A Practical Approach to Sustainable Hosting

01.09.2009

Changing Business Context

In recent years, heightened awareness of online presence means that Hosting and Data Centres have moved from an afterthought to front of mind for businesses - driven by the perfect storm of the squeeze on real-estate, higher server densities and availability of power. Furthermore, the transition towards a carbon-aware economic framework, facilitated by compliance to NGER and CPRS will fundamentally change the way Australian businesses operate. Businesses who fail to take carbon into consideration will find it difficult to navigate the future business environment.

One of the most visible effects of the impending changes is the uncertainty surrounding energy supply and costs. From a risk management and operational efficiency perspective, as a hosting provider who operate our own Data Centre, Macquarie Hosting has a vested interest in committing our Data Centre (Intellicentre) to sustainable design practices. Our approach seeks to balance sustainability concerns without sacrificing the stringent performance and availability criteria demanded by mission-critical application hosting customers.

Macquarie Hosting started on this journey in 2006. This whitepaper details the practical measures that have been put in place at the "Intellicentre" to reduce our environmental footprint and begin the journey on the path to sustainability.

A holistic approach

We accepted early on that there was no one broad-brushed "silver-bullet" solution to achieve sustainability in a Hosting context. Achieving sustainability outcomes is the result of interaction between efficiencies on many different areas, relying on continuous improvement over time, to ensure we are at the forefront of sustainable Data Centre practices.

We call this philosophy "*Sustainable Hosting*" – which is the combination of best-practice *Green Data Centre* infrastructure and *leveraging technology* on a hardware level. Green Data Centre infrastructure involves integrating best-practice Data Centre operations on a physical level with economies of scale and security inherent in the Data Centre design and certifications. Our differentiator lies in the application of innovative technology on top of the efficiencies achieved in the Data Centre infrastructure. One such example of this leverage is our aggressive use of virtualisation.

The "plan of attack"

The energy budget of a Data Centre can be roughly divided into energy consumed by IT Equipment (demand) and the energy consumed by the Data Centre infrastructure (support). Figure 1 illustrates this simplified energy flow in a Data Centre. Notice that a reduction or energy efficiencies on the demand side will provide a considerable saving on the total energy budget due to reduced demand itself AND the corresponding decrease in the support infrastructure needed to support this decreased demand load. The savings add up – typically a 1 Watt reduction in demand will yield an equivalent amount of saving on the support infrastructure.

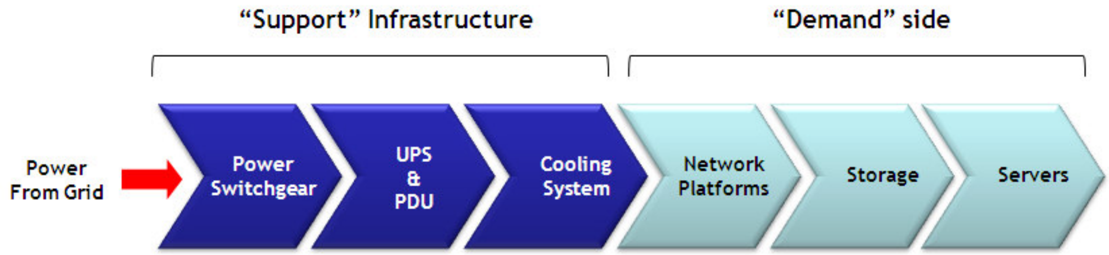

Figure 1 Data Centre Energy Flow

From this energy budget analysis, the main priorities and the largest opportunities for improvements were identified as follows:

- Perform energy audit
- Revise cooling strategy
- Operating more efficient hardware
- Leverage technology to multiply physical infrastructure efficiencies

Energy Audit

The adage "you can't control what you can't measure" rings true, before we can seek to affect any change we had to establish a baseline by carrying out an energy "audit". The energy "audit" is a crucial first step because it enabled us to understand what is consuming the energy. One of the first steps we took was to install per rack power monitoring systems. This ensured we knew exactly how much power each rack was drawing.

In order to get a more granular handle of energy consumption in the Data Centre, we are also looking at more detailed power monitoring measures which provide readings at the Blade enclosure, Per Blade and even Resource pool level.

Having holistic perspectives means more than implementing efficiency initiatives on the "support" side but also looking at what we can do on the "demand" side. To this end, we are looking to change behaviours by providing information on the actual customer energy consumption. The transparency will help implement more effective "demand-side management".

Revised Cooling Strategy

From the energy audit we knew that A/C was one of the big hitters in terms of energy usage. We examined three components of the A/C system separately – cabinets, return air temperature and underfloor airflow. The components were then combined to implement the revised A/C strategy. We turned the existing approach on its head by directly cooling the hardware and in-rack customer equipment rather than cooling the room. This meant channelling the available cooling capacity to the racks themselves instead of a general diffusion of cooling capacity into the room cavity.

Increasing Return Temperature

The previous convention was to cool the Data Centre to a temperature between 17° - 20° C. This meant that the Data Centre is run relatively "cold" from an ambient temperature perspective and a lot of energy is required to maintain this temperature at these low levels. Most of the equipment can be operated safely at slightly higher temperatures.

While this may seem insignificant, small changes in temperature can make a big difference to the energy consumed for cooling. One of the first decisions was to investigate the possibility of increasing ambient temperature. After careful monitoring and controlled testing, the return temperature was increased to 22° C - without adversely affecting equipment function.

Deploying Closed Cabinets

The next area to come under the microscope was how the cold air was actually circulated to the equipment. We worked closely with our rack supplier to re-examine how to re-engineer this.
The concept was to create a closed-loop cooling environment to directly focus cold air on the equipment and to reduce any cooling energy loss by diffusion into the room cavity. This concept relies on blowing pressurised cold air from the CRAC units to the equipment racks.
New cabinets with a closed front and open rear were introduced. The closed rack front creates a plenum to channel the cold air from the underfloor space. Hot air is then expelled via the rear of the cabinet which is open into the room cavity. Note that this expelled hot air is prevented from circulating to other racks by the closed rack front doors.

Optimising airflow

A key component of the closed-loop A/C strategy is maintaining pressure of the cold-air in the underfloor area. Pressurised cold-air from the underfloor space ensures the correct directional flow of cold-air through the equipment. This meant ensuring the airflow path between the CRAC and equipment racks is contiguous by eliminating obstructions and ensuring there are no gaps in floor-tiles at unoccupied rack locations using blanking panels to ensure high air-pressure is maintained. This concept was also extended within the racks by tethering cables to the rack risers and cable-tray runs clearing potential air-flow obstructions. While these measures seem insignificant, they ensure high-pressure and maximum airflow for the chilled air which is a prime ingredient in the A/C strategy.

The closed-loop A/C strategy is implemented by combining the components discussed above. A raised-floor design meant we had a good platform to start from but we had to progressively change-out open-racks to the new fully enclosed cabinets to implement this closed cooling loop concept. Referring to Figure 2, the closed cooling loop consists of CRAC unit, blowing pressurized cold-air into the underfloor space, channeled through closed racking cabinets and expelled into the room cavity itself (as the return path).

This closed cooling loop system is an elegant solution, without the major reconfiguration needed for arrangements such as hot aisle – cold aisle.

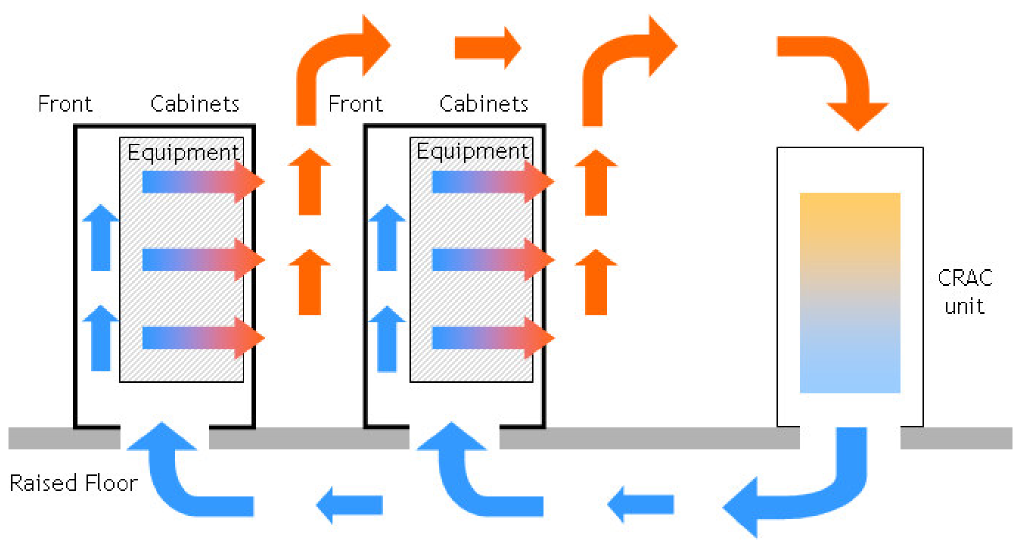

Figure 2 Closed-Loop A/C System

Deploying Latest Hardware

While cooling provides the supporting working environment of a Data Centre, the hardware represents the actual working components themselves and is typically the other major consumer of energy in a Data Centre. In fact energy consumption, rather than space restrictions, of hardware and in-rack equipment is currently the major cause of capacity issues at most Data Centres – a direct consequence of server miniaturisation and higher in-rack equipment densities.

Newer generation hardware has the advantage of incorporating the latest generation power saving and thermal dispersion measures. We systematically review Data Centre infrastructure hardware to investigate if there are further opportunities to change-out less energy efficient hardware.
Additionally, we worked with our hardware vendors to test the hardware at different operating loads to discover the optimum energy configuration and ensure we were operating our hardware at the optimum energy consumption band (for more detail please refer to the following section).

Leveraging Technology

Virtualisation is a key technology which reduces carbon footprint on multiple fronts – less waste due to fewer devices, better energy efficiency, and higher densities (maximizing real-estate). Virtualisation allows consolidation of multiple servers onto fewer physical servers, thereby driving the server load into the efficient zone. The Virtualisation layer allows different Virtual Machines (logically distinct server instances) to share physical server hardware. In this way, the utilisation of resources is actually multiplied several times due to the interaction of two effects – consolidation (running multiple virtual servers on the one physical server, so less physical servers are actually needed) and optimizing server resources (running the server more efficiently at higher loads). Because the typical server energy vs. load curve is non-linear, a typical server will use around 50% of total peak utilisation power even when idle.

We can use a hypothetical example and Figure 3 to illustrate the energy saving benefits of virtualisation. As detailed in Figure 3, the efficiency of servers actually increases with increasing load. In fact, typical non-virtualised servers are running at around 10% of utilisation (which equates to around 20% efficiency) – Point A on the blue Power curve. If we virtualise this server and load the server with 5 VMs, the operating point will move up to Point B on the blue Power Curve.
In moving the operating point from Point A to Point B, we have increased the server efficiency by over 300% (comparing the Red bars) but only increasing average power consumption by 28W (or 115% relative to Point A).
Instead of operating 5 physical servers consuming a total of 5 x 177 = 885W, with virtualisation we consume merely 205W. Less power consumed also means less heat produced hence less cooling required.

Reduced energy consumption and improved server utilisation is only part of the benefits of virtualisation – for a more detailed treatment of the other benefits of virtualisation refer to the companion article “Making Virtualisation work for your Business”.

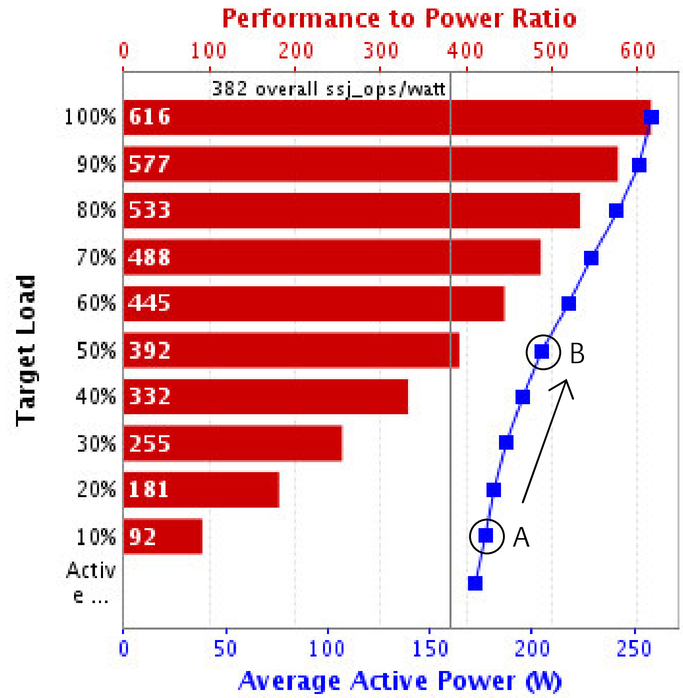

Figure 3 Example Server Performance to Power Ratio Characteristics

SPECpower_ssj® benchmark for HP ProLiant DL380 G5 (3.00GHz, Intel Xeon 5160 processor, Quad-core), 8 April 2009.
SPEC® and the benchmark name SPECpower_ssj® are registered trademarks of the Standard Performance Evaluation Corporation. For the latest SPECpower_ssj® benchmark results, visit http://www.spec.org/power_ssj2008/results/power_ssj2008.html.

Maximising Economies of Scale

For many businesses, a large part of the decision to use an external Data Centre lies in the ability to access carrier-grade physical infrastructure such as power feeds from dual-power grids, redundant CRACs, redundant power sources (UPS/batteries) to ensure performance and security for online-critical hosting ecosystems.
The Data Centre environment fully exploits economies-of-scale effects on the energy (server power and cooling) consumption front, by further multiplying the efficiencies achieved through virtualisation on an individual customer level – analogous to individual commuters taking public transport instead of everyone jumping in their own cars.
To this end, we are pursuing the following initiatives to further maximise the economies-of-scale effect by improving floor space utilisation through higher-density layouts:

> Unless there is a specific requirement, all equipment is installed in uniform racks instead of caged-areas. This also increases the overall density of the layout.

> Having a raised floor system means that the available floor space is fully utilized for equipment racks (in some flat-floor systems useful space is used up with in-suite distributed CRAC units).

Tangible Results

Nearly 2 years from the inception of our sustainable hosting platform, total Data Centre energy consumption is 15% lower, despite organic and incremental customer growth. A conservative estimate of the running total reduction in energy consumption is equivalent to 3,600 metric tonnes of CO_2-e.

For comparison, this is equivalent to running a fleet of 600 cars (15,000km per vehicle over the period of one year) or operating 600 servers (400W servers 24x7 over the period of one year).

The initiatives implemented also optimised the energy budget by reallocating energy used for A/C to powering IT hardware. Before we embarked on this program, 43% of the energy consumed in the Data Center was used for cooling purposes. The revised cooling strategy has reduced the amount of energy consumed for cooling purposes by 25%.

Conclusion

Together the initiatives detailed in this whitepaper have allowed us to reduce our energy consumption profile despite an increase in customer number trajectory. Our "360 GREEN" Hosting platform has modulated the Energy Proportionality equation from a linear relationship to a more sustainable relationship.

Carbon awareness is changing the context businesses are operating in. Macquarie Hosting is committed to make Hosting part of the solution by improving the "carbon productivity" at our Data Centre.

We understand that Data Centre sustainability is not a one-off exercise and requires constant vigilance and fine-tuning. The initiatives detailed in this whitepaper represent only the first step in our journey.

In the broader scheme of things, reducing energy consumption is one of the "low-hanging fruits" of sustainability. As we move forward, sustainability improvements will come from making more and more granular adjustments to existing services and adopting an end-to-end lifecycle framework (i.e. cradle to grave) when designing new services.

About Macquarie Hosting

A division of Macquarie Telecom; Macquarie Hosting is the Australian leader in mission-critical application hosting for companies who rely on their websites for their business. Macquarie Hosting owns and operates Australia's most highly accredited Data Centre, the Intellicentre, and has the people and processes in place to provide the highest levels of security and uptime. Macquarie Hosting enables greater flexibility, agility and on-demand scalability for organisations to respond to spikes and increases in web-site performance.

To find out more www.macquariehosting.com