Benefits and Methods for Integrating IT Service Utilization Data with a CMDB

Executive Summary

Configuration Management Database (CMDB) projects aim to address many issues with IT service provision. However, rarely are they seen as a way to directly reduce the amount spent on IT. Instead the offer softer routes to cost containment based on efficiency improvements, more predictable service, reduced unplanned outages, more satisfied users etc. By introducing Asset Utilization information into a CMDB, executives can get direct cost savings through the identification, recycling or retirement of under utilized assets. Such information can also be used to maintain momentum in CMDB projects by showing early and measurable financial successes. In addition, the shift to cloud-based models of computing for datacenters, i.e. the so called private-cloud model, absolutely requires that executives can see the extent to which the fixed assets, within the service provision, are being used in order to ensure on-demand charging models are appropriate

This white paper details how Service Utilization intelligence can be blended with CMDB projects; and how organizations can use this superset of information to dramatically and rapidly take non-labor costs out of IT while ensuring service is maintained at optimum levels.

Types of Asset Utilization Data

Service Monitoring

Service monitoring tools are excellent at detailing how hardware resources are being exploited and identifying bottlenecks that can degrade service. Over the past decade these tools have been enhanced to the point where end-to-end service can be monitored and mapped onto discrete IT assets in order to facilitate root-cause analysis. Root cause analysis is the holy grail of IT performance management. It can spot service degradation before it affects the user and identify the failed or failing components causing the degradation. When combined with CMDB capabilities, identification of these failing components can flow through other service definitions within the CMDB to ensure service desk personal are equipped to pro-actively deal with other services that may be impacted.

Capacity Planning

Capacity Planning tools use information gathered from Service Monitoring tools to plan future expenditure on IT equipment based on expected workload growth and the historical impact of certain workloads on IT asset utilization. These tools traditionally focus on expenditure associated with server hardware and software and network equipment. Rarely will a capacity planning tool concentrate on end-user IT asset requirements.

Service Utilization

Service monitoring will not tell you the extent to which the service is being used by the business, and therefore the business benefit it is conveying. In all application environments there will be an ultimate consumer of the business benefit. This can be a person or another system, and that consumer should be benefiting from the investment in the application system. Measuring the return on this investment is a difficult when viewed from the system out towards the consumer, since it is very difficulty to correlate workloads with amount of time individual or groups of consumers spend user the service. For example, it may be that a system, deployed to 500 users, is generating 80% of its workload from 20% of those 500 users (the 80/20 rule often encountered in IT systems deployment). When viewed from a service monitoring perspective everything is OK and the system is operating as planned, however the organization may be paying for infrastructure to support 400 users unnecessarily. It can be argued that certain infrastructure items exist regardless of the actual number of users exploiting a business application that depends on this items; this is no doubt true of certain classes of server and network device, but the closer the purchase gets to the end-user the greater the scope for savings. Software licensing is perhaps the most well known example of this.

Usage Auditing

Usage auditing is mentioned here for completeness, but is rarely considered as part of a CMDB project. Usage auditing is more event driven discipline that forms part of security or efficiency studies; often part of human performance analysis. It is mentioned because there is an overlap in service utilization and usage auditing in some tooling. We will not be considering the issues of usage auditing integration with a CMDB as part of this white paper.

IT Cost Management

IT Cost Management is the discipline of accurately associating the costs of IT assets with the business services they support. Using the inherent relationships in most CMDB structures, the individual costs of each discrete Configuration Item (CI) can be aggregated according to the services they support. For example, if a particular business application runs on a cluster of web application servers backed by a single Oracle database and is delivered to a user department over a specific network topology, it should be possible to add the costs of each of the CIs (routers, WAN links, servers, database software, application software etc), which support the application, to determine how much IT asset expenditure it is consuming.

Private Cloud Considerations

The View from The Datacenter

For organizations considering provisioning IT assets to the business in the form of a private cloud the relationship between what those physical assets cost and how they are being used is essential, because private clouds break the traditional chain of association between IT asset and consumer. It is not correct in a cloud model to say, "This business service uses these specific IT assets". Therefore it is not possible to allocate the cost of an IT service to its consumers in the traditional sense. Instead a system of classification is required, largely determined on an organization-by-organization basis, for the types of IT assets available to be consumed by the business, according to certain characteristics they may have. For example, servers can be grouped into small, medium and large or high-availability and standard availability. Database resources may be classified by some of the headline features they have such as clustered database, transactional database, data warehouse and so on. In all cases a usage metric for each of the cloud items will need to be associated with a cost, and that usage metric must be easily measurable and understandable. If the metric is not understandable, end-user groups will resist it when presented with the bill.

The View from the Business

For consumers of cloud resources the issues are equally as challenging. Common charging models for cloud-provisioned assets will be based on a fractional resource consumption model. If this charging model is presented in isolation it will ultimately create work for both the business and IT as cost-reduction initiatives are pursued, because it does not represent the business view of an on-demand model. Business costs for IT in an on-demand model are increasingly looking at a cost per consumer time unit. For example, charging based on the time of user hours per application.

A typical on-demand charging model, as mentioned above, will be based on a count of "compute" resources. A compute resource is an arbitrary definition that includes costs for hardware and software time-slices across users of those resources. It is virtually impossible, in any systematic way, to get down to a level

of granularity that a typical business user would require using traditional methods of service monitoring. The reason is quite simple, in most organizations server computing comprises Commercial Off The Shelf (COTS) tools and homegrown applications. The number of applications that run server-side can easily run into the hundreds and may span dozens of different implementations. To implement technical introspection of the server application portfolio down to the user level in a consistent way, across all implementations is impossible. However, it is certainly possible if the problem is approached differently, which is where Service Utilization monitoring comes in.

Linking IT Cost Management to ROI with Service Utilization

Service Utilization looks at the usage of services provide by IT from the perspective of the consumer. While some IT services have no user dimension whatsoever – batch jobs are a great example - many do. And where they do, it is possible to determine the extent to which individuals, groups, departments, business units (or any other arbitrary grouping of end-users) are using the application, in a consistent, low-overhead, straightforward manner using Service Utilization tools such as Survey™ from Scalable.

By measuring the amount of time a user interacts with a service at the point of delivery, which in most cases means the workstation, it is possible to calculate the relative usage of the service across the various groups using the service. The definition of what constitutes a service is something the business and IT groups can agree on up-front and the utilization can be aggregated accordingly. For example, and service could be considered to be something as granular as an individual web-page buried deep in a web-application that exports data to a "fat-client" CRM system, or as high-level as "all MS Office Applications provisioned via Citrix". In both cases, and everything in between, both IT and the business can agree a common definition and see the extent to which each discrete group of users is using the service relative to each other. All costs associated with service, which as we have seen should be measurable by virtue of the CMDB can be allocated fairly and understandably.

What the business is now able to do, far more easily, is factor IT costs accurately into the ROI models for the business services used to generate revenue.

Integrating Service Utilization Data with a CMDB in this way can provide a very quick win for the CMDB projects and introduce more momentum for the remainder of the implementation program.

What follows is a more technical explanation of how this can be achieved.

Practical Considerations

This section deals with the practical business of integrating Service Utilization data with the CMDB in order to extend the ROI features of IT Cost Management. By way of example it focuses on the particular implementation of Service Utilization within Scalable's Survey™ tooling. We'll look at how data can be integrated with a CMDB and how a good CMDB would take advantage of that information.

Survey™ Service Utilization

Survey™ is Scalable's Service Utilization Intelligence tool. It can measure in detail the extent to which human's are interacting with applications, and the impact that utilization has on the underlying hardware. Survey can measure usage of all types of application; including traditional desktop, virtual desktop, thin-client, browser-based, java-based, virtual applications and more. It can determine whether the user is power-user or simply an ad-hoc, read-only user and it can drill down into application component use. It can group discrete application use into high-level service use in order to match the charging and ROI models of a customer.

Integrating with a CMDB

CMDB schemas rarely have native support for service utilization information. This omission is not a shortcoming; it is simply that as originally conceived a CMDB has no need for dynamic, near real-time, system-derived data. Consequently, to gain benefit from service utilization data it is important to be in possession of a CMDB that supports scheme extensions. Most modern CMDB tools have the capability, as the current conventional wisdom for CMDB implementations is one of federation; where no single CMDB holds all the configuration information (CI) for an organization.

Logically, software is the deciding factor when considering how IT assets are being used. Without running software the equipment lies idle. Most CMDB's will identify software as a CI. However, this is usually a static footprint of a software installation. In a world of virtualized applications, cloud-computing and browser based computing these existing software representations are inadequate since an increasing percentage of the applications that run on hardware are not actually installed there. Instead Scalable advocates two fundamental new classes of CI called Running Task Definition (RTD) and Running Task Instance (RTI). In some cases a RTD will also be represented in the CMDB by an installed software CI, but in many important cases it will not. A single RTD entry will exist for each unique running task encountered with the estate. Naturally for a large organization we are talking about many thousands of CIs but crucially this number will certainly be a lot less than the number of installed software CIs within an organization. For example there would only be one CI for Microsoft Project 2007 in an organization, regardless of how many machines have the software installed. RTI CIs represent instances of RTD usage and link to the CIs that define both the user and machine upon which a running instance of this RTD

was encountered. In Scalable's implementation there is a maximum of one RTI CI per RTD per user per machine that would be refreshed as often as the minimum measurement period required

The concept of a measurement period is important and is largely organization and business purpose dependant, and requires further explanation. Suppose the business requirement is to measure usage of assets over the prior day, week, month and year. The RTI would need to be refreshed each day, and would have a revised set of figures for asset utilization for each of the measurement periods.

Of course if no running instance of the RTD were encountered there would be no RTI.

Again, although the data volumes are not trivial, they are still not materially different to the total data volumes associated with installed software CIs. By way of illustration, how many different core software packages would an average user work with per day? Probably more than 5 but almost certainly less that 15. Also, machines are rarely shared, and when they are it tends to be within specialist parts of an organization or in response to a unique set of circumstances. Consequently service utilization monitoring is unlikely to generate more than 20 new CIs per machine per day.

The data volume issue is not the central issue though, it is the transactional nature of the information. CMDB's are not typically used to store this kind of information and are often deployed in ways that would make a daily update of 200,000 CIs (a reasonable expectation in a 10,000 node organization) a problem. But this must be recognized as a philosophical objection rather than an inherently technical one. An instance of a running task in the world of virtualized applications and cloud provisioning has just as much right to be called a CI as a statically installed application. And in terms of benefit to the organization, the RTI CI has far greater potential. Consequently, we have found the biggest barrier to inclusion of service utilization information within a CMDB is political rather than technical and stems from a traditional, purest perspective on the use-cases for CMDB projects. We therefore recommend strong executive support for based on the scale of financial benefits outlined in an earlier section of this white paper.

Survey can export information, via CSV, that matches the RTD and RTI definitions above. These structures can be configured to use the same key fields found in CMDB structures populated by common network discovery tools. As previously mentioned, modern CMDB's support schema extensions and also the integration of CSV type data on a scheduled basis. Consequently it is a relatively trivial exercise to begin the process of integrating Service Utilization data with other aspects of a CMDB project.

Conclusion

For any executive supporting a CMDB project, it is vital to demonstrate some "quick-wins". CMDB projects can be lengthy and require organizational as well as technical change, which makes quick-wins something of a rarity under normal implementation scenarios. Incorporating Service Utilization data into the CMDB project at an early stage resolves this challenge as it provides an actionable, beneficial model that business sponsors can consume and see real value in.

Furthermore the effort to incorporate this data is not significant. Modern CMDB's are geared up to for extensions; data volumes are not significant in reality, and key fields are common.

The biggest barrier to getting the benefit of Service Utilization data into a CMDB is political; a weak executive sponsor will run into technical objections based on an out-moded, purist approach to a CMDB. However, those with the strength to represent the needs of the business, and carry the support of the technicians implementing the CMDB, will be able demonstrate ROI much faster and secure momentum for the support for the long-term CMDB organization benefits.