

January 2005

What goes around comes around: Lessons learned from ICT outsourcing & the shift to the new utility.

Abstract

Today, as enterprises focus on business innovation and on the bottom line, the outsourcing of information and communications technology (ICT) services has become more strategic and selective while at the same time resulting in a broader mix of service providers and involving a myriad of relationships. This has led to many outsourcing contracts being scrutinised over 'value' versus real cost of service provision. This complicated ICT backdrop has given rise to additional debate over just how agile, responsive and flexible an organisation can be in a multi-sourced ICT environment, and whether service delivery from multiple suppliers under shorter contract terms is, in fact, of high value.

A new model of ICT outsourcing – utility computing – is evolving to deliver purportedly improved benefits to both client and vendor with a renewed focus on flexibility in contract terms and pricing. Utility computing arrangements work on the premise that computing capacity (in the form of infrastructure and applications) can be provided to customers to meet their peaks as required on a 'user pays' basis. Like traditional electricity provision, customers of utility computing benefit from savings derived from no capital outlays on infrastructure, they can tap into updated technology and businesses can expand or reduce their resources without long-term contractual tie-in or wasted ICT capacity. This enables them to upscale or downscale as their requirements fluctuate, leading to potential cost savings.

For economies of scale to be realised, however, the service provider must initially have the critical ICT mass of either a number of smaller customers or one or a few major customers with considerable processing load. The service provider benefits from improved utilisation rates of their equipment across the board, through 'load sharing' across all customers' fluctuating demand cycles, and better streamlining of their systems across standardised and consolidated ICT environments.

This paper discusses the evolution of ICT outsourcing in Queensland and explores the theoretical and practical applications (and challenges) that the utility computing outsourcing model poses for a traditional outsourcing provider like CITEC, the commercialised ICT arm of the Queensland Government.

The paper provides preliminary research findings on the utility computing concept gathered from focus group discussions with 'C'-level executives in the public and private sector, and presents the framework to CITEC's own utility computing model. Finally, the paper provides recommendations to early adopters on the advantages and challenges of transitioning to the new computing model, and identifies key issues that need to be addressed by the broader ICT community in order for the utility computing model to be matured.

Introduction

Businesses today are increasingly reliant on computers and emerging technology to assist them in providing products and services to their customers as well as performing day-to-day operations.


Enabled by the remarkable innovation in information technology, we are fast approaching a time when the world will always be connected, always open for business. It will be a quick-paced environment, marked by the drive for increasing efficiency and decreased cost. But it will also be intensely competitive and without boundaries.

This increased reliance on computers, however, does not come without all the challenges of using the technology.

The evolution of ICT outsourcing

In the mid to late 1990s organisations everywhere began to outsource virtually any function that could be performed more cheaply with demonstrated effectiveness outside of the firm.

And, with the ever-increasing complexity of technology as well as tight labour markets for IT personnel, business for the IT outsourcing company boomed. What we witnessed was the evolution of IT outsourcing, which began with full



outsourcing, where often a single service provider supplied and managed the customer's infrastructure on their behalf. The perceived loss of control' over their IT department however, saw a shift to the popular second generation model for Managed Hosting Services, where customers in most cases did not own their own infrastructure, and services were delivered by a third party. Today, as enterprises outsource more strategically and selectively, they are using a wider mix of external service providers, in addition to internal service providers.

The added complexities associated with managing the myriad of relationships involved in the service delivery has led to a complicated IT backdrop.

With an increased focus on cost cutting during the past few years, many outsourcing contracts have come under renewed scrutiny. Intense debate over value provided versus cost has been one of many major points of contention.

According to Gartner, many enterprises have not focused on creating the agile, responsive, flexible organisation necessary to effectively manage the multi-sourced environment, leaving them vulnerable to low levels of service value.

Moving to the next generation ICT outsourcing model: Lessons learned

The most important lesson learned from the past decade is that an outsourcing deal can only be successful if both client and vendor benefit. Effective negotiations and outcomes result not from the negotiating prowess of one party or the other, but rather from a fact-based and objective appraisal of individual and mutual goals of the parties involved. Another, and related, lesson learned is that a business can't successfully outsource what it doesn't understand or can't define.

What we are seeing is the emergence of a new model of ICT outsourcing that keeps relationships on track where the contract term is shortened, thereby enabling renegotiations sooner rather than later. Another is to build in more flexibility in terms and pricing.

The emerging second and third generation sourcing deals mean strategic sourcing is moving up the management agenda. The shift to selective outsourcing has seen all-embracing contracts being split into smaller contracts,

sometimes with multiple suppliers (multi-sourcing).

The benefit of selective outsourcing is that the IT department can see fixed-cost service provision from a specialist provider, and companies can benefit from the service providers' greater economies of scale. The downside however, includes multiple contracts and greater load to manage the relationships.

While there is no magic fix to arrive at a 'perfect' ICT outsourcing model, what we have learned is that in the outsourcing realm, there is no single perfect relationship model and no single choice for everyone and forever. It's an evolving business model that is likely to change, and continue to adapt as micro and macro environmental trends shape and influence the way in which we conduct business.

Nevertheless, the different outsourcing models that have been tried and tested, failed or proven, clearly demonstrate that an outsourcing strategy should always be based on an individual case.

New challenges for outsourcing

The current economic climate requires organisations to respond to the most pressing business challenges even quicker than before, while at the same time constantly evolving their ICT resource strategies. We are increasingly seeing the role of information technology shifting from its former status of serving as an "enabling" function -- where business strategies could be more quickly and effectively realised -- to that of a true "driver," where a company's strategic positioning is propelled by considerations concerning information technology.

Enterprises are also facing new imperatives to deliver value for every dollar invested. CEOs, CIOs (and increasingly, CFO's) and ICT organisations generally, are looking for "on demand" solutions that can minimise business risk, manage uncertainty, enhance agility, and control costs, while reducing the complexity of their IT environment. This shift is akin to what the manufacturing sector experienced in the 1990s with the emergence of Total Quality Management processes such as Just in Time (JIT) inventory management.

With these shifting goal posts of business impacting on IT delivery models yet again, vendors have been quick to act according to

evolve and adapt their products to keep up with the changing IT landscape and its requirements.

The rise of the utility computing concept

It is a deep economic truth that throughout history, people who focus on an activity and develop a capability for doing it best end up doing that activity for everyone else.

Most of us no longer grow our own food or make our own clothing, nor do businesses typically generate their own power. The ICT world is not immune to this economic reality either.

We know that few companies today gain competitive advantage from their applications software, with the return on investment for “custom” application development even more difficult to measure. Eventually, that will be provided in the lowest cost, most reliable manner. Hence, the rise of the next generation outsourcing model - utility computing - will also be driven by this economic reality.

Gartner defines the IT utility as an evolving framework for the remote delivery of standardised IT services or business processes to enterprises. This definition is primarily used to describe external service provider (ESP) offerings, but enterprises may build internal IT utility infrastructures. IDC research holds a similar view of utility computing, and define it as a next generation of ICT outsourcing.

So what is different about this model compared to its predecessors? Utility computing goes one step further than traditional outsourcing, requiring a shift in the way in which IT products and services are delivered and charged for.

In addition, it also solves many of the challenges and pitfalls that traditional or second generation outsourcing experience. The fundamental difference lies in standardization and consolidation which is not dissimilar to the classic utility offering of electricity, where users pay for what they use. In Australia, for example, everyone uses the same voltage offerings (eg. 240 volt and 415 volt). If someone wishes to use a different voltage then they are required to install their own generators at a capital cost plus pay ongoing costs to run them. It is unlikely that utility computing in the practical, on-demand sense would be much different.

Secondly, utility computing arrangements offer far more flexibility in terms of contract relationship and length, leaving businesses with greater freedom to expand or reduce resources without long-term contractual tie in or wasted IT capacity, enabling businesses to upscale or downscale as their requirements fluctuate.

Predicated uptake of utility computing

Analysts believe that the vendors have done a good job in selling the utility computing idea, but lament that the industry as a whole has not done a good job on educating the market on the readiness of the concept. It's also generally recognised – both by analysts and the business community – that there are theoretical economic benefits of the model, but it is seen as a distant, rather than imminent, trend.

The market appears to be largely uncomfortable in making the quantum leap to utility computing, and don't appear to fully understand the true benefits against an application service provision (ASP) or traditional outsourcing model. In 2004, IDC found that just 7% of Australian companies ($n = 263$) intend to move to a utility computing environment, and they predict that on the whole, on-demand, subscription-based IT functionality is at least five to seven years away.

In our own online focus group ($n = 9$, Australia-wide) into C-level executives' perceptions of utility computing we found that awareness was very low, and was dosed with healthy scepticism of the true cost savings that promise to be delivered with a utility computing model.

While the survey results are by no means exhaustive or indeed, prescriptive, to defining a utility computing solution, the selected sample did give us a taste of the general sentiment towards the level of understanding, the likely buying centre, and future acceptance of shifting to such a model. The benefit of conducting a focus group for us was to enable us to further define and better articulate, the practical pros and cons of the model.

The executives' key questions and concerns plaguing acceptance of the utility computing concept relate to:

- Associated costs – implementation, ongoing and future;
- Reliability of the systems and service providers; and
- Proven applications.

Our preliminary findings from the focus group revealed that the CFO would be a vital influencer in the analysis of the financial benefits for the concept. Along with the CIO, a proposal to the CEO and Board of Directors would be a key requirement to convince an organisation to adopt a utility computing model.

Importantly, many of those surveyed equated utility computing with the need to make the commitment to selectively outsource their ICT to access a utility model. For those who have not yet been exposed to an outsourcing environment, the perceived risks (eg. Loss of control) are no different to the barriers to entry to traditional selective outsourcing.

While noting that articulating the benefits of the model would greatly assist an organisation to understand the practical advantages of the model, executives expect that benefits should be clearly defined with supporting documentation and evidence of the savings and benefits to strengthen the effect. Key benefits of the model as highlighted by executives include:

- Flexibility – upscale and downscale;
- Pricing parallels to the flexibility;
- Capacity to keep up to date with IT developments; and
- IT effectiveness.

Implications for IT Customers and Service Providers

If utility computing comes to commercial fruition, it will effectively represent a new way of conducting business without owning the IT assets. Granted, it will provide the delivery of standardised processes, applications and infrastructure over the network, as a service, with both business and information technology functionality. Utility computing services are designed to run on a shared infrastructure, in which these resources can be dynamically shared among customers, bringing the financial advantage of economy of scale.

Conversely, there are three primary challenges to utility computing implementation: migration, integration and vendor management. The issues are that with migration comes complexity. If an IT environment is to be consolidated and standardised, how will be this achieved?

Whether it is middleware (or some kind of systems integration to bring together the disparate applications) remains to be seen. Moreover, the challenge of systems integration is further

exacerbated by the difficulty of bringing customers' infrastructure together to work seamlessly and compatibly. The utility computing services and the underlying physical IT infrastructure resources (including network routers, firewalls, server machines and storage devices for example) collectively form the managed utilities environment.

Lastly, vendor management will be a key performance criterion in making utility computing a success. Since no assets are owned by the customer, this leaves the way open for reduced contract duration, smaller switching costs and customer tie-in, unpredictable and fluctuating customer demand and reduced product lifecycles to keep up with technological advances.

The CITEC utility computing model

One of the biggest misnomers about the utility computing model is how the charging methodology will be applied to customer billing. The main questions surround the topics of transparency and cost allocation, particularly because of the many interdependencies between resources, which are often hidden and not well understood.

As an IT services provider and major outsourcer, CITEC has been developing and refining its utility computing charging methodology to enable customers to pay for what they use, for a variable workload. What CITEC has discovered in the process is that while hardware is available in the utility sense, the software vendors have been slow to adopt their utility charging methodology.

The underlying premise of the model is simply the provision of computing capacity to meet customer peaks as required and charged on a capacity basis. Obviously in the real world, there are going to be some limitations on the minimum amounts of capacity that can be utilised, as well as some financial limitations on the peak capacity.

For example, the minimum requirement might be 1 CPU, so the customer might be able to purchase 1 CPU per month as their minimum committed capacity. Their maximum peak capacity, however, may be in fact 4 CPUs. This means that the customer's billing in any one month will be somewhere between 1 and 4 CPU usage. Even though that customer might be on a 16 CPU box and have a 4 CPU share, we do have to provide some upper limits so that we can accommodate all customers with their respective variable loads. To take the capacity on demand

model further, instead of “reserving” 3 CPUs at a certain agreed rate, the customer may be able to access these CPUs at a lower cost by sharing that capacity with other users on that same resource.

We envisage that at the end of the day, billing will resemble something like a mobile phone bill where the user will not be restricted in what they use, but they will know the usage costs up front and be able to plan for their peak and off-peak load. It might even be similar to stand-by travel, whereby there will be a minimum IT capacity requirement, but there will be flexibility to jump on board ‘spare’ capacity on a first in best served basis to meet unplanned peak load fluctuations and/or spikes.

What goes around comes around

Our model is by no means a flash in the pan concept. In reality, pay-as-you-go IT services began decades ago with time-sharing and service bureau businesses. More recently, they have simply been reinvigorated by the universal connectivity of the Internet, by ubiquitous Web browsers that enable a simple, standard access method for people, and now by Web services that deliver a similar simple, standards-based access mechanisms to access these services.

Indeed, CITEC first established time-share and bureau service facilities more than 25 years ago when it operated the Queensland Treasury's IT arm as part of the Queensland Government. Since then, CITEC has provided on-demand and pay-as-you-go services including mainframe services, data storage and back-up, and transaction-level services including payroll and other based business process outsourcing services as well as its Confirm information brokerage offerings.

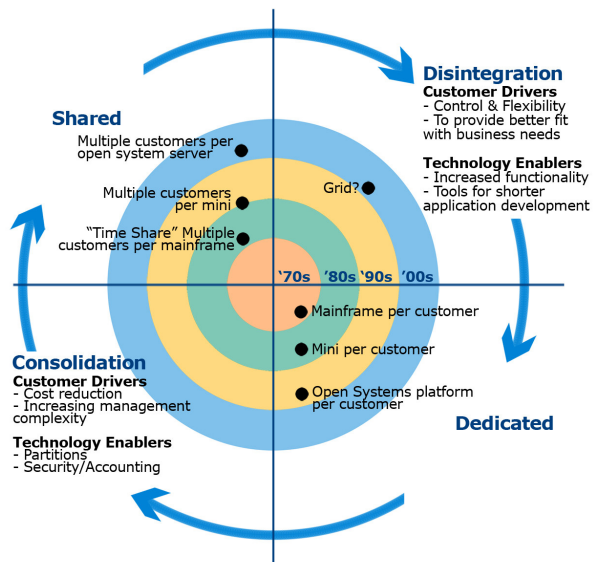


Figure 1 – The ICT lifecycle diagram

The tried and proven “pay as you go” concept was in widespread use in the late 1960s when computers were very large and very expensive systems. Because of this they necessarily had to run a mixture of workloads. Back then, IT was not cost effective for most individual agencies and departments to buy and operate themselves. For this reason, the Queensland Government invested in setting up a bureau service (operated by CITEC) based on a number of these mainframes. The type of processing that took place on those shared environments was typically called “time sharing” processing.

It was a form of utility computing insofar as the charge-back regimes were based on actual resources utilised by individual customers. They were charged on a ‘user pays’ basis for items such as CPU cycles consumed, pages printed etc. IT capital investment and operational expenditure was the responsibility of the bureau, and it was the responsibility of the bureau operator to recover those costs through charge-back regimes.

These operating systems were specifically designed for this time sharing mode of operation. Individual users had their own accounts on the machine, and the operating system tracked the various usage elements for each account.

For the mainframes that are still in use today, this is precisely what still happens, representing a form of utility computing since it is usage based charging and on demand computing.

With the rise of “mini computers” in the 1980s (which were generally dedicated to individual applications), organisations and/or departments could now afford to purchase their own computer systems. This often meant that they became responsible for providing capacity, all of the investment, capital and operational upkeep, technical support and forward capacity planning generally. This also corresponded with the shift towards second generation outsourcing where agencies chose to acquire and manage the capacity themselves, as well as carry the cost of the spare capacity.

Concurrently, in the 80s and 90s in Queensland at least, we saw the market move away from proprietary mainframe centralised computing to open system small to midrange platforms which were far more affordable. With open systems comes choice, and unless you keep a tight rein on systems architecture and technology, the result can be a ‘grab bag’ of technology that leads to expensive support, along with troublesome interoperability, cost and rework and ultimately, unreliability.

Now, companies want more coherent technology, better economies of scale and scope, and no lock in – which leads back to the model of cost reduction through standardisation, or at least a controlled number of key platforms.

These trends indicate the inevitable shift to consolidation and back to a bureau-type, on demand computing environment.

True utility is not leasing

A key issue in moving forward, of course, is making sure vendors, who are critical partners in the relationship, get behind the utility computing model.

In a true utility model, the vendor shares the risk with the customer and assumes responsibility for providing capacity, technology upgrades as required, retains ownership of the asset and the customer simply pays for usage on their facility.

Utility in action

What CITEC currently does with its mainframes does in some part resemble utility computing in practice. And, while many vendors have a utility offering for hardware infrastructure and can point to success stories in both Australia and overseas,

accurate, independent results for the actual adoption and success in Australia is difficult to obtain in this introductory stage - particularly as many of these Australian implementations have not yet been in production for a full 12 months.

Over time, we foresee larger customers realising the benefits of utility without necessarily outsourcing their processing load. For the smaller customers who do not have the critical mass in their own right, access to utility computing benefits really means making the decision to outsource their hardware in some form or fashion.

As a long-established ICT outsourcer, CITEC has been exploring how we can take these smaller customers on board, and what products and services might be attractive to them.

To fully understand where a service provider such as CITEC sits with its practical utility offering, we need to look at the broad spectrum utility computing concept.

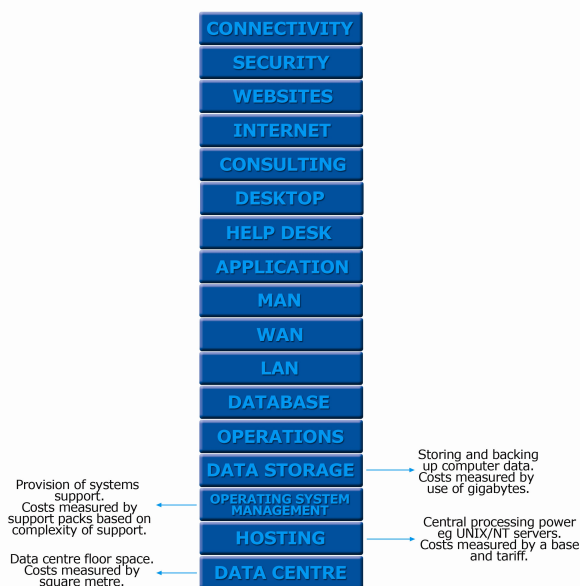
At one end of the spectrum there is what we describe as a ‘single component level’ utility model - which could be any IT component. Solution complexity is highest at this end as it consists of many single components from multiple sources. Conversely, service provider management of the total solution presents the least complexity. This is because at this most basic single component level the charging regime isn’t too technically different to how that component has been provisioned.

At the other end, we have the business level transaction utility, or ASP which is a total business transaction utility where a customer can buy payroll at a price per person, or a per purchasing transaction for a finance solution. We see this as the pure “plug in the wall” concept with the metric at the business level transaction, not the technical level.

Solution complexity is lowest at this end; however, service provider management of the whole solution is of a highly technical and complex nature because of the cross-metrics required to calculate customer billings.

What CITEC has devised for one customer that has entered into a contract for utility computing services lies somewhere in the middle. We have a ‘four component level’ utility, or ‘solutions stack’ comprised of data centre, hosting, operating systems management and data storage, which is

around half way through the journey of functionality.



The IT Solutions Stack represents all the possible levels of delivery involved in an integrated IT solution

Figure 3 – CITEC ICT solution stack

At this level our metrics are still technical but there is a more complex relationship between the components. For example, if a customer was to purchase more storage and it could fit into the existing footprint of the configuration that is being applied to their customised solution, the other elements might not necessarily change.

On the other hand, if a customer was to purchase more storage and it couldn't fit into the existing configuration, it would mean that their data service and other utility measures would be affected. As a customer, these complex component interrelationships need to be taken into account. There is no doubt then, that a certain level of maturity is required around the buying of these services.

If we were to revisit the single component level again, this is comparatively far more complex because there are multiple models required to resource the IT solution. For example, there might be one charging system for the CPU, but the customer still owns their own disk but it is being maintained by a third party (which requires management of resources).

We believe an ICT outsourcer is best positioned to be able to deliver 'more than just the boxes' in the true utility computing model. CITEC, for

example, is finding a middle ground where some of the complexity is taken out of the solution and we will provide the building blocks via utility and then work our way up the solution stack to provide other components.

At our current 'four component level' there is more transparency, and there is ability to understand the interrelationship between components and in other cases, lack thereof, and the customer can clearly see where their usage has increased or decreased.

Barriers to entry

For customers with flexible peak capacity requirements (eg. batch processing), they may be able to share their processing load with other customers. The challenge will be configuring, metering, and monitoring this dynamic allocation of resources to ensure that customers' business needs are met in a cost effective manner. It will take a very experienced service provider or large internal providers with bureau or time-sharing experience to achieve this.

Not all firms will benefit equally from utility computing but to begin with, early adopters are likely to span the full spectrum of customers. Those businesses that can benefit from consolidation or have the critical mass are primary candidates for utility computing. The Queensland Government, with its shared services model, would certainly benefit from this model.

Other organisations that consist of multiple business units (such as multinationals with affiliates or subsidiaries) who can see benefit in consolidation will move to trial this concept. Other organisations whose usage levels reflect directly with "seasonal" IT peaks and troughs such as companies in the financial services, media or manufacturing industries that require that flexibility without carrying spare capacity will also benefit.

In any case, in order to achieve the cost saving promise of utility computing, there does appear to be consensus that a high level of standardisation is required to reach the economies of scale. Large customers may actually have enough critical mass in their own right to achieve this and they may decide to consolidate their applications onto standardised platforms.

Where CITEC can add value is by bringing a number of small users who agree to standardise

together and consolidate them using a utility model.

Moving forward

While at the most basic level CITEC has developed a 'solutions stack' of technology offerings, we acknowledge that utility computing requires an even more solid foundation to realise the benefits of agile and efficient computing.

Utility computing is really about finding a balance between transactional computing (fee for transaction) and arriving at business transformation as it relates to vital technology needs. It involves using technology as a business strategy to remain responsive to the market, and to improve operational efficiency to gain competitive advantage.

So how can business afford to move forward and stay competitive if the required infrastructure for growth is unobtainable? If what we are seeing is technology at its very early stages of the product life cycle, then utility computing (or the organisation's integrated IT infrastructure) will become the "technology driver" to underpin corporate strategy and ultimately, enable a successful business transformation with operational gains to boot.

Security, scalability, ROI and future functionality including reliability and availability, will be the ultimate test of utility computing. The outsourcer that has the business skill and technological nous to successfully transition and standardise an organisation's technology and processes (though still including some element of customisation to meet unique client requirements) on a utility computing platform serves to position both themselves, and their customers, on a competitive pedestal if it means they are better able to adapt and evolve their business strategies more effectively.

Critical success factors

Across the board we believe customers of utility computing will demand:

1. clearly defined timetables for implementation;
2. clearly defined pricing;
3. system functionality that lives up to its billings;
4. short-term ROI;

While these are not new concepts for IT solutions providers, relatively few software and services firms have consistently delivered on all four of these objectives. In many circumstances, software and services providers that have deep vertical industry knowledge will be best positioned to compete on a going forward basis.

Those firms that have the ability to understand the business processes, business issues, and most frequently utilised legacy systems within a given industry are better positioned than a generalist to present all of the aspects of an IT utility computing solution that will generate ROI for a prospective customer. Further, knowledge of existing systems will give solutions providers a leg-up in understanding how to leverage existing investments through network integration. Knowledge of existing systems also gives them deep knowledge of the gaps that exist in traditional solution and where the opportunities lie for future improvements.


Relatively few enterprises have truly mastered strategic sourcing and those that have are expected to enjoy a competitive advantage because of it. Strategic sourcing, in a utility computing sense, will not be about technology, but about leadership and superior management capabilities (market analysis, relationship and project management through to contract negotiation etc).

These enterprises already understand the importance of strategic sourcing and the difficulty of internal and external management to achieve business results. They are already organised in semi-independent or independent business units that often rely on external partnerships to differentiate and innovate in their market.

The importance of a sourcing strategy needs to be stressed at the board level and should be leveraged to enhance the business/IT relationship and the IS organisation's participation in business decisions.

Summary

From an industry point of view, agreement on technological standardisation among the vendor community will be extremely difficult to achieve. Differentiation in the market is largely driven by proprietary innovation and we can expect to see significant investment among both software and hardware vendors as they race to develop their own 'plug and play' versions of utility computing convenience products.



Whatever products, services or metrics system is decided upon, we believe that strategic change management will be a critical competency required by service providers in helping facilitate and transition customers onto a true utility computing environment.

An organisation such as CITEC, with its 40 year history as an outsourcing expert, is in a unique position to become an expert broker of IT utility computing solutions for business innovation.

As the market becomes more ready, CITEC's experience with early adopters of utility computing and our proven track record in managing multi-vendor environments, consulting expertise and technical prowess, means we will be seen as a master in strategic sourcing with a depth of understanding in sourcing options and models, including the ability to attract and retain the best partners. As a strategic sourcer (or "multisourcer"), we will ensure that all the utility computing service providers work together (where collective services are required) to achieve business goals.

Content contributors (CITEC): Mike New, Carlos Cienfuegos, Darryl Patterson, Bob Gurnett.

Written and edited by Nyree McKenzie (CITEC). Copyright CITEC February 2005.

About CITEC

CITEC is a national information and communications technology (ICT) service provider with offices in Brisbane, Sydney, Canberra and Melbourne. As Queensland's largest ICT provider CITEC's clients include local, national and multinational organisations. Founded in 1964, CITEC has today established itself as the second largest Australian-owned information systems outsourcer and is ranked eighth in IDC's "Australian IT Outsourcing Market, Top 10 Players - 2003" (Source - IDC, 2004).

CITEC manages several high security data centres and has approximately 600 employees located across Australia. CITEC owns, through its investment company CSI Holdings Pty Ltd, a majority interest in human resources and payroll software developer - Aurion Corporation. CITEC's core solutions encompass the following areas, information brokerage, business process outsourcing and applications management, infrastructure management and ICT professional services.