs and barriers, and speed time to market, which enables IT organizations to become more competitive against external cloud offerings and service providers.
- **Improve internal operational efficiency.** Today, many IT staff members spend huge amounts of their time on the manual effort associated with report generation and analysis, including running reports, normalizing and aggregating multiple data sets and so on. This effort represents a huge ongoing time commitment, and a continued distraction from the organization's key strategic objectives. With effective capacity management platforms, organizations can significantly streamline this effort, and so boost internal operational efficiency.

Section 4:

Conclusions

Capacity management is a critical endeavor today, and it's only going to continue to grow more so in the months and years ahead. Fundamentally, organizations need to get more out of their IT investments and services, and capacity management is the way to make that happen. By taking the six steps outlined in this paper, IT organizations can begin to establish the comprehensive, intelligent capacity management capabilities they need to more effectively address their operational and strategic objectives.

Section 5:

The CA Technologies Advantage

CA Capacity Management harnesses unique predictive IT management capabilities to deliver effective insights and decision support. The solution combines real-world performance data, modeling, simulation, financial information and intelligent reporting that enable accurate and dependable projections of the performance and service levels key applications and IT infrastructure will deliver. The business insights derived from this unique set of inputs give you the information needed to help effectively plan capital budgets, make investments that deliver maximum innovation, avoid costly downtime and manage risk across your portfolio of IT applications.

With CA Capacity Management, you can leverage your existing monitoring capabilities to get the prescriptive answers and business justification you need to make informed investment decisions. No additional performance monitoring agents are needed, which means you can deploy the solution while minimizing costs and potential risks to your production environment.

CA Capacity Management can also help you get more out of your investment in CA Application Performance Management (APM). By leveraging the solution's real application performance data, CA Capacity Management enables you to perform simulations of a variety of load patterns across a number of architectural options, so you can easily identify the best-suited environment configurations. The solution enables IT to deliver the right-sized infrastructure that supports business critical application delivery and helps to consistently meet or exceed user expectations.

Critical differentiators

CA Capacity Management provides:

- Enterprise level scalability and performance that has been proven to support the world's largest IT environments.
- Capacity prediction for your multi-tiered applications, including those running in mainframe, distributed and cloud environments.
- An enterprise-class centralized repository that collects and normalizes performance data from multiple sources, without the need for agents or additional instrumentation.

- An extensive catalog that allows you to compare and evaluate alternative platforms. The catalog encompasses over 7,000 infrastructure components and services, including public and private cloud offerings, operating systems, hypervisors, distributed systems, and mainframes.
- A robust normalization methodology that can aggregate performance data and enable you to understand new hardware's power, capacity consumption and workload impact.
- An advanced modeling function that can perform non-linear growth calculations, so you can assess the effect new demands will have on available capacity.

CA Capacity Management uniquely provides advanced scalability and capacity analysis with end-to-end, cross enterprise support and capabilities in order to provide the prescriptive insight needed to make informed business decisions.

Next steps

To learn more about CA Capacity Management solutions, please contact your CA Technologies representative or visit ca.com/capacity.



Connect with CA Technologies at ca.com



CA Technologies (NASDAQ: CA) creates software that fuels transformation for companies and enables them to seize the opportunities of the application economy. Software is at the heart of every business, in every industry. From planning to development to management and security, CA is working with companies worldwide to change the way we live, transact and communicate – across mobile, private and public cloud, distributed and mainframe environments. Learn more at ca.com.