I D C E X E C U T I V E B R I E F

Future Proofing Your Datacentre Facilities Services

Best Practice Tender Document Creation

March 2013

Matthew Oostveen, Head of Research, IDC Australia.

Sponsored by NEXTDC

Introduction

This document focuses on the requirements that an end-user organisation should include in a tender document when acquiring services from a datacentre facilities provider. The document focuses on the changing requirements of the datacentre and on a medium-term view of priority requirements. This is not an exhaustive list of datacentre selection criteria. It includes IDC research that looks at the adoption of datacentre-centric technologies and the impact it has on operating environments.

The document also investigates the direction of IT and takes a view on the future of IT and what the datacentre facilities providers will need to offer the market to cater for the change and growth ahead.

Based on this research, IDC makes recommendations to highlight the priorities for defining selection criteria in the request for proposal (RFP) for finalising a datacentre facility provider and or co-location services.

Definitions

IDC has conducted some of Australia's most detailed research into the Australian datacentre market. In 2012, IDC investigated over 300 Australian CIOs and CTOs to ask detailed questions about their datacentres, asking whether they be enterprise operated or operated by a third party. This not only provided an excellent insight into the current state of the Australian datacentre landscape but also enabled IDC, with the wealth of historical data, to assess the changing consumption patterns and predict important emerging trends in the local market. This holistic view of the Australian datacentre market enables understanding of current and future challenges.

Trends

IDC research revealed some interesting facts about Australian datacentres. First, Australian datacentres are the oldest in the Asia/Pacific region. The average age of an Australian datacentre is over 18 years, which means its construction dates back to 1994. It is

Level 20, 8-20 Napier St, North Sydney NSW 2060 Australia P.61.2.9925.2298

little wonder then that we observe many CTOs and CIOs grappling with several issues pertaining to running and maintaining these facilities. Historical market data shows that the majority of datacentres were populated by large mainframe systems with internal water cooling capabilities and RISC Unix systems, sparsely positioned across the datacentre floor. In fact, there was not even a need to configure the floor space into hot and cold aisle configurations. What was missing from the datacentre of 1994 is the now ubiquitous x86 server.

The datacentre facilities architect back then did not have the advantage of the decades of research and data we now have available to us. It was practically impossible for them to foresee the "enterprisation" of the x86 server which today dominates the datacentre floor commanding a lion's share of space. Nor could they have had the slightest inkling of the exponential growth in storage to be the behemoth it is today.

Since 2000, the average number of servers per rack has increased over 50%, primarily driven by the increased adoption of more dense factors. This development calls for increased and effective management of the energy and thermal conditions of IT racks.

What are the challenges faced today?

IDC identifies four major challenges which it considers mission or business critical in the day-to-day operations of Australian datacentres. These, if left unchecked, can potentially disrupt and adversely impact business operations. Let's examine each of them closely.

Space. Whilst the cost of acquiring a server has decreased significantly over the years, the cost of keeping that server housed has continued to grow. Australian CIOs and CTOs today are being constantly challenged by demand for more space due to expansion of their organisations' datacentres.

Power. Aligned with the increased uptake of x86 servers with high thermal outputs and energy hungry CPUs and RAM, IDC believes power will soon be the number one concern for datacentres in Australia.

The major challenge faced by datacentre managers today is that they lack a comprehensive understanding of their datacentres' energy environment. To operate efficiently, there is a need for visibility into how and where power is being consumed by IT systems across the environment and whether or not there is adequate cooling capacity that is correctly provisioned to the IT loads. Energy efficiency extends beyond the IT systems to the datacentre infrastructure in support of those systems. Today, many datacentres have over-provisioned or ineffectively provisioned their power and cooling, resulting in stranded power that does not make full use of the datacentre capacity, thus creating hotspots that risk system failures.

Weight. Modern storage and compute is far denser than previous iterations. This has led to spikes in weight as older and lighter systems are replaced with densely populated racks. IDC research

©2013 IDC

reveals that close to 20% of end-user organisations struggle with weight in their datacentres. The worldwide trend also shows a similar trend of increasing rack density (please refer to Figure 1).

It's interesting to note that there are some compute racks weighing over 1 tonne, with a footprint of less than 1 square metre.

The significance of this increase in weight and its impact means that end-user datacentres located in non-specially constructed buildings may be exceeding the weight limitations and floor loading capacity.

Figure 1

Worldwide Datacentre and Rack Density, 2000–2014

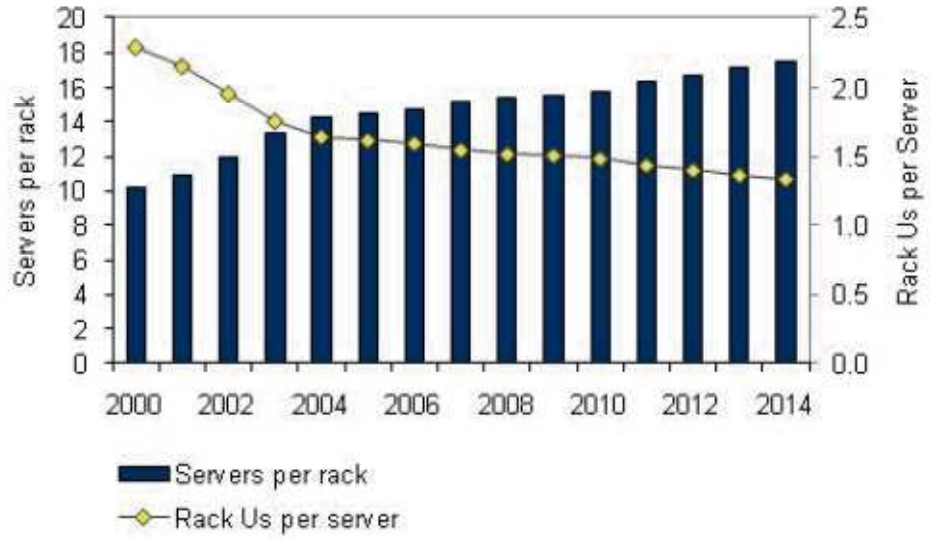

Source: IDC, 2013

Cooling: Modern server and storage infrastructure radiates large amounts of heat, which places pressure on datacentre cooling facilities. As density increases in line with power consumption, according to the law of thermodynamics, more heat will be created. Cooling facilities in Australian datacentres currently are inadequately equipped to overcome the several issues poised by this challenge.

What will be the challenges of tomorrow?

As part of the *2012 Australia Datacentre Study*, IDC surveyed 300 CTOs, CIOs, and facilities managers across Australia to investigate the current state of the datacentre landscape and gather data that allows us to forecast technology trends across the nation.

The study showed an increasing trend for organisations to move away from owning and managing their own datacentre facilities. Instead, they prefer to lease floor-space from third-party providers specialising in co-location services. The main drivers for adopting

©2013 IDC

third-party datacentres were rising costs, redundancy, and increased power demands (please refer to Figure 2).

Figure 2

Drivers for Adopting Third-Party Datacentres

Q. What drove you to use a third-party datacentre?

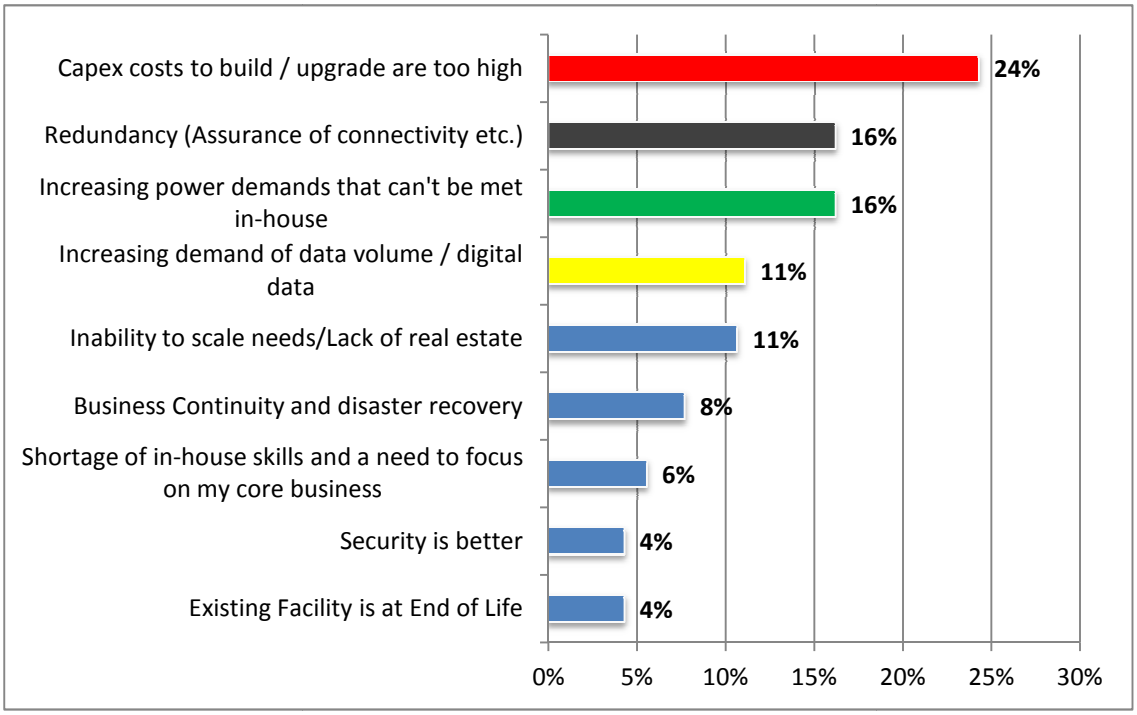

Source: IDC's *2012 Australia Datacentre Study*, (n = 300)

On the supply side, scalability, delivering on SLAs, accessing carrier services, and securing additional power were some of the top factors that influenced the end-user organisations' experience with the services of the third-party datacentre facility (please refer to Figure 3).

Connectivity is of critical importance when selecting a datacentre facility services provider, as it allows for interconnection between a variety of network carriers. IDC views this as a primary service delivery from a datacentre provider; however, research reveals (please refer to Figure 3) this basic need is not being met by much of the marketplace. 23% of end users reported difficulty in accessing carrier services. This can result in such problems as the organisation not having appropriate access to telecommunications providers, increased communications charges, and causes "lock-in" as hosting in a neutral location allows providers to be changed without relocating infrastructure to another location.

The research validates anecdotal evidence found by IDC that uncovered large shortcomings in many third-party datacentre providers; these being their ability or lack thereof to supply enough

©2013 IDC

power for their clients' needs. Additionally, it should be noted that some respondents may have indicated an inability to scale infrastructure as needed in reference to the service providers' inability to provide the power for newly installed servers and storage.

Figure 3

Factors Influencing End-User Customer Experience

Q. What are the challenges you faced with a third-party datacentre?

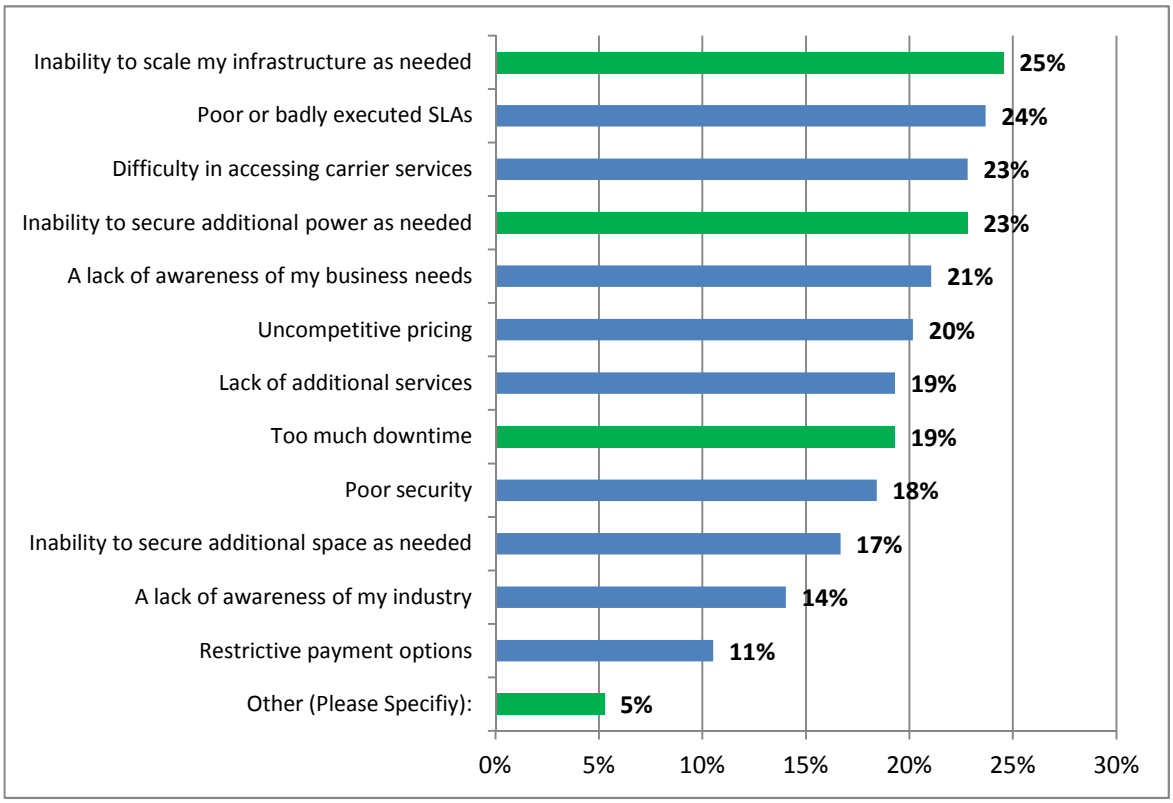

Source: IDC's *2012 Australia Datacentre Study*, (n = 300)

Considerations

The changes that have taken place in the IT industry from both supply and demand perspectives have led IDC to update the advice we give to end-user organisations for the acquisition of datacentre facilities services. In Australia, datacentre facilities services are typically acquired though a tender and RFP process, which necessitates a change in the priority requirements included in the document.

Based on IDC's extensive research on the future of the IT department and utilising data sources from multiple research streams, we recommend inclusion of the following criteria in future tenders and RFPs.

©2013 IDC

Management and DCIM capability

Over the past number of years, the market has experienced marked increases in the cost of electricity. Whilst many factors have contributed to the escalation of operational costs, the issue of energy consumption and the associated costs are certainly on top of the priority list for organisations that operate their own IT infrastructure.

Energy-efficient designs for server and storage technology have been available to the marketplace for a number of years, and most server vendors offer options for lower-voltage CPU and memory configurations for their device offerings. Whilst IDC advocates strongly for the deployment of energy-efficient equipment, we also advocate for a more holistic approach to the reduction of power draw at the IT layer.

Energy management software can help IT organisations gain a better understanding of their datacentre environment, as today many customers admit they have little or no visibility in that space. IDC believes that when customers make energy management a key component of the datacentre management portfolio, it can be utilised to effectively balance the needs of controlling operating expenses and delivering the required SLAs. The measurement, monitoring, and management of energy can enable customers to not only optimise the datacentre by matching IT utilisation to the appropriate power and cooling resources but also serve as a capacity planning tool to effectively manage future change in the environment.

IT availability: The primary driver of energy efficiency

The energy expense of datacentres is most often discussed in terms of the financial impact to companies, yet IDC finds that the primary driver to improve datacentre energy efficiency is to ensure the availability of IT to business.

IT today has become an integral part of the business and its performance is directly linked to business results. In this context, the financial impact of the escalating energy costs are a significant consideration; a lackluster performance of the IT operations can be even more detrimental to the business.

Whilst stated initiatives may include how IT organisations are striving for improved energy efficiency, seeking to reduce or control operating expenses, or even looking to go "green," the IT availability overrides all other factors and will not be compromised to achieve other efficiencies.

Business impact of inefficient power and cooling in datacentres

IDC research reveals the business impact of datacentre inefficiencies pertaining to power and cooling that affect customers. Here are some of the main ones:

- **Increased operational costs.** As previously discussed, the significant increase in the expense to power and cool servers is taxing datacentre budgets. Over the past five years, the customer expense of server power consumption has grown at a rate four times that of new server spending. Server energy is

©2013 IDC

draining a sizable part of customers' budget that could be used for innovation and value-added initiatives. Additionally, IT staff time gets regularly diverted to identify and address issues arising from power and cooling challenges.

- **Unplanned downtime.** Almost half of end-user organisations have reported to IDC that server or storage downtime in their datacentres is a direct result of power and cooling issues. More alarming is that almost a third of the respondents have admitted experiencing datacentre outage. Guaranteeing sufficient power delivery to the datacentre as well as managing the temperatures within the racks and room has therefore become a significant priority.

- **Constrained expansion.** One in five customers have reported that power and cooling issues are constraining the deployment of new servers or hampering new applications and projects. Datacentres that are exceeding, or close to exceeding, their energy thresholds are unable to deploy additional systems and applications to keep pace with the business requirements.

- **Loss of revenue/customer satisfaction.** Today, IT is more closely tied to business results; companies are leveraging IT to drive revenues from new products and services and to gain a competitive advantage in the market. The need for access to real-time data has increased from external (and internal) clients. When IT systems go down, there is a direct impact to revenue from loss of services and a decrease in customer satisfaction from lost access. The situation is exponentially worse when the entire datacentre goes down.

IDC considers energy the key design point for the datacentre — now and in the future. As discussed, it is frequently the case that energy has become the primary limiting factor for financial budgets and the expansion of IT capacity. Having sufficient power and adequate cooling determines where systems can be deployed within the datacentre and even within the rack. Energy management is a primary factor to efficiently and effectively operating a datacentre by reducing risk from hotspot failures and power overloads, as well as making full use of the often untapped power capacity.

IDC estimates that the energy expense associated with powering and cooling the worldwide server installed base increased 31.2% over the past five years. The scale of this increase is more apparent when placed in the context of the overall server market. In 2009, the server energy expense represented $32.6 billion, whilst the server market generated $43.2 billion.

Conclusion

Future-proofing the datacentre should be a core focal area for end-user organisations. As IT becomes more critical, more dependence will be placed on the datacentre facility. It is vital that appropriate steps are taken to ensure the selected provider can offer a facility that meets the needs of the organisation today, whilst allowing for growth and the inevitable advances of technology.

IT infrastructure such as servers and storage have evolved inexorably into ever denser, hotter, and more power-hungry systems that will continue to place increasing burden on the facilities that support them. This has caused a shift in the selection criteria paradigm where space is no longer the leading requirement for an organisation and placed power as the most important need for today and into the future. The supply of, the ability to scale it, and sourcing it sustainably and environmentally are critical.

With power playing such an important role, efficiency has also become paramount and the ability to monitor, measure, and manage the equipment from a facility, row and rack level are now mandatory best practice requirements for any tender document sourcing datacentre facilities. Coupled with the always-on, 24 x 7 nature of today's IT department, this management capability should be accessible in any location, from any device and time.

IDC therefore recommends the following inclusions into a tender document:

- **Datacentre infrastructure management software:** This needs to be built into the facility design and offered by the facility operator as part of the service and not be acquired through an add-on. The facility provider can tie into the system core capabilities such as HVAC and power distribution. The level of granularity of the information is paramount to its effectiveness, and IDC recommends that the tender document specify data is available down to the rack level.

- **Mobility:** With the move towards mobility, IDC recommends that management capability be accessible by smart handheld devices such as smartphones and media tablets.

- **High levels of power availability:** The key metric in this section of the tender should be either power per rack or power per square metre. When calculating this figure, considerations must be placed on the future energy required, not the current energy needed. If you operate a heterogeneous environment that includes non-x86 and x86 equipment, consider the power that would be needed if a wholesale migration took place from legacy to scale out server infrastructure.

- **New workloads:** Future work load requirements, especially disruptive high-performance workloads such as big data and analytics which place high demands on server and storage performance, need to be factored from a scalability standpoint of the datacentre facility.

- **Floor weight loading:** With data revealing that the weight of racks is continuing to increase, it is important to include language within the tender that states a rack limit that will cater to new dense computing solutions where populated racks can be up to 1,200kg.

- **Sustainability:** Environmentally sustainable power supply options should be on the facility provider's radar. Currently, the majority of Australian power is supplied by coal fire power stations. Whilst coal is a reliable and cheap form of power, the

©2013 IDC

market is beginning to embrace alternative energy sources such as wind and solar. As such, it is prudent to future-proof options in the supply of energy to the business by ensuring the datacentre chosen to supply your organisation at least offers alternative energy offerings.

With a large proportion of Australian organisations rationalising and centralising their facilities, attention must be paid to business continuity and disaster recovery. Centralisation strategies provide better operating costs through economies of scale; however, risk must always be mitigated appropriately through careful site selection in line with the needs of organisational requirements. With this in mind, IDC recommends that organsations consider datacentre facilities providers with multiple geographic locations dispersed across broad regions to avoid downtime in the event of a disaster — natural or otherwise.

It's important to note that apart from the above-mentioned criteria, the fundamental basic requirements (for example, floor space, geographic location, SLAs, security, uptime record, competitive pricing, access to carrier services, and so forth) must also be addressed in the tender document.

C O P Y R I G H T N O T I C E

The analyst opinion, analysis, and research results presented in this IDC Executive Brief are drawn directly from the more detailed studies published in IDC Continuous Intelligence Services. Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from IDC. Contact IDC Go-to-Market Services at gms@idc.com or the GMS information line at 508-988-7610 to request permission to quote or source IDC or for more information on IDC Executive Briefs. Visit www.idc.com to learn more about IDC subscription and consulting services or www.idc.com/gms to learn more about IDC Go-to-Market Services.

Copyright 2013 IDC. Reproduction is forbidden unless authorized.