

## Cost and pricing of shared IT services

Based on discussions at the Common Solutions Group Meeting, 14 May 2009 (Edited and expanded by Jerry Grochow, MIT, 10 August 2009)

### INTRODUCTION

The goal of our discussion was to describe an approach for developing cost and pricing models to support sustainable IT shared services across the multi-university environment. Some of the core principles we feel are important to this discussion are:

- Must reflect the financial goals of the university or other entity providing the service, e.g. cost reduction, recovery, or break-even.
- Must be open and subject to review
- Must be “sellable” within the context of the university financial environment

There are several general considerations that everyone must understand:

- The cost of providing a service and the price that someone pays to use the service may be related but are separately determined. Costs are determined on incremental, partial, or full-costs basis. Price can be set to fully recover cost, to partially recover cost, and may be set proactively (based on an estimate of expected costs) or retroactively (after all costs have been incurred) to effect different goals.
- Different service providers (specific universities or other non-profits) will likely have different goals with regard to pricing and cost recovery. Some may only be interested in recovering the incremental cost of providing a service to another entity, while others will want to recover at least part of their full cost (i.e. contribution toward investments already made in supplying the service to their own university), or, most likely in the case of a separate “for purpose” non-profits, to recover all costs.
- Prices (which, of course, are costs to users of the service), must be acceptable to potential service users within their own context. In the university environment, price basis has to be understandable and seen as reasonable in order to attract customers to the service being offered. It doesn't matter if the price seems reasonable to the service provider based on their financial objectives if it does not also seem reasonable to the service customer based on their financial environment.

### COSTS

There are two basic types of cost models, full and incremental. There will likely be some discussion about what costs to include in full cost models (do you include the salaries for time spent by administrators all the way up the organization?) which may lead to variety of oxymoronic “partial” full

cost models, which we call a “partial cost model.” Since the price of a service is generally determined to recover some or all costs, it is important to be explicit about the cost model being used.

True full cost models include all costs of providing the service as if starting from zero. For example, if a new entity was created to provide data center services, it would have to lease or buy land or a building, outfit the space, hire staff (including managers, financial people, etc.), pay for utilities, network connections, and much more. Universities already have a great deal of that infrastructure, and may in fact not actually charge the central IT organization for some categories of cost (few universities charge for space and utilities, although that may be changing). Even if a university does charge for space, there will be the question of whether the space charge is an average across the campus or a charge based on the cost of building and maintaining a particular building. Determining full cost in this environment isn’t always possible and various estimates have to be made. [Estimates will be acceptable as long as they are explicit and can be justified.] Of course, an entity may not be concerned with the true full cost and can be quite satisfied with only identifying some of its costs. We will call this a “partial cost model” to distinguish it from an “incremental cost model.”

Incremental cost models only look at the incremental costs incurred in providing incremental service. For example, it costs essentially zero to provide email service to one more person if you are already providing email service to 10,000 people (even this can be argued, but it is close enough to zero to be inconsequential). On the other hand, providing email service to an additional 10,000 people will likely require the purchase of additional servers, with all the attendant costs that go along with that (space, power, cooling, maintenance). It may or may not require the hiring of additional personnel. In an incremental cost model, only the costs actually incurred in adding capacity to provide the service would be counted.

Regardless of whether a full, partial, or incremental cost model is being developed, various categories of costs may be included. We created a list of cost categories for the case of providing data center services, although many of these also apply to providing other IT services:

- Network (cabling within the facility, switches and routers, connectivity to the external network)
- Power and cooling infrastructure (power distribution, generators, UPS, cooling units)
- Utilities (electric power, heating, water)
- Software licensing
- Personnel for management and system administration
- Building (rent or construction or renovation)

Capital expenditures, whether for building and renovation or infrastructure and equipment may be recorded in full (“expensed”) at the time of actual expenditure, but more likely (and more usefully) ratably expensed (amortized) over the useful life of the asset, with interest charges accrued on the balance not yet expensed.

A number of factors affect the costs seen by one service provider versus another, the most important being:

- Volume: many IT services experience decreasing cost to volume relationships, including software licenses, network connections, equipment purchases, and construction costs (up to a point). One key component of cost that generally is not affected by volume is hired personnel. Hiring three programmers generally costs three times as much as hiring one. This is only true for employees, however, as personnel acquired via consulting services do generally have a decreasing cost per unit for higher volume contracts.
- Location: land, buildings, construction, utilities (especially electricity) are all dependent on location. We all know the cost differences between our major urban universities and those located in less costly parts of the country. And we all know that cost per kilowatt can vary significantly based on location – it is a key reason that Google, Amazon, and Microsoft have all located mega-data centers on the Columbia River.

One of the big questions (mostly unanswered) is the shape of the cost/volume curve for various IT services. Is a 10,000 square foot data center half as expensive to operate (on a per square foot basis) as one of only 5,000 square feet? What about data centers of 20,000 or 40,000 square feet? What does the curve look like? Many important decisions about what and how much to share depend on the shape of cost curves that we don't know.

## PRICES

Volumes have been written on setting prices for products and services. The issues we discussed relate to the specific situation of universities (or other non-profits) providing shared services. Here the issues are somewhat limited. First, the service provider has to determine what costs (full, partial, incremental) and what percentage to recover via the pricing model (the "cost recovery goal"). Second, the provider has to decide how to charge users of the service (e.g. flat or per-unit price, discounts for longer term contracts, etc.) to best achieve that cost recovery goal. And finally, the service provider has to decide how much risk to take that either the pricing model will not recover the expected amount or that the expected costs will be different than planned, either way leading to over- or under-recovery of costs (in a for-profit enterprise, profit or loss).

If the service provider absolutely needs to recover full cost, then the only sure way of doing that is to set the price after all the costs have been expended and charge the users of the service retroactively. Some universities have done this internally, although most have abandoned it as a generally unsatisfactory approach. Under this method of setting prices, service users won't know how much the service will cost them until after it has been used (typically for a year) and the provider has determined all the costs and divided them among all the users. In this case, the provider is not willing to take any risk of underrecovery and has effectively shifted the "cost risk" to the service users (who takes the risk that their costs will be different than might have been originally estimated when deciding to use the service).

If exact cost recovery is not a requirement (which it generally is not), then service providers have the flexibility to estimate costs as well as expected user volumes at the beginning of the year (or other period), and set prices accordingly. Under this approach, the price charged will only fully achieve the

cost recovery goal if both estimates are 100% accurate, a highly unlikely situation. The service provider has to decide if risk of under-recovery is acceptable given the accuracy of cost and usage estimates.

Various other factors also enter into how prices are set in order to reduce risk. One such factor is the volume or time-period lock-in and discount. A provider may be willing to charge less per unit for the certainty of knowing that the service user will buy a certain volume of service or stay with the service for a period of time (e.g. one or more years). This is similar to the idea that buying in bulk generally costs less per unit volume of most commodities, or that a multi-year lease costs less per month than a simple month-to-month rental arrangement. From the service provider's point of view, both arrangements trade off per-unit price for certainty of selling the total volume, generally seen as reducing long-term risk.

Conversely, and somewhat uniquely in the university shared services environment, a buyer may be willing to pay more per unit for the certainty of knowing that the provider will continue to offer the service for a designated length of time. Finding the right price point will be a challenge with multiple buyers having different goals and needs.

Another factor in pricing is that the primary goal of a university service provider is providing service to local users. This may lead to a desire to charge different prices for members of local university community (faculty, staff, and students) than to users from other universities. [In our discussion, this was raised as an issue without recommendation.] Perhaps a better way of handling this situation (if it is indeed desired) is to not to have two prices, but instead to have subsidies for members of the faculty and staff that come from school or university funds. That way, everyone sees what the true cost and pricing models are, with local staff seeing that they are getting a clear subsidy from some other funding source. Note that there may be other issues regarding university policy or government accounting rules that have to be examined before deciding on any such subsidy arrangement.

Bartering was also raised as a possible way to pay for shared services. Bartering is just another form of determining value, although that form of currency may have desirable characteristics (it doesn't show up on anyone's budget or finance statement!) and advantages in some cases. Certainly, there are already cases of universities bartering services but this could be problematic on a larger scale.

## CONCLUSIONS

While the general consensus was that universities should always be talking about shared services, we are talking about them because of the heightened issues of the world financial situation. We believe that shared services reduce costs, whether directly or indirectly. All other reasons come back to cost reduction: having fewer services to run in each university reduces complexity and the need for management oversight, which in turn results in reduced cost. So two questions now appear:

- What level of cost savings is necessary to persuade a university to stop providing a particular IT service for itself and start buying it from someone else?

- What level of revenue (cost recovery) is necessary to persuade a university to start providing a service for others, or to persuade some group to set up an organization to provide that service?

If all services were identical wherever they were provided, and if all decision making were made on a totally rational economic basis, then the first question would be easy to answer: every university would buy services from someone else as soon as the on-going costs over some time period were less than your switching costs. Almost by definition, however, a shared service isn't the same as a dedicated service. You immediately have to make adjustments for the value of having control of your own service versus casting your lot in with others in a shared service environment.

The question for the service provider to answer is what level of revenue would make the effort (and responsibility) of being a provider of a shared service worth the hassle. "If I can't make at least \$x, it isn't worth my attention to deal with the issues versus all the other things I have on my plate." Some people will say that any revenue is better than no revenue. Others will set a fairly high threshold because their circumstances are different. And, in the case of setting up a new entity to provide the shared service, minimum commitments will be necessary in order to get the enterprise going.

There are several key issues that have to be addressed if we are to develop sustainable shared services:

- How to provide for capital expenditures, often ranging into multiple millions of dollars?
- How can we "sell" the idea of providing a shared service on our campuses when that is not necessarily seen as part of any university mission?
- How do we balance the need to recover costs with having prices that are seen as reasonable when universities account for costs in different ways?
- Are we better off creating a separate entity, e.g. a separate non-profit organization (perhaps an existing one) to provide shared services? Would this help attract funding?

To be successful over the long run with shared services, we have to tackle all these questions.

## ACKNOWLEDGEMENT

This paper is a result of the discussion begun at the Common Solutions Group meeting held in May 2009 at the University of Indiana, and continued via email. Primary participants were

- Jerry Grochow, MIT
- Michael Mundrane, University of California at Berkeley
- Greg Jackson, University of Chicago
- Donna Tatro, Princeton University
- Tammy Closs, University of Texas at Austin
- Andy Palms, University of Michigan
- Steve Cawley, University of Minnesota
- Marilyn McMillan, New York University
- Ron Kraemer, University of Wisconsin at Madison