The IT Utility Model—Part I

Emlyn Pagden, Sun Professional Services

Sun BluePrints™ OnLine—July 2003

http://www.sun.com/blueprints

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95045 U.S.A.
650 960-1300

Part No. 817-3177-10
Revision 1.0, 6/25/03
Edition: July 2003

Please
Recycle

The IT Utility Model—Part I

Providing mission-critical IT resources for a growing environment requires a significant increase in compute resource utilization. Such an increase is typically limited by capital expenditure. By implementing a well-developed utility model, you can decrease total IT costs and optimize the capital investment of IT resources.

This article is Part I of a two-part series that describes the current business requirements for a utility model, and discusses the current commercial and political issues faced when implementing one. Both financial and technical aspects are covered, from detailing what a utility model is and why it is needed, to describing the mechanism required for capturing compute resource consumption to accurately bill customers.

This article addresses the following topics:

The intended audience for this article is IT Architects, Finance staff, and Executive officers.

Why is a Utility Model Required?

The business requirements for a utility model are:

- Decrease total cost of IT capital expenditure
- Maximize resource utilization
- Minimize resource waste
- Increase the transfer of spending to current expenditure utility charges
- Increase accurate accountability of resource cost against business units

This article is written from the perspective of running a utility model from within either a service provider or a data center, so certain assumptions are made regarding the required technology. However, if an alternative technology meets the business and functional requirements of the solutions defined in this article, then it is capable of being integrated into the environment.

What is a Utility Model?

A utility model is a combination of utility computing and utility pricing that results in minimizing total IT cost expenditures and maximizing IT resource usage. This section defines the utility model components and addresses common questions and statements regarding utility models.

What is Utility Computing?

Utility computing is a method by which an end user consumes the resource of a computer based on utilization rather than hardware ownership. Utility computing solutions are implemented to:

- Collect resource utilization data
- Aggregate and apply a pricing plan to the utilization data
- Create eXtensible Data Records (XDRs) or XML Records

XDRs are synonymous with Telco CDRs (Caller Data Records), for the purposes of billing customers. XML and XDR records are used as a primary mechanism to integrate the billing data into commercially available billing applications such as SAP or Amdocs Horizon.

Utility computing and utility pricing can be differentiated by viewing one as a technical solution and the other as a financial solution. If a business does not employ a financial structure to support utility computing, the success of such a technical solution is limited.

What is Utility Pricing?

Utility pricing is a method used by a supplier of a computing resource to bill a customer based on resource utilization. This method may be used alone to support the requirements of a business that chooses not to employ a utility computing architecture to capture resource consumption. In this case, customers are charged for resource based on a financial model that employs previously agreed estimations of resource utilization.

Utility pricing provides a customer with a financial plan to enable a monthly cost for using a resource based on an agreed usage. There are caveats for this pricing structure that include an up-front flat-rate charge such as a utility standard charge (USC) (often seen as the Telco equivalent of line rental), and the possibility of over- or under-utilization. Because a financial solution might not include a data capture or metering technology solution, estimates are put in place to accommodate peaks and troughs in resource consumption.

Therefore, the supplier will need to recover the costs over a given time period, and is likely to seek to do so within the agreed contract period. This implies that the contract period needs to be set with capital cost recovery as one of the factors.

To summarize, as an IT supply function you can introduce a utility pricing solution without utility computing. However, with the lack of a financial structure in place to support utility pricing, it might not be possible to successfully realize the financial benefits of implementing a utility computing solution.

Maximizing IT Resource Utilization

The primary components of maximizing resource utilization are:

- Resource Analysis – Clarify who, where, when, and what is consuming resources
- Consolidation – Based on the outcome of resource analysis, allocate IT resources to maximize utilization and generate a *spare pool* of resources that can be used for future capacity
- Ownership of IT resources – Move ownership from the individual business units into the IT supply function (either a service provider or internal IT department); therefore managing overall capacity for the organization.

There are prerequisites for maximizing IT resource utilization. The most important prerequisite is persuading the business that the IT supply function is capable of accommodating this task. The IT supply function must be changed to not only provide technical support, but to also provide business and service support to maximize IT resource utilization. This task, discussed in the next section, can enable the IT supply function to be both technically and commercially astute. The IT employees must understand the business and service level requirements for supplying resource in addition to understanding the latest available technologies to satisfy service delivery.

Integrating the Utility Model Technology

A technology stack is used to deliver utility computing. It is important for customers to understand the impact that utility computing and utility pricing have on a business. To successfully integrate a utility model, the business must first implement the financial structure and processes.

For a service provider, integrating a utility model is fairly straightforward. A service provider builds an infrastructure, connects a utility computing architecture to the hardware and software resources, and introduces a price plan to sell the resource to a number of customers. However, integrating a utility model into an organization with multiple business units that require IT support, can initially be both financially and politically sensitive from an internal point of view. This area is discussed during the remainder of this article.

One route to alleviate this possible issue is to leverage service level management to allow discussion between the suppliers of IT and the business users. Using deployment / development methodologies that explicitly define the required service levels (such as SunTone^{SM}) ensure that the language used between business units and infrastructure providers is objective and not subjective, which permits the joint

discovery of common goals. The elimination of subjective criteria by measuring and using service levels, minimizes the scope for bad political and economic conversation. This method provides a mature approach to ensure that the consuming organization wins as a whole and that the divisional goals are subordinate to corporate goals.

The ideal utility model does not require capital expenditure or capital assets in respect to big servers or big storage. In effect, the IT supply function purchases and manages compute resources based on the planned resource utilization of individual business units and not on an assumed mass of resources suggested by the overall business. Thus, the IT function and the business must work together to plan for business growth that may involve expanding existing services or supporting the provision of new business.

However, the IT supply function will face potential problems if they want to introduce a utility model for purchasing compute resource, and offset the cost of this resource by introducing a chargeback solution to the business users.

These potential problems can prove to be significant because most business cost centers or business units, through capital expenditure, are perceived to own the IT resource. They then require the IT supply function to maintain the resource at zero cost. This causes an inherent problem that the IT supply function constantly faces—how to manage the IT resources if proper funding is not provided. In addition, it is often difficult for IT departments to justify the purchase of enterprise management tools that would allow them to manage the IT resources and thus provide a service management solution to the business.

One possible solution to these problems is to shift the ownership of the IT resources to the IT function. Ownership of the resources enables the IT department to proactively assist the business to gauge future growth requirements, and allows the business to concentrate on business-related issues rather than IT specific issues.

With this solution, an organization is able to purchase resources by using the utility model and employ chargeback to the business cost center. This can be achieved by purchasing resources directly from the manufacturer or indirectly through a service provider. To purchase directly or indirectly is a decision that must be made between the customer and solution supplier.

Typically, a manufacturer has a number of key customers that buy services and goods directly. Manufacturers also have some customers that buy goods and services through indirect channel partners. For a utility solution, these partners would be utility service providers.

Chargeback Utility Model

A utility model is used when a customer purchases a resource from a supplier on a utilization basis, such as gas and electricity. If a business has already purchased IT resources on a utilization basis, the method applied to the business unit purchasing the resource is called chargeback. The chargeback method allows the IT supply function to recover costs, hence paying for the resource.

FIGURE 1 illustrates where the chargeback method resides in a utility model:

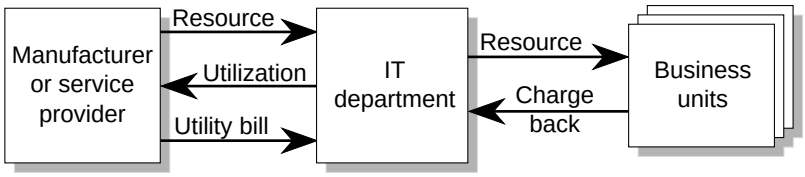

FIGURE 1 Chargeback Utility Model

FIGURE 1 makes it clear that a supply chain for resource exists. For example, in the power supply industry, the grid bandwidth must be able to meet demand; correspondingly, in an IT utility model supply chain, the available IT resources must be able to meet demand. This means that system resources that are deployed as available, but currently unused, have a necessary and valuable role to play in the supply chain. If IT resources are deployed to meet short term demand fluctuations, they have value to the consumer and need to be priced and billed.

System vendors operate a manufacturing 'products offered' business model, and in order to offer a true utility price to a consumer, someone in the supply chain needs to take risk regarding demand and price. It is the carrying of risk that earns much of the reward in implementing a utility model solution; and while some system vendors including Sun, are agreeing to carry some of this risk, it is open to question whether companies wish to lose the opportunity to earn the returns on carrying the risk. It is clear that higher value service providers, including internal IT providers, will wish to carry this risk as a way of broadening their revenue earning capability.

Service Provider Utility Model

For service providers, the utility model works in two ways, providing a possible three-way utility model. The manufacturer sells the resource to the service provider, who, in turn, (through domain technology) provides a utility model solution to a customer of the service provider. FIGURE 2 illustrates a service provider utility model.

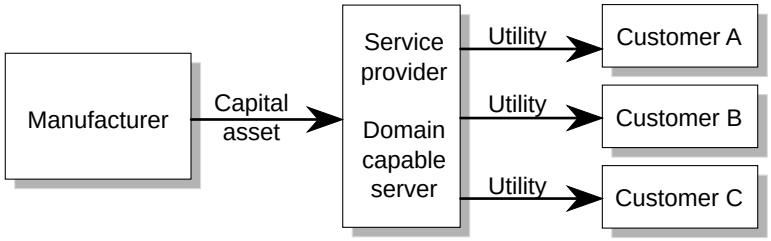

FIGURE 2 Service Provider Utility Model

This method is attractive to all parties involved. For example, the manufacturer is able to sell at an agreed discounted price; the service provider is able to sell the resource at three times 70 percent of cost; and the end-user is able to purchase their required resource at 30 percent less than the list price.

This example is based on a two-million dollar purchase of resource and intends to solve the widespread problem of resource wastage, which is based on a typical customer typically never using more than 30 percent of purchased resource.

On a $2 million purchase of resource, 70 percent waste equates to $1.4 million.

In this example, the customer or business unit is in a position to capitalize on service management, and impose agreed financial penalties on an IT resource supplier that doesn't meet service level agreements. The customer or business unit is not in a position to best manage the utilization of IT resources. This responsibility belongs to the IT function.

If a business unit or cost center (for example of a financial institute) owns five high-end servers each using 30 percent resource, the organization will not realize the full return on investment that is initially expected. Based on the previous example, this would be a waste of 3.5 million dollars.

A utility model affords a way of providing controlled and accountable computing resources for organizations, departments, and business units within a company, or to a customer sourcing its resource from a supplier. However, charging organizations or departments for such utilization has, in the past, proved inherently difficult to gauge accurately. The solutions offered in Part II of this article provide an answer to this issue.

Combining a number of commercially available products to create a robust and flexible control, regulation, and metering solution provides the interface between the business support systems/operational support systems (BSS/OSS) and the service platforms that actually run the business. In effect, these software products extract the resource utilization data, and turn this data into billing information that is both accurate and real-time.

In addition, there is a need to proactively manage the availability of such resources and report the availability data in an accurate and effective manner. The management of availability can be accomplished through the use of resource analysis, workload consolidation, and capacity and service level management.

Capacity and Service Level Management

The use of service level management to support the utility model will be introduced in Part II of this article.

The following products, described in detail in Part II of this article, are used as an example to support the utility computing architecture.

- Solaris™ Operating Environment (OE)
- Solaris™ Resource Manager (RM)
- Teamquest
- Centauri
- Xacct
- Amdocs
- Ejasent

Where the Utility Model Applies

The examples within Part I and II of this article cover a data center and a service provider. However, the following uses are also applicable.

- Residential – A service provider might wish to charge residential users for consuming computing resources through a network device such as a thin client workstation delivering office productivity or gaming applications.
- Supplying to a service provider – The provider of a utility architecture might wish to create a partnership with many service providers, thus providing a utility architecture to each service provider from a remote location. The network between the service provider locations and the utility architecture would carry the utilization data and the return flow would provide the billing information.
- Manufacturer – For a manufacturer to supply hardware on a utility basis provides both negative and positive points. From a positive point of view, this enables the manufacturer to recognize a constant revenue stream. However, it is more suited for the manufacturer to supply resources, at a preferred discounted rate, to a service provider; thus obtaining the immediate revenue of hardware, rather than over a possible three year period. A negative aspect is the investment that is required to support the implementation of a utility architecture and the network involved to extract data from the data center. However, with Sun it is possible to take advantage of the Sun[SM] Remote Services Net Connect infrastructure to achieve the investment and network support.

- Commodity – Due to the pricing and delivery structure that the utility model offers, as a supplier of a utility, you would be in a position to offer the following based upon available components to deliver a quality product or 'IT utility unit':
  - Software unit
  - Network unit
  - Computing unit
  - Storage unit

  In this case, the 'IT utility unit' becomes the 'raw material' so the customer doesn't have to focus on concerns with the technology supporting IT; and instead can focus on choosing the best quality unit that allows the most cost effective way to 'consume resources as and when required' to support the business process.

Software for Implementing a Utility Model

Part II of this two-part series offers solutions for implementing a utility model by incorporating the following commercially available software packages:

- Solaris RM – Resource management tool for the Solaris OE
- XACCTusage – Resource aggregation and mediation tool used in the Telco market
- Amdocs Horizon – Billing and provisioning system
- Ejasent – Resource aggregation, mediation, and billing tool for the utility market
- Teamquest – Performance and capacity management tool
- Centauri – Service level management tool