



Submission to the Australian Government Information Management Office

In response to the

CLOUD COMPUTING STRATEGIC DIRECTION PAPER

Opportunities and applicability for use by the Australian Government

January 2011

## Contents

Definitions.....	4
Establishing Cost Value & Risk analysis.....	5
Cloud economics – an analysis.....	5
Historical context .....	6
Economic drivers for cloud .....	6
Supply-Side Economies of Scale.....	7
Demand-Side Economies of Scale.....	8
Multi-tenancy Economies of Scale.....	13
Overall Impact.....	14
Implications.....	14
Public Sector Cloud Benefits .....	14
Private Cloud.....	15
The Cost of Private Clouds .....	17
Finding Balance Today: Weighing the Benefits of Private Cloud against the Costs .....	18
The Long View: Cloud Transition Over Time .....	19
Confronting and Managing Risk in the Public Sector Cloud scenarios .....	20
Current practice .....	20
Taking a Principles & Risk approach .....	20
Example of APRA.....	21
Role of Interoperability & Standards .....	21
Interoperability in the Cloud.....	22
The Promise of Cloud Interoperability.....	22
The Hard Work of Interoperability.....	22
Openness, Industry Collaboration to Pave the Way.....	23
AGIMO whole of government role in JTC-1 Cloud standards work.....	24
Additional cloud examples in Australia.....	24
MYOB .....	24
Brisbane City Council .....	24
Additional comments.....	26
Additions and omissions .....	26
Industry involvement.....	26
Appendix A .....	28
Public Sector related (Non Australia).....	28
Commercial (Australia) .....	29
Commercial (Non-Australia) .....	30

## Introduction

Microsoft welcomes the release of the discussion paper issued by the Australian Federal Government and the opportunity to provide feedback on this important area. We also note the alignment this Cloud Computing Strategic Direction Paper has with the government's strategic reform of the Public Sector outlined in [Ahead of the Game: Blueprint for the Reform of Australian Government Administration](#). Microsoft believes five of the nine key recommendations in the blueprint can be directly supported by a carefully developed and implemented Cloud Strategy:

- 1. Delivering better services for citizens*
- 2. Creating more open government*
- 7. Strengthening the workforce*
- 8. Ensuring agency agility, capability and effectiveness*
- 9. Improving agency efficiency*

Microsoft agrees with AGIMO that cloud computing primarily represents a new form of delivering and consuming ICT related services rather than a new technology per se and that it does hold promise for improving the efficiency and effectiveness for delivering services to citizens. In particular, the significant changes in citizen expectations of their public sector organisations, uncovered through the course of the Gov2.0 Taskforce, will see agencies struggle to deliver with their present IT systems. We offer some local Australian examples later in this submission that demonstrate how cloud can provide an alternative means of supporting consumer/citizen demands in the 21<sup>st</sup> century.

Microsoft also agrees that although the technological and business application options for the use of cloud services is expanding rapidly, a number of important aspects required for the appropriate selection, deployment and use of any cloud service remains at an early stage of development or understanding. This is particularly true of the legal and contractual arrangements supporting the operation of cloud services and the understanding of the economic drivers and trade-offs between traditionally operated IT and cloud-delivered variants. We discuss these later in the document.

Lastly, Microsoft believes that cloud computing represents an opportunity for organisations generally to reassess the way in which not just security but overall risk is managed to ensure that the Australian Government is able to continue to protect its people, information and assets, at home and overseas. As the most recent Protective Security Policy Framework<sup>1</sup> states: "How the Government protects its people, information and assets is critical to effective engagement with the Australian people". Microsoft therefore advocates a move away from purely prescriptive advice to agencies in favour of a more Principles and Risk based approach which provides the necessary resilience, flexibility and comprehensiveness given the increasingly distributed nature of information technology and the ever-changing threat landscape within which it must operate.

Microsoft Corporation has been a world leader in the development and operation of services that are cloud or cloud-like and has been operating global scale Internet services since 2000. The following outlines some key data about Microsoft Global Cloud experience as context for our submission:

---

<sup>1</sup> Protective Security Policy Framework V1.1 – September 2010

## Operational capacity

Microsoft operates multiple data centres across the globe with our most recent designs spanning more than 65,000 m<sup>2</sup> with critical power of 30 MW

- Hotmail – Microsoft has owned and operated the Windows Live Hotmail, formerly known as Hotmail, since the acquisition of Hotmail in 1997. There are over 350 million active Hotmail accounts today around the world.
- Tellme – Microsoft purchased Tellme Networks in 2007, and has been hosting the Tellme Service for the past 4 years. Tellme handles over 2 Billion unique calls each year.
- Xbox Live – Xbox Live has been available for 9 years, and currently has 23 million active members in over 35 countries around the world.

## Microsoft Consumer Services

- 2 billion Bing Search queries per month
- 460 million Windows Live IDs in current use
- 300 million Messenger accounts in current use
- 5.7 million Sky Drive accounts in current use
- 1.5 million Security Essentials (OneCare) accounts
- 550 million visitors to Microsoft sites per month across 42 markets in 21 languages

## Microsoft Online Services (for businesses)

- 15,000 LiveMeeting customers for web conferencing
- 36 million LiveMeeting attendees per year
- 10 million Exchange Hosted Services seats for email filtering
- 4.5 billion messages filter per day by Exchange Hosted Services

## Public Sector Cloud Services

- 229,000 active Microsoft Online Productivity users today across 851 customers
- 593,000 users of Exchange Hosted Services for email filtering
- Over 4M Outlook Live and Live@edu users in Education market
- 81,600 LiveMeeting users for web conferencing in Government market

## Definitions

As a result of significant experience gained in building and operating globally distributed cloud computing operations, Microsoft has gained a deep insight into the relative dynamics of “Public” and “Private” clouds. We have observed a trend globally to mis-use the terms, particularly when discussing issues of data sovereignty where “Private clouds” become a proxy for reduced sovereign risk. We feel that this lack of precision hinders a fuller understanding of the true deployment and operational models for these options thereby reducing the accuracy of any business case assessment that an agency might make.

Specifically, Private clouds are often cited when what is really meant is a cloud that operates applications and data *within* a specific country or geography.

Burton Group provides a good description of Public vs. Private cloud<sup>2</sup> which is worth highlighting:

---

<sup>2</sup> Cloud Computing: transforming IT-v1, April 20, 2009, Burton Group

**Public cloud:** An IT capability as a service that cloud providers offer to any consumer over the public Internet. Examples: Salesforce.com, Google App Engine, Microsoft Azure, and Amazon EC2.

**Private cloud:** An IT capability as a service that cloud providers offer to a select group of consumers. The cloud service provider may be an internal IT organisation (i.e., the same organisation as the consumer) or a third party. The network used to offer the service may be the public Internet or a private network, but service access is restricted to authorised consumers. Example: Hospitals or universities that band together to purchase infrastructure and build cloud services for their private consumption.

The two key take outs from this then are:

**Tenancy:** Public clouds by definition operate on a multi-tenancy basis whereas Private clouds service only a subset of available users and may have only one tenant. This has significant implications for their ultimate potential for efficiency, resiliency and availability.

**Data location:** Both Public and Private clouds may operate inside or outside a national boundary. This independence between the physical location of its resources and the allowable user base is very important and has significant implications for user experience and the economic models possible as we shall expand on later in this submission.

## **Establishing Cost Value & Risk analysis**

Cloud computing represents an opportunity to extend the enterprise beyond the existing data centre walls and is part of a long term trend towards a more distributed computing model. Although we are at an early stage of this development, it is already fair to say that by externalising parts of the IT function, cloud computing is already supporting an improved level of efficiency for users and enabling a greater level of support for interactions between the organisation and its citizens, and suppliers.

However, determining if, when and how an organisation should embark on this journey requires a rigorous, fact-based assessment that answers three vital questions: What will this cost me? What will I get out of it? What are the risks? After all, cloud computing offers increased choice for internalising or externalising resources right across the IT operations and therefore a holistic approach is critical to support strategic and tactical decision making.

Although the *Cloud Computing Strategic Direction Paper* wisely identifies certain risks, calls out the requirement for vendor selection criteria and a governance process, Microsoft believes that in order to fully equip agencies to develop sound cloud computing plans and execute on them, a rigorous and complete decision framework will be required. Our own experience in developing and operating global scale cloud services and consuming them within our own organisation over a number of years suggests that the dimensions of Cost, Value and Risk together provide a good starting point. These three inter-related aspects enable a suitable analysis of what is possible, its economic logic to the organisation and any changes in the risk profile in so doing.

## **Cloud economics – an analysis**

To date, much of the economic discussion about cloud has revolved around the potential savings from data centre consolidation and improvements in hardware utilisation through virtualisation. While these have certainly delivered significant savings, these aspects represent only a small subset of the drivers for cloud economics and Microsoft believes that a deeper understanding of the underlying economics driving cloud services will enable government to make better decisions.

In the following section, we assess the economics of the cloud by using in-depth modelling. We then use this framework to better describe the long-term IT landscape, with the goal of providing useful context for public sector IT leaders as they make decisions about how to take advantage of cloud.

### **Historical context**

Historical analogies are often a useful tool in clarifying present and future changes in market dynamics – particularly when technology is an important driver and for this reason we think there are useful similarities between the development of the auto industry and what is currently unfolding across the IT industry.

The first cars that emerged in the early 20<sup>th</sup> century were initially called “horseless carriages”. Understandably, the new invention was initially viewed through the lens of the paradigm that been dominant for centuries; the horse and carriage. The first cars also looked very similar to the horse and carriage (just without the horse), as engineers initially failed to understand the new possibilities of the new paradigm, such as building for higher speeds, or greater safety. In fact, engineers kept designing the whip holder into the early models before deciding that it really wasn’t necessary anymore.

Initially there was a broad failure to fully comprehend the new paradigm. Banks claimed that *“The horse is here to stay but the automobile is only a novelty, a fad”*. Even the early pioneers of the car were unable to envision the impact their work could have on the world. In an attempt to estimate the long term market opportunity, the Daimler Company concluded there could never be more than 1 million cars, mainly due to their high cost and the lack of capable chauffeurs. By the 1920s the number of cars had already reached 8 million and today there are over 600 million cars. What the early pioneers failed to realise were the profound reductions in both cost structure and complexity of operating cars, eliminating prior constraints and bringing cars within everyone’s reach.

Today, IT is going through its own seismic shift: the shift from the client/server datacenters of old to the cloud. Cloud promises not just “faster, cheaper” IT, but also “easier”, “more flexible”, “more effective” IT.

Just as in the early days of the car industry, it’s currently difficult for to see where this new paradigm will take us. The goal of this analysis is to help build a framework that supports governments to plan effectively for a cloud transition. We take a long-term view in our analysis, as this is a prerequisite when evaluating decisions and investments that could last for decades. As a result, we focus on the economics of cloud rather than on specific technologies, as economics often provide a clearer understanding of transformations of this nature.

In the section below, we outline the underlying economics of cloud, focusing on what makes it truly different from client/server. We then assess the implications of these economics for the future of government IT. We look discuss both the current situation as well as the long term changes.

### **Economic drivers for cloud**


Economics are a powerful force in shaping industry transformations. Today’s discussions on the cloud focus a great deal on technical complexities and adoption hurdles. While we acknowledge that such concerns exist and are important, historically, underlying economics have a much stronger impact on the direction and speed of disruptions, as technological challenges are resolved or overcome through rapid innovation. During the mainframe era, client/server was initially viewed as a “toy” technology not viable as a mainframe replacement. Yet, over time the client/server technology found its way into the enterprise (Figure 1).

Similarly, when virtualisation technology was first proposed, compatibility concerns and potential vendor lock-in were cited as barriers to adoption. Yet underlying economics of 20 to 25 percent savings<sup>3</sup> compelled CIOs to overcome these concerns, and adoption quickly accelerated.

The emergence of cloud services is again fundamentally shifting the economics of IT. Cloud standardises and centralises IT resources, bringing the core IT infrastructure into large data centers that take advantage of significant economies of scale in three areas:

- **Supply side savings.** Large-scale data centers (DCs) lower costs per server.
- **Demand aggregation.** Aggregating demand for computing smooths overall variability, allowing server utilisation rates to increase.
- **Multi-tenancy.** When changing to a multitenant application model, increasing the number of tenants (i.e., customers or users) lowers the application management and server cost per tenant.

FIGURE 1: Beginning the Transition to client/server




Source: "How convention shapes our market" longitudinal survey, Shana Greenstein, 1997

### Supply-Side Economies of Scale

Cloud computing combines the best economic properties of mainframe and client/server computing. The mainframe era was characterised by significant economies of scale due to high up-front costs of mainframes and the need to hire sophisticated personnel to manage the systems. As computing power – measured in MIPS (million instructions per second) – increased, cost declined rapidly at first (Figure 2), but only large central IT organisations had the resources and the aggregate demand to justify the investment. Due to the high cost, resource utilisation was prioritised over end-user agility. Users' requests were put in a queue and processed only once needed resources were available.

FIGURE 2: economies of scale (illustrative)



Source: Microsoft

With the advent of minicomputers and later client/server technology, the minimum unit of purchase was greatly reduced, and the resources became easier to operate and maintain. This modularisation significantly lowered the entry barriers to providing IT services and switched the trade-off from efficient resource utilisation to agility. The result is the current state of affairs: datacenters sprawling with servers purchased for whatever specific purpose end-user needed at the time, but running at just 10% utilisation.

<sup>3</sup> Source: IDC Datacentre Trends Survey, 2007. Savings calculated on total cost of ownership for servers.

Cloud computing is not a return to the mainframe era as is sometimes suggested, but in fact offers users economies of scale and efficiency that exceed those of a mainframe, coupled with modularity and agility beyond what client/server technology offered, thus eliminating the trade-off.

The economies of scale emanate from the following areas:

- **Cost of power.** Electricity cost is rapidly rising to become the largest element of total cost of ownership (TCO),<sup>4</sup> currently representing 15 to 20 percent. Power Usage Effectiveness (PUE) tends to be significantly higher in large facilities than in smaller ones. While the operators of small data centers must pay the prevailing local rate for electricity, large providers can pay less than one-fourth of the national average rate by locating its DCs in locations with inexpensive electricity supply and through bulk purchase agreements.<sup>5</sup> In addition, research has shown that operators of multiple data centers are able to take advantage of geographical variability in electricity rates, which can further reduce energy cost.
- **Infrastructure labour costs.** While the cloud significantly lowers labour costs at any scale point by automating many management tasks, larger facilities are able to lower them further than smaller ones. While a single system administrator can service approximately 140 servers in a traditional enterprise,<sup>6</sup> in a cloud data centre the same administrator can service thousands of servers.
- **Security and reliability.** While often cited as a hurdle to public cloud adoption, increased need for security and reliability leads to economies of scale due to the largely fixed level of investment required to achieve operational security and reliability. Also, large data centres are often better able to bring deep expertise to bear on this problem than a typical corporate IT department.
- **Buying power.** Operators of large data centres can get discounts on hardware purchases of up to 30 percent over smaller buyers.

In the future, there will likely be many additional economies of scale that we cannot yet foresee. The industry is at the early stages of building data centres at a scale we've never seen before. The massive aggregate cost of these mega data centres will bring intense scrutiny and R&D to bear on running them more efficiently, and make them subject to a process of continuous improvement. Providers of large-scale data centres, for which running them is a primary business goal, are likely to benefit more from this than smaller data centre providers that run inside enterprises or public sector agencies.

### **Demand-Side Economies of Scale**

The overall cost of IT is determined not just by the cost of capacity, but also by the degree to which the capacity is efficiently utilised. We need to assess the impact that demand aggregation will have on costs of actually utilised resources (CPU, storage, and I/O).<sup>7</sup>

In the non-virtualised data centre, each application runs on its own physical server.<sup>8</sup> This means the number of servers scales linearly with the number of server workloads. In this model, utilisation of

---

<sup>4</sup> Not including app labour. Gartner suggests that three-year spending on power and cooling already outstrips three-year hardware spending.

<sup>5</sup> Source: U.S. Energy Information Administration (July 2010) and Microsoft. While the average U.S. commercial rate is 10.15 cents per kilowatt hour, Microsoft pays 2.2 cents in Quincy, Washington.

<sup>6</sup> Source: James Hamilton, Microsoft Research, 2006.

<sup>7</sup> Here we talk generally about "resource" utilisation. We acknowledge there are important differences among resources. For example, because storage has fewer usage spikes compared with CPU and I/O resources, the impact of some of what we discuss here will affect storage to a smaller degree.

<sup>8</sup> Multiple applications can run on a single server, of course, but this is not common practice. It is very challenging to move a running application from one server to another without also moving the operating system, so running multiple applications on one operating system instance can create bottlenecks that are


servers has traditionally been extremely low, around 5 to 10 percent.<sup>9</sup> Virtualisation enables multiple applications to run on a single physical server within their optimised operating system instance, so the obvious benefit of virtualisation is that fewer servers are needed to carry the same number of workloads. But how will this affect economies of scale? If all workloads had constant utilisation, this would entail a simple unit compression without impacting economies of scale. In reality, however, workloads are highly variable, often demanding large amounts of resources one minute and virtually none the next. This opens up opportunities for utilisation diversification.

We analysed the different sources of utilisation variability and then looked at the ability of the cloud to diversify it away and thus reduce costs.

We distinguish five sources of variability and assess how they might be reduced:


1. **Randomness.** End-user access patterns contain a certain degree of randomness—for example, people checking their email at different times (Figure 3). To meet service level agreements, capacity buffers have to be built in to account for a certain probability that multiple people will do particular tasks at the same time. If servers are pooled, overall variability can be reduced.
2. **Time-of-day patterns.** There are daily recurring cycles in people’s behaviour: consumer services tend to peak in the evening while workplace services tend to peak during the workday. Capacity has to be built to account for these daily peaks but will go unused during other parts of the day, causing low utilisation. This variability can be countered by running the same workload for multiple time zones on the same servers (Figure 4) or by running workloads with complementary time-of-day patterns (for example, consumer services and enterprise services) on the same servers.
3. **Industry-specific variability.** Some variability is driven by industry

FIGURE 3: Random variability (exchange server)




Source: Microsoft

FIGURE 4: time-of-day patterns FOR SEARCH



Source: Bing Search volume over 24-hour period

FIGURE 5: industry-specific variability




Source: A

difficult to remedy while maintaining service, thereby limiting agility. Virtualisation allows the application plus operating system to be moved at will.

<sup>9</sup> Source: The Economics of Virtualisation: Moving Toward an Application-Based Cost Model, IDC, November 2009.

dynamics. Retail firms see a spike during the holiday shopping season while U.S. tax firms will see a peak before April 15. There are multiple kinds of variability—some recurring and predictable (such as tax season and the Olympic Games), and others unpredictable (such as major news stories). The common result is that capacity has to be built for the expected peak (plus a buffer), most of which will sit idle the rest of the time. Strong diversification benefits exist for industry variability.

FIGURE 6: Multi-resource variability (illustrative)




Source: Microsoft

4. **Multi-resource variability.**

Compute, storage, and I/O resources are generally bought in bundles: a server contains a certain amount of CPU, storage, and I/O. Some workloads like search use a lot of CPU but relatively little storage or I/O (“CPU bound”), while other workloads like email tend to use a lot of storage but little CPU (“storage bound”—see Figure 6). While it’s possible to adjust capacity by buying servers optimised for CPU or storage, this addresses the issue only to a limited degree because it will reduce flexibility and may not be economic from a capacity perspective. This variability will lead to resources going unutilised unless workload diversification is employed by running workloads with complementary resource profiles.

5. **Uncertain growth patterns.** The difficulty of predicting future need for computing resources and the long lead-time for bringing capacity online is another source of low utilisation (Figure 7). For start-ups in the private sector, this is sometimes referred to as the “TechCrunch effect.” Public sector organisations must get approval for IT investments well in advance of actually knowing their demand for that infrastructure, resulting in a constant push and pull between legislators/administrators, and IT pros within government agencies. By diversifying among workloads of multiple customers, a public cloud data centre can reduce this variability, as higher-than-anticipated demand for some workloads is cancelled out by lower-than-anticipated demand for others.

FIGURE 7: Uncertain growth patterns



Source: Microsoft

A key economic advantage of the cloud is its ability to address variability in resource utilisation brought on by these factors.<sup>10</sup> By pooling resources, variability is diversified away, evening out utilisation patterns. The larger the pool of resources, the smoother the aggregate demand profile, the higher the overall utilisation rate, and the cheaper and more efficiently the IT organisation can meet its end-user demands.


Using Monte Carlo simulation, we modelled the impact of **random variability** of demand on server utilisation rates as we increase the number of servers.<sup>11</sup> Figure 7 indicates that a pool of 1,000 servers can be run at approximately 90% utilisation without

violating its SLA – assuming no other form of variability exists. Note that higher levels of uptime (SLAs) become much easier to deliver as scale increases.

Clouds will be able to reduce **time-of-day variability** to the extent that they are diversified amongst geographies and workload types. Within an average organisation, peak IT usage can be twice as high as the daily average. Even in large, multi-geography organisations, the majority of employees and users will live in similar time zones, bringing their daily cycles close to synchrony. Also, most organisations do not tend to have workload patterns that offset one another: for example, the email, network and transaction processing activity that takes place during business hours is not replaced by an equally active stream of work in the middle of the night. Pooling organisations and workloads of different types allows these peaks and troughs to be offset.

**Industrial variability** results in highly correlated peaks and troughs throughout each firm (that is, most of the systems in a retail firm will be at peak capacity around the holiday season (e.g., web servers, transaction processing, payment processing, databases). Ideally, we would use the server utilisation history of a large number of customers to gain more insight into such patterns. However, this data is difficult to get and often of poor quality. We therefore used web traffic as a proxy for the industry variability. Figure 9 shows industry variability for a number of different industries, with peaks ranging from 1.5x to 10x average usage.

FIGURE 8: Diversifying random variability



Source: Microsoft

<sup>10</sup> Cloud computing is also distinctly different from the traditional server model in its ability to lower labour costs (automated management), lower hardware costs (homogenous and lower-end hardware), and provide scale-out (a more efficient application model).

<sup>11</sup> To calculate economies of scale arising from diversifying random variability, we created a Monte Carlo model to simulate data centers of various sizes serving many random workloads. For each simulated DC, workloads (which are made to resemble hypothetical web usage patterns) were successively added until the expected availability of server resources dropped below a given uptime of 99.9 percent or 99.99 percent. The maximum number of workloads determines the maximum utilisation rate at which the DC's servers can operate without compromising performance.

Microsoft services such as Windows Live Hotmail and Bing take advantage of **multi-resource diversification** by layering different subservices to optimise workloads with different resource profiles (such as CPU bound or storage bound). It is difficult to quantify these benefits, so we have not included multi-resource diversification in our model.

Some **uncertain growth pattern variability** can be reduced by hardware standardisation and just-in-time procurement, although likely not completely. Based on our modelling, the impact of growth uncertainty for small and midsize enterprises (up to 1,000 servers) is 30 to 40 percent overprovisioning of servers relative to a public cloud service. For smaller companies (for example, Internet start-ups), the impact is far greater.

So far we have made the implicit assumption that the degree of variability will stay the same as we move to the cloud. In reality, it is likely that the variability will significantly increase, which will further increase economies of scale. There are two reasons why this may happen:

- **Higher expectation of performance.** Today, users have become accustomed to resource constraints and have learned to live with them. For example, users will schedule complex calculations to run overnight, avoid multiple model iterations, or decide to forgo time-consuming and costly supply chain optimisations. The business model of cloud allows a user to pay the same for 1 machine running for 1,000 hours as he would for 1,000 machines running for 1 hour. Today, the user would likely wait 1,000 hours or abandon the project. In the cloud, there is virtually no additional cost to choosing 1,000 machines and accelerating such processes. This will have a dramatic impact on variability.
- **Batch processes will become real time.** Many processes—for example, accurate stock availability for online retailers—that were previously batch will move to real-time. Thus, multi-stage processes that were once sequential will now occur simultaneously, such as a manufacturing firm that can tally its inventory, check its order backlog, and order new supplies at once. This will amplify utilisation variability.


We note that even the largest public clouds will not be able to diversify away all variability; market level variability will remain. To further smooth demand, sophisticated pricing can be employed to shape demand. For example, similar to the electricity market (Figure 10), customers can be incentivised to shift their demand from high utilisation periods to low utilisation periods. In addition, a lower price spurs additional usage from customers due to price elasticity of demand. Demand management will further increase the economic benefits of cloud.

FIGURE 9: industry variability

Industry	Peak-Season Traffic / Average Traffic
Tax Services	10x
General Retail	4x
Sports (NFL)	2.5x
Travel (Airlines, Hotels)	1.5x
News	1.5x-2.0x

Source: Corporate Strategy Group, Alexa Internet, Inc.

FIGURE 10: Variable pricing in electricity




Source: Ameren Illinois Utilities

## Multi-tenancy Economies of Scale

The previously described supply-side and demand-side economies of scale can be achieved independent of the application architecture, whether it be traditional scale-up or scale-out, single tenant or multitenant. There is another important source of economies of scale that can be harnessed *only* if the application is written as a multitenant application. That is, rather than running an application instance for each customer – as is done for on-premises application and most hosted applications such as dedicated instances of Microsoft Business Productivity Online Suite (BPOS-D) – in a multitenant application, multiple customers use a single instance of the application simultaneously, as in the case of *shared* BPOS, or BPOS-S. This has two important economic benefits:

- **Fixed application labour amortised over a large number of customers.** In a single-tenant instance, each customer has to pay for its own application management (that is, the labour associated with update and upgrade management and incident resolution). In a multitenant instance of the application, that cost is shared across a large set of customers, driving application labour costs per customer towards zero. This can result in a meaningful reduction in overall cost, especially for complex applications.

FIGURE 11: utilisation overhead



Source: Microsoft


- **Fixed component of server utilisation amortised over large number of customers.** For each application instance, there is a certain amount of server overhead. Figure 11 shows an example from Microsoft's IT department in which intraday variability appears muted (only a 16 percent increase between peak and trough) compared to actual variability in user access. This is caused by application and runtime overhead, which is constant throughout the day. By moving to a multitenant model with a single instance, this resource overhead can be amortised across all customers. We have examined BPOS-D, BPOS-S, and Microsoft Live@edu data to estimate this overhead, but so far it has proven technically challenging to isolate this effect from other variability in the data (for example, user counts and server utilisation) and architectural differences in the applications. Therefore, we currently assume no benefit from this effect in our model.

## Overall Impact

The combination of economies of scale in server capacity (amortising costs across more servers), workload demand (reducing variability), and the multitenant application model (amortising costs across multiple customers) leads to powerful economies of scale. To estimate the magnitude, we built a cost scaling model which estimates the long term behaviour of costs.

Figure 12 shows the output for a workload that utilises 10 percent of a traditional server. The model indicates that 100k datacentre has an 80% lower TCO compared to a 1k DC.

FIGURE 12: economies of scale in the cloud



Source: Microsoft

## Implications


We now outline the implications of the previously described economies of cloud on Public Sector. We recognise the dual function of governments as both a *consumer* of cloud services and as an *enabler* of cloud services. Within public sector agencies and departments, CIOs and other IT leaders will make decisions about what types of cloud computing to use to run their operations. As legislators and regulators, governments will impact how and to what degree their citizens will be able to take advantage of cloud technology. In this section, we focus primarily on considerations for governments as *consumers* of cloud but in the last paragraph we will also briefly allude to the broader *enabler* aspects.

## Public Sector Cloud Benefits

The economics of cloud will have a profound impact on IT. Many IT leaders today are faced with the problem that 80% of the budget is spent on “keeping the lights on,” maintaining existing services and infrastructure. This leaves few resources available for innovation or addressing the never-ending queue of new user requests. Cloud will free up significant resources that can be redirected to innovation. Demand for general purpose technologies like IT has historically proven to be very price elastic (Figure 13). Thus, many IT projects that previously were cost prohibitive will now become viable thanks to cloud economics. However, lower TCO is only one of the key drivers that will lead to a renewed level of innovation within IT:

1. **Elasticity is a game-changer** because with cloud, renting 1 machine for 1,000 hours will be equivalent to renting 1,000 machines for 1 hour. This will enable users to rapidly accomplish complex tasks that were previously prohibited by cost or time constraints
2. **Elimination of capital expenditure** will significantly lower the risk premium of projects, allowing for more experimentation
3. **Self-service** (e.g., provisioning servers through a simple web portal rather than through a complex IT procurement and approval chain) **lowers friction** in the consumption model, enabling rapid provisioning and integration of new services and allowing projects to be completed in less time with less risk and lower administrative overhead
4. **Reduction of complexity.** Complexity has been a long standing inhibitor of IT innovation. From an end-user perspective SaaS is setting a new bar for user friendly software. From a developer perspective Platform as a Service (PaaS) greatly simplifies the process of writing new applications

FIGURE 13: price elasticity of storage



Source: Coughlin Associates

These factors will significantly increase the value add delivered by IT. Elasticity enables applications like scientific research, public health data analysis, financial market monitoring, complex event processing, logistics planning, environmental simulation, and military tactical planning, as these workloads exhibit nearly infinite demand for IT resources.

This will enable governments to better to deliver on some of its key priorities:


- **Fiscal responsibility:** In times of tight budgets, cloud can help governments achieve necessary spending cuts without cutting into essential services.
- **Better serve its citizens:** Cloud can help make governments more responsive to the needs of its citizens. It can increase the collaboration and coordination between departments. Through projects like Data.gov and Recovery.gov in the US, cloud has demonstrated that it can increase interaction with the public, allowing for more feedback and easier contact with interested parties
- **Lower emissions:** New cloud facilities are less power-hungry than existing IT infrastructure and require fewer servers to generate the same output by running them more efficiently. This can dramatically reduce the carbon footprint of IT.

### Private Cloud

As we noted earlier, Microsoft distinguishes between public and private clouds based on whether the IT resources are shared between many distinct user groups (public cloud) or dedicated to a fixed group of users (private cloud) – and factor out physical location of the data. This taxonomy is illustrated in Figure 14. Compared to traditional virtualised data centres, both private and public clouds benefit from automated management (to save on labour) and homogenous hardware (for lower cost and increased flexibility). Due to the shared nature of public clouds, a key difference between private and public clouds is the scale and scope at which they can pool demand.

- **Traditional virtualised data centres** generally allow pooling of resources within existing organisational boundaries—that is, the corporate IT group virtualises its workloads, while departments may or may not do the same. This can diversify away some of the random, time-of-day (especially if the company has offices globally), and workload-specific

FIGURE 14: COMPARING virtualisation, private cloud, and public cloud



Source: Microsoft

variability, but the size of the pool and the difficulty in moving loads from one virtual machine to another limits the ability to capture the full benefits. This is one of the reasons why even virtualised data centers still suffer from low utilisation.

- **Private clouds** move beyond virtualisation. Resources are now pooled across the entire organisation rather than by department,<sup>12</sup> and workloads are moved seamlessly between physical servers to ensure optimal efficiency and availability. This further reduces the impact of random, time-of-day, and workload variability. In addition, new, cloud-optimised application models (Platform as a Service such as Azure) enable more efficient app development and lower ongoing operations costs. Again, it should be noted that in addition to the organisation owning and running a cloud, these clouds could also be located outside national boundaries and run by a 3<sup>rd</sup> party such as an outsource provider or hoster.
- **Public clouds** have all the same architectural elements as private clouds, but bring much higher scale to bear on all sources of variability. Public clouds are also the only way to diversify away industry-specific variability and bring multi-tenancy benefits into effect.

Private clouds can address some of the previously mentioned adoption concerns – but only if they remain in country. By having dedicated hardware, they are easier to bring within the organisational firewall, which may ease concerns around **security and privacy**. Bringing a private cloud on-premise can address some of the **regulatory, compliance and sovereignty** concerns - for example, the desire to avoid subjecting data to the subpoena authority of other countries. Finally, private clouds can sometimes address issues of **customisation**, wherein specific needs can be met.

Private clouds do not address other concerns, such as **legacy**. In order to get the full benefits of cloud for legacy applications, some additional work has to be done. Private cloud do not solve this problem.

While private clouds can alleviate some of the concerns, in the next paragraph we will discuss whether they will offer the same kind of savings described earlier.

<sup>12</sup> Aggregation across organisational units is enabled by two key technologies: live migration, which moves virtual machines while remaining operational, thereby enabling more dynamic optimisation; and self-service provisioning and billing.

## The Cost of Private Clouds

While it should be clear from the prior discussion that conceptually the public cloud has the greatest ability to capture diversification benefits, we need to get a better sense of the magnitude. Figure 15 shows that while the public cloud addresses all sources of variability, the private cloud can address only a subset. For example, industry variability cannot be addressed by a private cloud, while growth variability can be addressed only to a limited degree if an organisation pools all its internal resources in a private cloud. We modelled all of these factors, and the output is shown in Figure 16.

The upper curve (dark blue) shows how the cost of a private cloud declines as scale increases. The lower curve (light blue) shows the cost for a public cloud. The public cloud curve is lower at every scale due to the greater impact of demand aggregation and the multi-tenancy effect. Public clouds are likely to be extremely large, at least 100,000 servers in size, whereas the size of an organisation's private cloud will depend on its demand and budget for IT.

This leads to some interesting comparisons. The chart shows that for organisations with a very small installed base of servers (<100), private clouds are prohibitively expensive compared to public cloud.

The only way for these small organisations or departments to share in the benefits of cloud computing is by moving to a public cloud. For larger agencies with an installed base of approximately 1,000 servers, private clouds are feasible but come with a significant cost premium of about **10 times** the cost of a public cloud, due to the combined effect scale, demand diversification and multi-tenancy.


In addition to the increase in TCO, private clouds also require upfront investment to deploy; this involves separate budgeting and commitment, increasing risk. Public clouds, on the other hand, can generally be provisioned entirely on a pay-as-you-go basis.

FIGURE 15: diversification benefits

	Sources of variability				
	Random	Time of day	Industry	Multiple resource	Growth
Private cloud	✓	✓	✗	✓	✓
Public cloud	✓	✓	✓	✓	✓

Source: Microsoft

FIGURE 16: The impact of economies of scale in THE cloud: PUBLIC AND PRIVATE




Source: Microsoft

## Finding Balance Today: Weighing the Benefits of Private Cloud against the Costs

With reference to US Government agencies we've mapped today's picture in Figure 17. The vertical axis measures the public cloud cost advantage. From the prior analysis we know public cloud has inherent economic advantages which will depend on the scale. Since we are still in the early days of scaling, the cost benefits of

FIGURE 17: cost and benefits of private clouds



Source: Microsoft

public cloud have not fully materialised yet. The horizontal axis represents the organisation's preference for private cloud. The size of the circles reflects the total budget allocated to organisations of each type. The bottom right quadrant thus represents the most attractive area for private clouds.

We mapped some of the key US Public Sector entities on this chart. The IT operations of municipal governments are, on average, much smaller than those of state governments, thus resulting in a higher public cloud cost advantage. In many cases, they also are handling lower-security data and, due to their size, have less complex IT installations, meaning that their preference for private cloud is generally lower. On the other extreme are the various arms of the military, each of which has an extremely large IT department and a strong desire for security and secrecy. In the middle are most other federal agencies and state governments. We have separated out Health & Human Services, the Department of Justice, the Social Security Administration and Homeland Security as these agencies have particularly sensitive missions with private and highly confidential data. They also have some of the largest IT organisations in the US Federal government.


We've made a number of assumptions in Figure 17. For examples, it assumes that, at the Federal level, each agency is an independent customer.<sup>13</sup> Some governments have expressed a desire to construct a cloud for all departments and agencies. Pooling resources in this way could dramatically improve the economics and appeal of a private cloud.

However, building such shared or "community" clouds amongst various agencies must be balanced against the desire for a certain degree of independence (departments may want to maintain more control over their choice of services) and sovereignty (states and provinces, for example, may prefer more independence from the federal government for their IT resources). These projects also require coordination and agreement at each stage of planning and investment. This could delay the adoption of cloud, postponing its benefits.

<sup>13</sup> For state and municipal governments, we have assumed that all IT for each state or city is consolidated across departments. Even with this combination, these IT organisations are likely sub-scale for private cloud implementations; segmenting them by department only reduces the scale further

Another important nuance is that IT is not monolithic. Each organisation's IT operation is segmented into workload types, such as email or ERP. Each of these has a different level of sensitivity and scale, and Public Sector CIO surveys reveal that preference for public cloud solutions currently varies greatly across workloads (Figure 18).

FIGURE 18 Government CIOs: cloud-Ready Workloads (2010)



Source: Microsoft survey question "In the next 12-24 months, please indicate if a cloud offering would augment on-premise offering or completely replace it"


Before we draw final conclusions, we need to make sure we avoid the horseless carriage syndrome and consider the likely shift along the two axis (economics and private preference).

### The Long View: Cloud Transition Over Time

As we pointed out in above, it is dangerous to make decisions in the middle of a disruption, without a clear vision of the end state. IT leaders need to design architecture with a 5-10 year vision in mind. We therefore need to consider how the long term forces will impact the position of the bubbles on Figure 17

We expect two important shifts to take place. First, the economic benefit of public cloud will grow over time. As more and more work is done on public clouds, the economies of scale we described in Section 2 will kick in, **and the cost penalty on private clouds will increase over time**. Customers will increasingly be able to tap into the supply-side, demand-side and multi-tenancy savings as discussed previously. On Figure 19, this leads to an upward shift along the vertical axis.

FIGURE 19: expected preference shift for public AND private cloud



Source: Microsoft

At the same time, some of the barriers to cloud adoption will begin to fall. Many technology case studies show that, over time, concerns over issues like compatibility, security, reliability, and privacy will be addressed. This will likely also happen for the cloud, which would represent a shift to the left on Figure 19.

In summary, while there are real hurdles to cloud adoption today, these will likely diminish over time, while at the same time the public cloud advantage will only grow over time as cloud providers

capture the benefits of economics we discussed in the foregoing. While the desire for a private cloud is mostly driven by security and compliance concerns around existing workloads, the cost effectiveness and agility of the public cloud will enable new workloads and we believe public sector organisations will move many workloads to the cloud as a result.

## **Confronting and Managing Risk in the Public Sector Cloud scenarios**

In the previous section we have outlined the economic drivers that characterise cloud computing and in so doing exposed a risk dynamic that also requires discussion. The move away from highly centralised systems towards more a profoundly more decentralised model is fundamentally important when considering not only the security of those systems but the overall risk associated with the system within which they operate.

### **Current practice**

The approach to IT security risk in place today largely centres on formalising and managing perimeter security based on the historical practice of applying a top-down hierarchical control over an agency's assets and resources. This is instantiated and managed against a set of prescriptive requirements and guidelines maintained by Defence Signals Directorate and has served government well. However, the ability for prescriptive guidelines to continue address the existing varied risk environments in today's existing IT deployments let alone those we face with cloud computing, is extremely limited.

Microsoft is aware of a number of Federal Government activities that evidence an awareness that the overall environment is changing and that a new approach is required:

- i. We note that the Australian National Audit Office Review<sup>14</sup> undertaken by KPMG noted a need for greater alignment of its own Corporate ICT risks to ICT Strategic Plan.<sup>15</sup> We believe a similar alignment for other agencies would help to move towards a more integrated risk approach necessary for cloud computing.
- ii. The Attorney General's Department (AGs) is responsible for setting the Government's protective security policy. We note that they have recently drafted a revision of the Protective Security Manual called the *Protective Security Policy Framework*<sup>16</sup> and far from a simple name change, it seems that this represents an acknowledgement that both a framework is required and that it needs to be principles-based. There appears to be a clear move to establish a cultural change in the way agencies address risk in their operations and take ownership for this process. We commend the government for this and encourage further development in this direction.

### **Taking a Principles & Risk approach**

Microsoft has benefited from an internal pooling of experience across a range of countries in both the public and private sectors and we see a clear theme emerging in response to the way risk is being viewed and managed – across a variety of our lives.

The Global Financial Crisis caused a system destabilisation that impacted almost every part of the globe in large part due to a poor understanding of the way risk was actually distributed through that system. Electricity brown outs in a number of countries from the USA to Europe have demonstrated

---

<sup>14</sup> Australian National Audit Office IT Performance Review -

<http://www.anao.gov.au/uploads/documents/IT%20Performance%20Review.pdf>

<sup>15</sup> Recommendation 9.1 "It is suggested that the ANAO consider, in its future consideration of the ICT Strategic Plan creating a clear link from the identified ICT Corporate Risks to the ICTSP."

<sup>16</sup> Australian Government Protective Security Policy Framework - <http://www.ag.gov.au/pspf> -

the risks implicit in complex systems and their potentially pernicious and unintended consequences. Closer to home, many IT projects have failed due to a lack of understanding of the system and the impact a small change may have on it. Understanding the system and applying a set of well-constructed principles to managing it is critical.

Microsoft's own experience in designing, building and operating global scale cloud infrastructure which must respect a growing list of compliance requirements also supports the value in adopting a Principles and Risk approach and we encourage the government to move even more forcefully in this direction and with greater speed. Microsoft believes that this is key to ensuring that the true impact of decisions to implement cloud alternatives, or not, can be appropriately made.

We are also aware that in some cases, applying this approach will identify new risks that are untenable for government, but also experience has shown us that in other cases, the move to cloud can support immediate improvements in the risk profile of activities. A best-in-class managed infrastructure where the high cost of patching, managing and provisioning complex IT systems can be amortised over a large pool of users at a low unit cost can actually improve the net risk position of those adopting this alternative.

### **Example of APRA**

Finance is a sector that has long been a heavy user of IT and it is well known that the significant reduction in their operating costs as a percentage of gross profit over the last 10 years has been largely due to the deployment of IT. In fact much of a bank's activity now is IT based and for that reason the Australian Prudential Regulatory Authority (APRA) maintains a strong role in reviewing IT risk as part of its overall statutory duties.

APRA utilises a Principles and Risk approach to the way in which it monitors the banking system and individual deposit taking institutions within it and we encourage the government to consider what appear to be well functioning and globally respected practices managing security and risk in this sector.

### **Role of Interoperability & Standards**

As we have described in the previous section, the role of market forces together with technological enablement provides a set of constraints that informs the ultimate development path of any innovation. The concept of how this happens - called "Technology Diffusion" - has been evolving for many decades<sup>17</sup> and was popularised by Geoffrey Moore in his book "Crossing the Chasm"<sup>18</sup>.

Although Moore's text is largely directed at the question of how to market technology, the discussion of the context and underlying dynamics of technology diffusion are valuable as we attempt to frame the possible adoption of cloud computing and the high order bit from his analysis is the fact that the path of technology origination, development, market release, market acceptance and obsolescence, all follow a predictable cycle of intense innovation characterised by many divergent technology vendors and technology approaches, followed by a consolidation of platforms and vendors and then a period of stasis before a gradual decline as the next wave of technology evolution washes through.

Clearly ensuring end customers of this technology innovation process are served well is critical and a major requirement therefore becomes interoperability between existing deployed technology and the new alternatives to ensure that true value is realised.

---

<sup>17</sup> Technology adoption lifecycle overview, Wikipedia - [http://en.wikipedia.org/wiki/Technology\\_diffusion](http://en.wikipedia.org/wiki/Technology_diffusion)

<sup>18</sup> Crossing the Chasm, Geoffrey Moore, Wikipedia - [http://en.wikipedia.org/wiki/Crossing\\_the\\_Chasm](http://en.wikipedia.org/wiki/Crossing_the_Chasm)

However, although standards are a key interoperability tool, it is important that they not be developed or mandated too early in this diffusion cycle, as the diversity of vendor and technology approaches requires a degree of shake-out before standardisation can take practical effect.

If we apply this to cloud services, the economic, or benefit drivers for Governments and businesses alike are: consolidation of IT infrastructure, scaling IT systems for future growth, and enabling innovative services and activities that were not possible before.

Therefore to help organisations realise the benefits of cloud services, technology vendors, including Microsoft, are investing in identifying and solving the challenges presented by operating in mixed IT environments, and collaborating to ensure that their products work well together.

In fact, although the industry is still in the early stages of collaborating on cloud interoperability issues, there has already been considerable progress. But what does “cloud interoperability” mean, and how is it benefiting people today?

### **Interoperability in the Cloud**

As we have noted publically<sup>19</sup>, cloud is intrinsically about connectivity between people, technology, devices, processes and interoperability is at the heart of this. The present headline use-case is specifically about one cloud solution, such as Windows Azure, being able to work with other clouds such as Amazon’s but the broader requirement we believe is how interoperability between specific clouds can be facilitated to embrace other platforms and other applications, not just other clouds. Microsoft is also keenly aware that governments also want the flexibility to run applications either locally or in the cloud, or on a combination of the two – and in fact this was alluded to in the Cloud Strategy Paper.

If the IT industry and its customers are to realise the innovation, cost and flexibility benefits of cloud computing, Microsoft believes that software must be created that is more open from the ground up, supporting existing standards, enabling customers to use Microsoft cloud services together with open source technologies such as PHP and Java, and ensuring that the company’s existing products work with the cloud.” Our own development of the Azure Cloud platform followed this very principle.

### **The Promise of Cloud Interoperability**

A good example of real-world application of this can be seen a project called the Eye on Earth. Microsoft has been working with the European Environment Agency, to help the agency simplify the collection and processing of environmental information for use by government officials and the general public. By using a combination of Windows Azure, Microsoft SQL Azure and pre-existing Linux technologies, Eye on Earth pulls data from 22,000 water monitoring points and 1,000 stations that monitor air quality. It then helps synthesise this information and makes it available for people to access in real time in 24 languages. We think this is the realistic future for not only cloud-based but also device-based computing.

### **The Hard Work of Interoperability**

This level of openness and interoperability doesn’t happen by accident. The technical work of interoperability is challenging, and requires a concerted effort on multiple fronts and a measured, pragmatic approach in how technology is developed. Microsoft’s efforts in this area include

---

<sup>19</sup> Open Cloud, Open Data, Jean Paoli, Microsoft - <http://www.oscon.com/oscon2010/public/schedule/detail/15643>

designing its cloud services to be interoperable. The Windows Azure platform, for example, supports a variety of standards and protocols. Developers can write applications to Windows Azure using PHP, Java, Ruby or the Microsoft .NET Framework. They can utilise components of functionality developed and delivered from separate non-Microsoft clouds or source data from repositories resident in an agency's own data centre. The choice is theirs.

Many of these product developments are the result of diverse feedback channels that Microsoft has developed with its partners, customers and other vendors.

For example, in 2006 Microsoft created the Interoperability Executive Customer (IEC) Council, a group of 35 chief technology officers and chief information officers from organisations around the world. They meet twice a year in Redmond to discuss their interoperability issues and provide feedback to Microsoft executive which are then incorporated into specific work actions for Microsoft product groups.

In addition, Microsoft recently published a [progress report](#), sharing for the first time operational details and results achieved by the Council across six work streams, or priority areas. And the Council recently commissioned the creation of a seventh work stream for cloud interoperability, aimed at developing various standards related to the cloud, working through business scenarios and priorities such as data portability, and establishing privacy, security, and service policies around cloud computing. We think this kind of direct, intentional feedback is critical.

Microsoft also participates in the Open Cloud Standards Incubator, a working group formed by the [Distributed Management Task Force \(DMTF\)](#), a consortium through which more than 200 technology vendors and customers develop new standards for systems management. AMD, Cisco, HP, IBM, Microsoft, Red Hat and VMware are among a handful of IT vendors that lead the Open Cloud Standards Incubator, creating technical specifications and conducting research to expedite adoption of new cloud interoperability standards.

### **Openness, Industry Collaboration to Pave the Way**

Developers also play a critical role. Microsoft is part of [Simple Cloud](#), an effort it co-founded with Zend Technologies, IBM and Rackspace designed to help developers write basic cloud applications that work on all of the major cloud platforms.

Microsoft is also engaging in the collaborative work of building technical "bridges" between Microsoft and non-Microsoft technologies, such as the recently released Windows Azure Software Development Kits (SDKs) for PHP and Java and tools for Eclipse version 1.0, the new Windows Azure platform AppFabric SDKs for Java, PHP and Ruby, the SQL CRUD Application Wizard for PHP, and the Bing 404 Web Page Error Toolkit for PHP. Each is an example of the Microsoft Interoperability team's yearlong work with partners to bring core scenarios to life.

Though the industry is still in the early stages of collaborating on cloud interoperability issues, great progress has already been made. The average user may not realise it, but this progress has had a significant positive impact on the way in which we work and live today.

Cloud interoperability requires a broad perspective and creative, collaborative problem-solving. Looking ahead, Microsoft will continue to support an open dialogue among the different

stakeholders in the industry and community to define cloud principles and incorporate all points of view to ensure that in this time of change, there is a world of choice.

### **AGIMO whole of government role in JTC-1 Cloud standards work**

As interoperability requirements emerge and mature the need for standardisation will surface. Given the global nature of Cloud Services, standardisation in this area will rely on existing international processes to ensure that the needs of all stakeholders are heard and represented. Microsoft believes that it is critical that Australia's national interests are represented as part of this.

In recent months Standards Australia has been circulating a proposal that, if implemented, will ensure that Australia gets the opportunity to play a much more active role in international IT standardisation within the ISO system.

The proposal suggests the formation of a group in Australia to engage with JTC1, the principal group within ISO and IEC focused on IT standards work, and then within JTC1 several subcommittees such as the "*Distributed Applications Processing and Services*" Technical Committee (TC)- also known as JTC1/SC38 - which is currently establishing several projects that will lead to the standardisation of key cloud services practices and protocols.

Microsoft fully supports the proposal that Standards Australia has put forwards and should it be implemented we will work with industry associations to ensure that Australian industry and government is able to both track the development of international cloud service related standards and also firmly express any uniquely Australian requirements that could arise.

Given the necessary collaborative nature of cloud services we would encourage AGIMO to consider backing the proposal put forward by Standards Australia and becoming the central voice for the Federal Government within the Standards Australia system.

### **Additional cloud examples in Australia**

Microsoft commends AGIMO on developing an inventory of examples of cloud use and we also take this opportunity to supplement this list with several other examples which we believe are relevant:

#### **MYOB**

MYOB is an Australian success story with over 1 million businesses across Australia and New Zealand and over 100 products and services and working with over 85% of Australian accountants. They are moving their business model to cloud based services using the Microsoft Windows Azure platform as their online system to deliver cloud based services and back their customers.

In addition to this, they are opening up cloud based APIs for integration of third party products, reconciliation and eLodgement systems and this will drive not only a vastly improved outcome for many small businesses but also a richer ecosystem of partners who are also now encouraged to adopt cloud platforms as a way of growing their businesses.

#### **Brisbane City Council**

Microsoft believes that some of the most valuable aspects of cloud services will be realised by those many local Australian software and services businesses that have already begun to utilise cloud infrastructure and platform services to build and deliver their own SaaS offerings to customers. We outline below the particularly pertinent use case by Open Windows with Brisbane City Council as an example of this:

In 1994 Open Windows marketed and sold their start-up product to over 70% of the State of Victoria's local municipalities (of which there were 77 at the time) in the first 2 years of the business.

Open Windows was subsequently one of the first vendors (or the first) to take an Australian Public Sector organisation to Microsoft .NET when Open Windows redeveloped the CONTRACTS application

in .NET while it was still in beta release in 2001. Today Open Windows has over 150 Australian and New Zealand Government Departments and tier 1 corporations using the system (with 80% Public Sector clientele).

Open Windows CONTRACTS is an enterprise Contract Lifecycle Management application, which provides pro-active management of procurement processes and contract commitment management and that has become the standard in Australian Local Government.

Functionally the system can be broken down as follows:

- Contract Procurement & Sourcing,
- Contract Compliance & Management, and
- Contract Analysis and Intelligence.

Open Windows is the first vendor to take an ISV application on the Azure platform to a Public Sector client. Not only are Brisbane City Council the first Public Sector organisation in Australia to consume an ISV application from the Azure Cloud, they are the biggest Local Government Council in the Country and one of the biggest in the World.

Not only has Open Windows been able to take Brisbane City Council to the Azure Cloud, but Council has embraced it in a big way, already creating over 11,000 user profiles in the system, and looking to take the application to their whole organisation.

Open Windows CONTRACTS CLOUD is 100% browser-based and is offered as on-premises or via the Cloud using Windows Azure, SQL Azure, and Azure Storage.

By utilising the Windows Azure platform, Brisbane City Council, in partnership with Open Windows is able to lead the Australian Public Sector into the Microsoft cloud with the following significant benefits:

- Access a functionally rich application traditionally delivered on-premises from the cloud without compromise to performance or functionality,
- Manage a more rapid implementation at lower Total Cost of Ownership, Azure lowers the cost for quality hosting accessibility
- Provide integrated Infrastructure as a Service in the form of Azure Storage as a highly cost effective storage solution for commercially sensitive documents, and scale up and down as required fully linked to a Cloud application and database,
- Access the application from any Council desktop or laptop without the burden of administering and licensing server hardware and software SOEs,
- Scale up or down processing power to the application as required to meet peaks and troughs in demand without leaving resources idle and money on the table, and
- Eliminate internal IT burden of supporting vendor delivery.
- The nature of Azure's UI and design means that decisions around up scaling or downscaling the infrastructure is virtually immediate, as opposed to traditional hosting models
- We are already able to pass on a lower cost of infrastructure using Azure, and look forward to the future cost benefits that volume on the Contracts CLOUD will bring; cost of delivery and cost of maintenance (for both customer and vendor),
- Open Windows found the Azure portal easy to use and lowered our configuration encumbrance because it effectively comes fully configured.

In addition to the foregoing Australian examples, we have also provided a list of pertinent global Public and Private sector examples of cloud computing that we believe may be useful. These are located in Appendix A.

## **Additional comments**

### **Additions and omissions**

We note that in section 2.2, Cloud Service Capability, Microsoft should also be represented as a major current provider of SaaS as well as IaaS with examples including our *Office365* productivity/collaboration suite and *Windows Azure Virtual Machine Role* respectively<sup>20</sup>.

Similarly, in Attachment 3 which outlines a list of the major vendors for each of the 3 modes of cloud services, Microsoft, although represented as a provider of PaaS and IaaS, is omitted as offering SaaS. Naturally, we are keen to see this recognised given the significance of the *Office365* cloud service.

### **Industry involvement**

Microsoft commends the government on the suggested formation of a Cloud Information Community and we would welcome the opportunity to participate in an open collaborative way with AGIMO as deemed appropriate.

---

<sup>20</sup> [http://msdn.microsoft.com/en-us/WAZPlatformTrainingCourse\\_VMRoleLab](http://msdn.microsoft.com/en-us/WAZPlatformTrainingCourse_VMRoleLab)

## Further consultation

Microsoft thanks Department of Finance for considering its views on some of the proposals made, and would welcome the opportunity to further discuss any of the points raised in this submission. To do so, please contact:

Hala Batainah	Federal Director, Microsoft Australia email: <a href="mailto:hala.batainah@microsoft.com">hala.batainah@microsoft.com</a>
Greg Stone	Chief Technology Officer Microsoft Australia email: <a href="mailto:gstone@microsoft.com">gstone@microsoft.com</a>
Simon Edwards	Director of Corporate Affairs Microsoft Australia email: <a href="mailto:simone@microsoft.com">simone@microsoft.com</a>

## Appendix A

The following are some pertinent case studies that demonstrate the value of leveraging cloud computing. We have included both public and private sector examples :

### Public Sector related (Non Australia)

**Freedom Speaks:** <http://www.microsoft.com/casestudies/Microsoft-Visual-Studio-2008/FreedomSpeaks/Social-Network-Uses-Cloud-Computing-to-Connect-Citizens-with-Elected-Officials/4000007618>

#### Social Network Uses Cloud Computing to Connect Citizens with Elected Officials

Founded in 2006, FreedomSpeaks is one of the first political social networks in the United States. The organisation advocates open communication and transparency between elected officials and their constituents, and uses technology to support this endeavor. In addition to its political activism website, the company acquired a mobile product called CitySourced that allows citizens to send comments and concerns directly to local governmental agencies via smartphone devices. To support this offering, FreedomSpeaks needed to migrate its existing infrastructure to an on-demand, internet-based—or cloud-computing—technology platform. After evaluating competing cloud-based offerings from Amazon and Rackspace, FreedomSpeaks selected the Windows Azure technology platform for its cost-effectiveness, easy integration with the company's existing infrastructure, and scalability.

**NASA:** <http://www.microsoft.com/casestudies/Cloud-Services/National-Aeronautics-and-Space-Administration/New-NASA-Web-Site-Engages-Citizens-to-Help-Explore-Mars/4000008289>

#### New NASA Web Site Engages Citizens to Help Explore Mars

Researchers at the NASA Jet Propulsion Laboratory (NASA/JPL) wanted to solve two different challenges—providing public access to vast amounts of Mars-related exploration images, and engaging the public in activities related to NASA's Mars Exploration Program in order to encourage learning in science, technology, engineering, and mathematics. Using a variety of technologies, NASA/JPL created its new BeAMartian Web site . The site provides entertaining and engaging ways to view and interact with information delivered by Mars-based rovers and orbiters. The goal is to let the public participate in exploration, making contributions to data processing and analysis. It also provides a platform that lets developers collaborate with NASA on solutions that can help scientists analyze vast amounts of information that can be used to understand the universe and support future space exploration.

**Florida House of Representatives:** <http://www.microsoft.com/casestudies/Windows-Communication-Foundation/The-Florida-House-of-Representatives/State-Helps-Improve-Census-Counts-Avoids-300-000-in-Capital-Costs-with-Cloud-Services/4000007975>

#### State Helps Improve Census Counts, Avoids \$300,000 in Capital Costs with Cloud Services

Every 10 years, the Florida Legislature draws its district boundaries. The public is encouraged to participate, but previously had only a desktop application to submit recommendations for new district boundaries. For the 2010 census, the Florida House of Representatives wanted to develop a web-based application to allow even greater public participation, but the cost to build the infrastructure for the application was prohibitive.

The Florida House of Representatives decided to implement the Windows Azure platform, including Microsoft SQL Azure, but first tested the technology with another project: My Florida Census, a website that helps residents report whether or not they've been counted by the U.S. Census. By using Windows Azure for both projects, the House avoided capital expenditures, quickly scales up or down based on demand, and facilitated a positive impact on the state budget.

**European Environment Agency EEA:** <http://www.microsoft.com/casestudies/Microsoft-ASP.NET/European-Environment-Agency-EEA/Environment-Agency-s-Pioneering-Online-Tools-Bring-Revolutionary-Data-to-Citizens/4000006197>

Environment Agency's Pioneering Online Tools Bring Revolutionary Data to Citizens

An agency of the European Union, the European Environment Agency (EEA) provides independent and reliable information on the environment for policy makers and the general public. The agency is working towards raising environmental awareness across Europe by delivering easy-to-understand information about a number of environmental topics—among them, water and air quality. It also encourages citizens to contribute their own observations about the environment around them. Working with Microsoft, it developed the Eye On Earth platform, based on the Windows Azure “cloud” services operating system. Users can view water or air quality from the 32 member countries of the EEA, using high-definition Bing maps. The EEA has also launched the Environmental Atlas of Europe, which features stories told by eyewitnesses about their first-hand experiences of climate change. Both solutions can help broaden awareness of the impacts of environmental change and help people in Europe make better-informed choices about their environment.

**City of Miami:** <http://www.microsoft.com/casestudies/Cloud-Services/City-of-Miami/City-Government-Improves-Service-Offerings-Cuts-Costs-with-Cloud-Services-Solution/4000006568>

City Government Improves Service Offerings, Cuts Costs with “Cloud” Services Solution

The City of Miami, even when limited by a tight budget, looks for ways to improve the services it offers citizens. The city wanted to develop an online application to record, track, and report on nonemergency incidents, but the application's sophisticated mapping technology would require significant computing resources. Further constrained by long hardware-procurement cycles, the city needed a cost-effective, scalable solution that would maximise its available resources. The city developed its 311 application on the Windows Azure platform, taking advantage of scalable storage, processing power, and hosting provided by Microsoft. As a result, the city was able to reduce IT costs, improve the services it offers citizens, and deliver those services faster than before. It also now relies on a cost-effective disaster-recovery model, an important benefit in this hurricane-prone region.

### **Commercial (Australia)**

**Computershare:** <http://www.microsoft.com/casestudies/Windows-Azure/Computershare/Investor-Communications-Specialist-Floats-IPO-Platform-into-the-Cloud-and-Cuts-Costs/4000008997>

Investor Communications Specialist Floats IPO Platform into the Cloud and Cuts Costs

Global investor services specialist Computershare provides communications services to 30,000 clients in more than 20 countries. One of its technology tools – HKeIPO – is a platform that enables companies in Hong Kong that are about to float on the exchange to manage online prospectus requests and receive share applications. Computershare hosted HKeIPO at its own data centres. However, the large server capacity the platform required was idle between floats, and bandwidth spikes during the short float periods were enormous.

In 2010, Computershare exported this platform into the cloud with Windows Azure. The move took six weeks and within a month the platform was used to host investor enquiries for one of Hong Kong's biggest initial public offerings (IPOs). The cloud-based platform-as-a-service easily accommodated peaks in service demand and cost less than five per cent of the nominal costs of onsite hosting.

## **Commercial (Non-Australia)**

### **T-Mobile:**

[http://www.microsoft.com/casestudies/Case\\_Study\\_Search\\_Results.aspx?Type=1&ProTaxID=3446&IndTaxID=1196&LangID=46](http://www.microsoft.com/casestudies/Case_Study_Search_Results.aspx?Type=1&ProTaxID=3446&IndTaxID=1196&LangID=46)

#### Mobile Operator Speeds Time-to-Market for Innovative Social Networking Solution

T-Mobile USA, a leading provider of wireless services, wanted to create new mobile software to simplify communications for families. The company needed to implement the application and its server infrastructure while facing a tight deadline. T-Mobile decided to build the solution with Microsoft Visual Studio 2010 Professional and base it on Windows Phone 7 and the Windows Azure platform. By taking advantage of an integrated development environment and cloud services, the company completed the project in just six weeks. Using a cloud platform instead of maintaining physical servers has also simplified management. As a result, developers have more time available to focus on enhancing the application. Customers will benefit from a streamlined, reliable communications solution with strong security, and T-Mobile is already designing new features for users to enjoy.

**Advanced Telemetry:** <http://www.microsoft.com/casestudies/Microsoft-ASP.NET/Advanced-Telemetry/Startup-Uses-Cloud-Computing-to-Change-Business-Model-Becomes-Instantly-Profitable/4000009076>

#### Startup Uses Cloud Computing to Change Business Model, Becomes Instantly Profitable

Advanced Telemetry began by manufacturing and marketing its flagship remote energy-monitoring software—and paying for the servers required for its customers' data. The relationship between adding customers and paying for more rack space was a serious barrier to expanding market share. Advanced Telemetry then moved its IT infrastructure to the Windows Azure cloud-computing platform, using Microsoft SQL Azure data storage and Windows Azure AppFabric connectivity services. It also switched to a new business model whereby OEMs can license an instance of its telemetry software in Windows Azure, and then offer the product in new markets. This move reduced IT infrastructure expenses by 75 percent, introduced new revenue from customisation, reduced marketing costs by 80 percent and turned the company into an agile, profitable operation that can respond quickly to customers in any market.

**Associated Press:** <http://www.microsoft.com/casestudies/Windows-Azure-Platform/Associated-Press/Global-News-Provider-Gives-Newspapers-a-New-Business-Model-for-Success-on-the-Web/4000008009>

#### Global News Provider Gives Newspapers a New Business Model for Success on the Web

As the audience for newspapers shifts to online news, traditional publishers are seeking a commercially viable business model to bring their content to the web. The Associated Press (AP) hopes to help meet that need with its News Reader, a rich Internet application for delivering content that isn't bound by a browser or limited to online use. The News Reader is the result of collaboration among the AP, Microsoft, and Vectorform, a global digital-experience studio. It is built with Microsoft technologies from end to end, from the highly visual and interactive interface based on Microsoft Silverlight, to a back end based on the Windows Azure platform, which gives it virtually unlimited scalability. With Microsoft technologies, the AP is speeding time-to-market, enabling broad newspaper participation, enhancing the consumer experience, and creating a compelling environment for advertisers.

**Lockheed Martin:** <http://www.microsoft.com/casestudies/Windows-Azure/Lockheed-Martin/Lockheed-Martin-Merges-Cloud-Agility-with-Premises-Control-to-Meet-Customer-Needs/4000007971>

#### Lockheed Martin Merges Cloud Agility with Premises Control to Meet Customer Needs

Headquartered in Bethesda, Maryland, Lockheed Martin is a global security company that employs about 136,000 people worldwide and is principally engaged in the research, design, development, manufacture, integration, and sustainment of advanced technology systems, products, and services. The company wanted to help its customers obtain the benefits of cloud computing, while balancing security, privacy, and confidentiality concerns. The company used the Windows Azure platform to develop the Thundercloud™ design pattern, which integrates on-premises infrastructure with compute, storage, and application services in the cloud. Now, Lockheed Martin can provide its customers with vast computing power, enhanced business agility, and reduced costs of application infrastructure, while maintaining full control of their data and security processes.

**Total Computer:** <http://www.microsoft.com/casestudies/Windows-Communication-Foundation/Total-Computer/Law-Enforcement-Application-Provides-Officers-with-Rich-Functions-and-Data/4000006783>

#### Law Enforcement Application Provides Officers with Rich Functions and Data

Total Computer wanted to create a single commercial law enforcement application with complete local and national real-time data sharing across disparate database sources. The company wanted to provide both mobile and offline functionality across a variety of devices. Its developers used an array of Microsoft technologies to develop Total Enforcement—which runs on various client systems and also in the Windows Azure cloud—to connect law enforcement officers with critical information from diverse sources in real time using any device. With Microsoft software and services, Total Computer developed the system quickly, providing numerous functions on a flexible, powerful platform.

**VeriSign:** <http://www.microsoft.com/casestudies/Windows-Azure/VeriSign/Security-Firm-Helps-Customers-Create-Highly-Secure-Hosted-Infrastructure-Solutions/4000005923>

### Security Firm Helps Customers Create Highly Secure Hosted Infrastructure Solutions

VeriSign wanted to provide customers who deliver services on hosted infrastructures with the same highly secure encryption technology used in on-premises data centers. The company partnered with Microsoft to provide Windows Azure™ platform customers with enhanced protection for critical business data. With VeriSign® SSL Certificates on Windows Azure, customers receive an added layer of online security, while reducing management and maintenance costs.