



IT COST REDUCTION

Maximizing your initiatives with analytics

To fully satisfy not only cost, but essential risk and impact implications, IT organizations should seek to base their decision making on informed, quantified facts ...

Introduction

Businesses worldwide are facing a series of tough operational choices in the current economic climate. Until recently, IT organizations were enjoying a period of relatively stable budgetary increases as confidence in IT-enabled processes and innovations grew amongst business leaders. But given today's uncertain climate and the knock-on effects this presents to the bottom-line, many businesses will have no choice but to cut IT spend.

Navigating the crunch

The reality this brings is that CIOs will need to apply a more selective approach to their planning, with water-tight business justification that identifies only the most valuable innovations and cost-cutting initiatives. But although today's IT teams have made great strides in understanding their businesses' strategic needs and aligning IT services toward them, many still lack an "across the board" visibility into their estate—to easily identify where inefficiencies exist, where costs can be trimmed and, most importantly, where new investment should be targeted.

Visibility to improve decisions

Although most IT organizations will have the capability to compare and benchmark the like-for-like cost savings of initiatives under consideration, many will find it difficult to calculate the combined effect of how the initiative will result in terms of risk and impact on the business—and this is vitally important for a number of reasons. For example, when compiling a business case on the expected cost savings that could be obtained from datacenter consolidation, traditional approaches will typically fail to fully consider the risks and impact that could potentially result, such as forecasting the user-perceived application performance, latency and throughput, after the change. Without these key considerations, the consolidation might achieve its initial cost saving, but can be just as quickly swallowed up if, for example, the deployment results in IT performance degradation and damages the business' ability to function.

Therefore, to fully satisfy not only cost, but essential risk and impact implications, IT organizations should seek to base their decision making on informed,

quantified facts that relate the initiative's cost saving and value potential to the current operating IT environment and business objectives. It is with this requirement in mind that growth in the specialist discipline of analytics is increasing in use and reputation.

The role of analytics

Analytics provides the quantified evidence that enables IT organizations to fully maximize the successful outcome of cost management and reduction strategies. Unlike rudimentary approaches to cost reduction that only provide simplified cost comparisons, analytics provides the necessary multi-dimensional analysis of all three considerations (cost, risk and impact) that are required to ensure initiatives successfully realize their savings without compromising service quality.

By capturing and combining data from across the IT estate, business objectives and running costs, analytics establishes a "big picture" model of the current working environment, to expose the hidden correlations that indicate where change will bring about the most benefit from cost reduction, risk and impact perspectives.

Baselining and scenario modeling

To understand where best to concentrate effort, it is important to gain a first, baselined view of the IT estate to understand where initiatives will provide the most worthwhile returns. Analytics has the flexibility to take a platform-by-platform approach to baselining, taking samples of data from each to ascertain the current levels of performance and capacity, as well as current and forecasted running costs. By building these baseline assessments, analytics enables organizations to fully understand the "before" and "after" pictures of their cost reduction strategy through the use of advanced scenario modeling and change analysis. The advantage of using this approach is that organizations can model the comparative outcome of change before applying it, enabling initiatives to more likely achieve their cost reduction targets.

Analyzing estate utilization

To aid IT organizations' understanding of how an analytics approach identifies the most valuable cost reduction initiatives, this paper outlines commonly found scenarios platform by platform, with actual case studies from Sumerian customers. However, as datacenters form the largest concentration of IT, combining many platforms of the estate, we have addressed this area as an individual concern.

Datacenter consolidation

In the race to satisfy business demand, many organizations are faced with managing inefficient and siloed application architectures that heavily consume datacenter resources. Energy costs alone are fast outstripping other expenditure, and with technology improvements such as virtualization, SaaS (software as a service), and more energy efficient hardware on offer, datacenter consolidation rides high as a cost cutting imperative. But although the cash and environmental rewards can be great, so too can be the planning pitfalls. Two key considerations are particularly vital: the first is the assurance that any planned change won't negatively impact user-perceived performance, and secondly, that capacity is adequately addressed and not simply a "find and replace" exercise that falls into previous traps of over- or under-provisioning.

Cost savings can be garnered from an array of choices—from consolidating multiple datacenters, moving to cheaper locations, application migration and rationalization, or virtualization, to fully outsourced utility computing. But common to all initiatives is the need to ascertain current application utilization, performance and latency measurements, so that the new chosen datacenter architecture can be suitably provisioned—bearing in mind the cost, business impact and risk implications. To achieve this, analytics captures application performance data at the packet level to quantify the current state of performance, taking into account time of day variations. Through the application of cluster analysis, common interactions and end-to-end application behaviors indicate whether network delays will occur post change. By then applying scenario modeling, analytics can effectively predict the most beneficial datacenter architecture to satisfy current and future demand—supplying precise sizing and capacity requirements, and qualified calculations on what the associated costs savings are likely to be.

Case study 1: Datacenter consolidation

In a datacenter consolidation involving virtualization, Sumerian worked with a public sector organization to determine their consolidated architecture requirements by baselining the existing environment and running scenario modelling for a virtualized estate. The resulting analysis enabled the organization to reduce its server estate by 35% and reduce energy costs by 20%.

In other analysis for a large services company involving the consolidation of three datacenters into one, Sumerian established a base assessment of the current architectures and demand to accurately determine the sizing requirements for the consolidated estate—ensuring end-user performance would be maintained and capacity adequately provisioned. The final set of recommendations resulted in a consolidation initiative that will deliver not only significant cost savings, but productivity and environmental improvements by dramatically reducing server numbers and their associated energy requirements.

Network

For network optimization and rationalization initiatives, analytics determines the capacity needed to achieve acceptable levels of bandwidth after the change, taking into account future growth. By analyzing the impact of application latency between locations at various points across the network, analytics not only uncovers network bottlenecks, but also ascertains where bandwidth is over-supplied. For convergence/VoIP initiatives, analytics captures the current usage of data from IT services and legacy PBX systems to determine the required levels of bandwidth for both data requirements and voice quality (codec) standards. By then applying various architecture scenarios under consideration such as IP telephony and IP trunking, analytics uncovers the true costs, risk and impact of deploying such a solution.

Case study 2: VoIP convergence

In analysis involving network convergence, Sumerian worked with a large professional services company who was seeking to integrate voice and data onto a single MPLS network. Sumerian baselined the existing voice and data usage by analysing data from various network probes, determining average and peak busy hour bandwidth utilization, cross correlating this with one year's call records from PBXs to calculate busy hour call loads. Using models for various VoIP codecs under consideration, Sumerian could then determine the necessary bandwidth and provide a capacity estimate for

both the physical bearer bandwidth and the class 1 & 2 bandwidth required at each site. Armed with these results Sumerian was able to predict the conversion, upgrade and operational costs based on figures provided by various suppliers under consideration. The results showed a massive difference in the cost-benefit picture depending on the option chosen, with the most expensive supplier being three times more expensive than the least expensive. The eventual outcome was a business case that gave the customer confidence in the cost and quality of their chosen solution, resulting in an overall operational cost reduction of 16%.

Storage

Gaining the right evidence to quantify individual service and macro capacity requirements for enterprise storage is a complex exercise, but is one that can be dramatically improved by applying analytics. Storage capacity requirements can be influenced by the outcome of other consolidation activity, such as database and application rationalization; however, regardless of this, many organizations will be using an inefficient mixture of storage architectures that will always be sub-optimal from a cost standpoint. Consolidation opportunities may arise from direct attached storage (DAS) migration to network-attached (NAS), or centralized storage area networks (SAN), depending on individual application requirements.

While capacity is clearly an issue in storage consolidation, risk and performance are also paramount. Understanding the impact of a failure in a shared storage system on services and applications is crucial, as is ensuring the system has enough throughput capacity to deal with rate of data being added and changed. In addressing solutions under consideration, analytics can determine either individual application/service or macro capacity requirements by capturing the actual levels of usage from across the business and scenario modeling expected data growth rates. The same modeling approach can also be used to size compliance legislation archiving and disaster recovery requirements, again ensuring capacity is not overprovisioned, but configured to ensure acceptable levels of service and that network bandwidth can cope with forecasted increases in traffic.

Databases

Cost savings from database consolidation and rationalization can be achieved through a variety of strategies, for example, standardizing existing systems, through retiring of legacy ones to rationalizing database workloads. The difficulty, however, is determining the level of usage from each system so that a convincing case can be compiled for busi-

ness unit sponsorship. Through data capture and cross correlation with actual business consumption and demand, analytics is able to reveal not only which databases are commonly accessed across the business, but where requests and workloads are coming from, and over which period of time. Armed with this precise quantification, accurate modeling can identify the outcome and impact of any database changes under consideration, determining whether cost savings are likely to be achievable for each chosen scenario.

Applications

Cost savings from improved application management and deployment can have a profound effect on releasing funding for new IT innovations. Many IT organizations are faced with a legacy set of applications that are a drain on the budget, but lack the necessary understanding on their usage and architectural setup to recommend their standardization or retirement. By applying analytics and User Profiling analysis, an accurate understanding of application usage can be reliably formed, mapping out which users across the business are accessing particular applications, and what infrastructure is utilized to do this. The resulting intelligence can help application owners to work with infrastructure teams to re-engineer application architecture and tailor applications to particular user groups, optimizing IT cost per head and improving deployment times.

Mainframe

Although mainframe technology has been around for over forty years, it remains an important component of many enterprises' computing needs. Mainframe services are often outsourced and, therefore, cost savings can most likely be garnered by assessing the current usage, typically paid for per CPU cycle, and licensing arrangements, whether they are unsuitable or over provisioned. By understanding their current level of usage, IT organizations can take a proactive position in negotiating more favorable arrangements with their current supplier, or enter negotiations with other suppliers using a precise set of requirements based on actual usage metrics.

Additionally, analytics can identify inefficient code changes to mainframe applications—an issue that can heavily impact resource utilization and result in expensive upgrade costs. By establishing a profile that normalizes mainframe performance and utilization parameters against transaction rates (demand) and application code changes, analytics can enable development teams to understand where inefficiencies in newly released code are pushing up utilization, therefore, facilitating improved code configuration management practices and cost savings.

Case study 3: Right-sizing applications

A Sumerian customer in the professional services sector was planning to replace a number of in-house ERP applications with a European-wide SAP rollout. The challenge was identifying the features, functionality and target users so that the SAP solution could be appropriately targeted and developed with minimal risk and cost. Sumerian's User Profiling analysis identified the key users and features of the existing applications, allowing the customer to target training and development resources to the correct areas first time, thereby minimizing costs and maximizing savings. Sumerian subsequently helped manage the acceptance testing of the SAP deployment by determining target SLAs for specific functions based on the existing applications, and analyzing the SAP performance against those SLAs, thereby ensuring a successful deployment.

High-performance computing (HPC)

For organizations running HPC/grid computing in-house, ensuring utilization is maximized and not over provisioned is vital. Grids sitting idle 50% of the time will have extremely high running costs, and the pay-off to owning in-house, therefore, has to be questioned. Improvements to security have made the outsourcing of HPC a viable option, even for security-sensitive businesses—so for organizations seeking to reduce costs, outsourcing may prove more cost effective. In exposing the correlations between grid usage and business demand, analytics enables a precise quantification of its performance, capacity and costs, enabling IT organizations to optimize application architecture via scenario modeling.

Case study 4: Optimizing grid processing

Sumerian worked with the investment arm of a large UK retail bank that was experiencing a number of issues with the capacity and performance of its critical trading services. Due to volatility in resource utilization and capacity, the trading systems had experienced a number of application problems, all of which negatively impacted trading performance and increased running costs. By applying analytics techniques to the bank's grid processing function, Sumerian revealed that there was sufficient capacity to complete the overnight batch process more quickly, but that the database server would quickly become a bottleneck as increased grid activity had a heavy impact on database load. As a result, the IT team has reconfigured capacity at numerous points within its trading services, improving overall performance and capacity, and optimizing running costs.

Ongoing management

Once cost reduction initiatives have been deployed, analytics has the ability to track their relative progress and business impact by taking regular samples of performance data for ongoing management and optimization. The main benefit in using analytics to monitor ongoing service delivery and its associated running costs is in its ability to dramatically reduce outages and identify opportunities to improve application architecture. By identifying deviations from the baseline normal profile, analytics uncovers long-term emerging trends that indicate instability from impending issues. This baseline view can also be used to optimize existing infrastructure and "sweat assets" by identifying opportunities to improve application architecture and capacity provisioning in areas such as load balancing. By carrying out steady, pragmatic improvements to application architecture, cost savings will be garnered from optimizing what is already in place and dramatically reducing outages and their associated support costs.

Realizing cost reductions

Today's testing financial uncertainty is having an undoubted knock-on impact for IT budgets. But where adversity appears, opportunities arise in equal measure. In order to maximize the potential of a constrained budget, CIOs should seek to gain an "across the board" view of their IT estate, to reevaluate its architecture and performance so that it can be optimized and release essential funding for new investments and innovation. Without adequate understanding of where inefficiencies lay, CIOs risk investment into areas where poor business value will result. Instead, analytics' quantitative approach to baselining the existing environment and scenario modeling using all three requisites outcomes (cost, impact and risk) will result in a set of initiatives that realize their cost reduction promises without placing the IT organization or business at risk. In applying an analytics approach to achieve this, CIOs can be confident in compiling a cost reduction strategy that not only delivers value, but optimizes the budget, not as a one-off exercise, but as an ongoing and considered objective.

▶▶ **More information**

To find out more about Sumerian Forward Thinking® predictive capacity planning and using Sumerian Capacity Planner, just give us a call on 0131 226 9300, drop an email to sarah@sumerian.com or visit our website at www.sumerian.com