HHAngus

INDUSTRY PERSPECTIVES

DATA CENTRES

Data Centre Efficiency – Why C-level Execs are the Problem and the Solution

Improving Data Centre Efficiency

The Future of the Data Centre is DCIM

How Much is Lost Capacity Costing You?

Photo Courtesy Uptime Institute

The theme of Uptime Institute Symposium 2012 was Digital Infrastructure Convergence, with special focus on modular data centres and DCIM.

What was clear was that business, IT and data centre managers will need to collaborate in order to meet expanding and evolving IT requirements.

Rod Mons, P.Eng., P.E., RCDD, LEED® AP BD+C
Director, Technology Division, HH Angus

Throughout the year, engineers, designers and project managers from HH Angus' Technology Division attend numerous conferences and symposia, gleaning pertinent information about new and emerging technologies and industry trends to benefit our clients and to inform our own designs.

This year's Uptime Institute Symposium provided a number of topics that we hope will be of special interest to you. From all around the globe, data centre developers, consultants, vendors and industry journalists gathered in Silicon Valley for the seventh annual Uptime Institute Symposium sponsored by its parent organization, The 451 Institute. Approximately 1700 professionals took in the 4-day event in Santa Clara, CA.

The theme was Digital Infrastructure Convergence and the rise of modular, prefabricated and containerized data centres. HH Angus and Associates is pleased to present an encapsulated look at several of the Symposium presentations which we feel have broad industry application.

We'd like to thank the event speakers for generously allowing us to share their findings.

Rod

HH Angus and Associates' Technology Division engineers systems to provide continuous service and near-zero probability of failure, delivering innovative and dependable technology solutions wherever total system reliability is a must. This includes such highly serviced buildings as data centres, hydro control facilities, laboratory and research buildings, telecommunications systems and transit systems, among others.

Features

Data Centre Efficiency – Why C-level Execs are the Problem and the Solution

All data centres are not the same, and there are different types of computing with different missions. Before you go 'duh', keep reading.

Ken Brill, Founder of the Uptime Institute, and Dr. Jon Koomey, author and Project Scientist at Lawrence Berkeley National Laboratory, take an interesting view of the main challenge in increasing data centre efficiency. The key points in their presentation, 'What's Next in Data Centre Efficiency", may surprise you.

Shooting from the lip, Koomey and Brill kicked off their presentation with the bold statement that you are never going to achieve efficiency in your data centre unless you have solid management buy-in, because the problems in doing so are far more about people than they are about technology. They contend that C-level executives need to be focused on total costs and changing the incentive structure and management structure so that people are looking at the total cost of making decisions about the data centre. The solution, they say, is breaking down silos to make sure that I.T. facilities and real estate managers are making decisions collaboratively.

Brill pointed out that all data centres are not the same, and there are different types of computing with different missions. Before you go 'duh', keep reading. This goes to operating expenses (OPEX) and capital expenses (CAPEX) and gets more interesting from there. He identifies three types of computing, each with its own unique OPEX, CAPEX and reliability requirements: "Historically, the first one was scientific computing – this goes back to the Second World War – trying to calculate the trajectory of the gun and the targets some distance away. And it was also related to trying to break codes. Today we would call that 'modeling'. That's a unique type of computing. It has its own OPEX/CAPEX. It's fundamentally different from other things. Then in the '60s, we started to do business data processing. And today we have a new kind of computing that I call consumer computing...the whole Facebook phenomenon, the apps that we have on our phones are consumer computing. They're not business computing, but it all goes through data centres. Each of these has a different combination of OPEX/CAPEX."

Brill drilled down to one of his key reasons why data centre efficiency and upper management support go hand in hand: comatose servers. "We postulated that there might be 10% or 15% of comatose servers. Comatose servers are ones that are dead. They're no longer doing anything but they're still plugged in because nobody knows what they are and nobody's willing to unplug them. Now how could that possibly be?"

Brill laid the responsibility directly at the feet of management in the form of mandating cutbacks to asset management staff. As jobs disappeared in data centres, no one was left to manage the assets. As a result, no one knows to whom the assets belong, so they keep running, year in and year out: "If you turn over 30 of your servers of every year and 15% of the servers for one reason or another don't get removed, after four or five years you've got a significant number of servers that are still installed."

Brill estimates that turning off 20,000 servers equates to 6,000kw. "That's a big data centre. That's a huge data centre. If we were to go build that data centre, we would spend approximately $120 million. Turning off what we don't need is the greenest solution you could possibly have. Then, in addition to your CAPEX, there's OPEX - $10 million per year in electricity saved, because you don't have these servers running. And then there's license savings, and in some cases there will be maintenance savings. All from turning off 20,000 servers."

Utilization also came under the microscope. Uptime Institute data estimates that 54% of the server population has a utilization rate of less than 0.6%. 24% of the population had an average utilization of 4.8% and another 22% had a utilization of 27.8%. "That's comatose", Brill asserts. "The only thing going on here is network updates and backup and disaster recovery and whatnot. And the servers that were most utilized, the peak utilization was still only 43%, and that was only for 6.1 hours a day. Clearly these servers are not working most of the day...this is terrible. If we built a manufacturing plant and had this kind of utilization, we would be fired. No doubt about it."

Another revealing statistic was power usage. There was no change. No one's using power saving features. Brill's advice is: "Number one, go out and kill those comatose servers. If you're building a new data centre and you're trying to figure out how big of one to build, the first thing you need to do is get rid of those comatose servers so you don't size in what's already obsolete. And then, of course, you need to consolidate and virtualize. And the next one - this is practically free - buy more energy-efficient power supplies. And turn power-saving features on."

Brill sees digital infrastructure convergence as IT and Facilities coming together. As CIOs face increasing budgetary pressure, this convergence can produce significant savings. "This message is becoming really clear – for large institutions with global presence, putting IT and Facilities in the same area, it's hundreds of millions if not billions of hours over a five-year period. That is the magnitude of savings that can occur when you get the two groups joined together, and you optimize around business mission. In the future, digital infrastructure conversion means that IT and Facilities are going to have to come up with the most optimum OPEX and CAPEX money using a combination of in-source for some stuff, out-source for some stuff, third party, managed services, all these different things. And if we think it's difficult now, it's only going to get worse. Those who do not embrace this are going to be swept away. It's as simple as that. The money is too big to not pay attention."

Dr. Koomey underscored the management problem of misplaced incentives, separate budgets for IT and Facilities, lack of common language for costs and components of costs, all leading to sub-optimal results. "We're afflicted by a management problem that's older than the data centre

industry, and there are ways to fix it, and the way we need to fix it is at the C-level. They need to fix the incentives and the responsibilities."

Anytime you're going to improve efficiency, you're going to gore someone's cow.

The Uptime Institute's first annual Server Roundup contest was won by AOL. Brill explains how high level support at AOL led to their win: "They took three runs at it to achieve those savings. The first two, they didn't have support. And the third time they got support from senior people, and that's what it took. Because anytime you're going to [improve] efficiency, you are going to gore somebody's cow. That is the nature of it; you're going to upset somebody. So you need to have sponsorship at a very high level to push through the bureaucratic resistance. Until you get that, it's not something you can do from the bottom up. It's something that must come from the top down."

According to Brill, efficiency is not a vendor problem: "The vendors, I think, have really done their part. It's the users now that have got to step up. And my message to the users is, if you don't do this, there are others who will, and your job is probably going to disappear. This is transformational. If these costs don't get covered or brought down, our economy will ensure over time that some people go out of business."

Koomey highlighted another example, the "Green Lights" program started by the EPA in the late '80s and the predecessor to Energy Star. "What they did was to get commitments from CEOs of major companies to retrofit their lighting, because the efficiency opportunity for lighting was huge, simple paybacks over one year, one and a half years. So they got a letter from the CEO saying, "We promise to retrofit our facilities up to a three-year payback," and they were able to use that commitment from the CEO to break through the institutional logjam. The enthusiasts who wanted to do something about efficiency, they could wave this letter around and say, "The boss wants it. Get out of my way." So having CEO-level or C-level buy-in is really critical to changing institutional behaviour inside big organizations."

If you'd like to participate in the Uptime Institute's annual Server Roundup contest, here is the link: http://blog.uptimeinstitute.com/2012/08/second-annual-server-roundup-contest-rules/

Improving Data Centre Efficiency

A new proactive approach and maturity in management is now required to plan, increase efficiency and deliver service.

The last decade has been marked with extreme growth and change for Data Centres. A new proactive approach and maturity in management is now required to plan, increase efficiency and deliver application services. The question becomes how do we advance with this rapid change and spend money wisely to address reliability and efficiency?

Clemens Pfeiffer, Chief Technology Officer at Power Assure and 25 year veteran in the software industry, suggested for reliability you "look at the failure rate of the equipment or by people or the hardware, the software, everything you have in your Data Centre, and you look at the maintenance side that goes with it as each influences outages and downtimes." Pfeiffer has also suggested that for efficiency an important first step would be in understanding and integrating IT and facility into "a common dashboard environment" to allow for a correlation between the two main parts contributing to the cost of a data center.

A key indicator of efficiency is the PUE measure (Power Usage Effectiveness) which tells you how much power consumption is allocated to cooling relative to IT. Reducing PUE, meaning power consumption allocated to cooling, effectively shifts available watts over to the IT side and provides more capacity. The key to obtaining efficiency in Pfeiffer's approach is increasing utilization of running servers, effectively allocating and de-allocating servers on the fly using intelligent automation across servers and cooling. This can be further improved using load balancing and virtualization and cross data centers load shedding and shifting.

At the end of the day, the way to look at Data Centre efficiency should be inclusive of IT, facility and application components and should enable effective resource utilization to maximize the efficiency with dynamic adjustments based on variable application demand.

Modular Schmodular?

Do we want to build data centres that are like Lego® blocks, that are all aligned, that are reliant on mythical software to perform remedial actions while our electrical and mechanical systems fail?

Photo Courtesy Uptime Institute

Dan Golding's Keynote Presentation, 'Waving the White Flag on Innovation', (or why modular solutions are not for everyone), was not only well attended, but drew distinct lines in the audience between traditional data centre supporters and modular proponents. Dan Golding is a Senior Data Center and Cloud Leader. His industry background includes roles as VP & GM at RagingWire Data Centers, Managing Director at DH Capital and Vice President and Research Director at Tier 1 Research.

Given the rise of modular data centre solutions in their various forms, Golding issued this challenge: "Are we ready to give up and say the data centre's baked, we've done the best we can, the data centre's never going to get better, it's never going to get more reliable, it's time to put it in a box and clone that box a million times, throw away the designs and do it like this forever? Are we ready to declare that our data centres will be uniform and uniformly mediocre in the years ahead? Do we want to build data centres that are like Lego® blocks, that are all aligned, that are reliant on mythical software to perform remedial actions while our electrical and mechanical systems fail?"

Golding summed up the current buzz about 'modular' this way: "We've heard a lot about how modularity and containers will take care of all our data centre problems and we can fly away in a world of fluffy clouds and unicorns. You've been told this by folks who sell containerized data centre solutions or who purchase containerized data centre solutions, or who are analysts who sell research papers about computerized data centre solutions."

Focusing on MEP, the guts of data centre infrastructure, the 'perfect world' approach is building out your data centre completely on Day One, with clearly planned phased builds before even breaking ground. Golding acknowledged that, to some degree, broad spectrum applicability is, in some ways, modular data centre design. But with phased design, the danger is ending up with a four-phase or four-module building that is, in fact, four separate data centres. (Full disclosure self interest note: This is something that experienced data centre infrastructure designers like HH Angus can help you avoid.) Each phase must enhance the capabilities of the existing raised or slabbed-floor data centre. Golding went on to say: "I think those are things that ten years ago were revolutionary, that were considered 'modularity'. They were things that most folks weren't doing. But now we do them, now we like them, now they're best practice. The issue then becomes that there's a very big leap between those broadly applicable areas of modularity... to move from those to the issue of prefabrication and testing of data centre modules. And then, if we go further to the idea of taking the mechanical and electrical components and stuffing them in some kind of a container, be it an ISO standard shipping container, be it some kind of customized box that you buy or create or build, as we increase the amount of modularity towards total containerization, then the number of data centre users who can gain great advantage from this approach starts to decrease."

What exactly does 'modular' mean?

There's a lot of confusion about what is and is not a modular data centre. For example, if you frame the online search question 'what is a modular data center', you get six million plus hits.

While the confusion isn't quite that dramatic, various stakeholders are often talking about different things when they talk 'modular'. Modularity can mean containerized servers, containerized MEP, phase-built modular MEP and prefab modular MEP. In its essence, modularity is about the efficient deployment of capital, but sometimes it's also about 'I don't have the capital to do what I really want to do, so I'm going with a modular solution because this is all the budget I have'. While that may make your accounting folks happy, that doesn't mean it's a great solution for your data centre needs.

Getting to Reliability

One of the benefits of going with a modular approach is ostensibly pre-selected and uniform equipment. Without this, attempting to re-commission a data centre with dissimilar MEP hardware between phases or modules is difficult. But according to Golding, the key question is really this: does pre-fabrication and pre-testing, and for that matter, containerization, enhance reliability? "That has not personally been my experience. It does catch some out-of-the-box failures, and it does catch some mis-wiring, but it doesn't catch other mis-wiring; it doesn't catch other out-of-the-box failures. There's always some significant risk of damage during transport, or improper installation. Essentially you can, you should, you must re-test everything when it gets to your site no matter how much testing and commissioning it had when it was on the equipment vendor's site. If you put it on your site and you don't re-commission and you don't re-test it, that is the trail of tears." Plus, correct and proper engineering and specification by a trusted and proven MEP supplier provide all the same benefits of pre-fabrication.

Modular or Snowflake?

Proponents of modular data centres argue that data centres shouldn't be like snowflakes, they shouldn't all be different just for the sake of being different. But does the modular approach mean an end to design innovation? Golding posits that adherents of the modular system assume that MEP design has reached its zenith, no further innovation required: "[But] there are better designs. Not all data centres are created equal, and I call upon you, I beg you to reject the idea that all data centres are the same – does everyone drive the same car? Do we all live in houses that look exactly alike? Are all our servers exactly the same? Have we commoditized everything about our industry? That is the risk. If we embrace this idea of extreme modularity, of extreme containerization, innovation goes out the window."

As proof, Golding points to improvements in PUE over the past decade, vastly improved UPSs, and the huge increases in data centre energy efficiency. And more is possible; he suggests there's a lot to do in the area of reliability and control systems.

Golding's key take-away? Modular data centres might be the right solution under the right circumstances but, for many customers, they're aren't and never will be.

Modular Data Centre FAQs

To paraphrase Mark Twain, reports of the death of the ‘snowflake’ – aka the custom designed data centre - are greatly exaggerated.

Q: **How do you define ‘Modular Data Centre’?**

A: There are generally three interpretations of ‘modular’:
a) an ISO shipping container populated with servers and packaged power and cooling,
b) a prefabricated solution that is manufactured off-site for installation in a shell and
c) a data centre that is expandable by modules.

There is an important role for the consulting engineer in all three, although HH Angus Technology Division Project Manager, Scott Brenner, reports that "the pre-built containerized 'SKU' is relevant for certain geographic locations or for organizations with specific performance, cost and scalability requirements but it is not the optimal solution for the majority of enterprise IT deployments. However, we do see benefits for some clients in using prefabricated modules, or using an incremental approach to capacity expansion to maximize ROI."

Q: **How modular should I go? What are some of the considerations in determining if my firm should go modular?**

A – The overarching answer is ‘it depends’, and what it depends on are the answers to the professional needs assessment you should have commissioned before choosing a solution.

The needs assessment will examine considerations such as: capacity planning, time to market, scalability, risk mitigation, capital budget, alignment with business goals, site geography and environmental factors, need for future flexibility, etc.

Q: **Do I need to engage a consulting engineering (CE) firm for a modular data centre project?**

A: The involvement of a consulting engineering firm in the initial stages of development is crucial to understanding the short- and long-term implications of the solutions you are considering. Consulting engineering firms represent the sole and unbiased interest of our clients, and will recommend the most appropriate solution.

We are well positioned to help you determine what configuration will best address your current and future needs. According to HH Angus Principal, Craig Sievenpiper, “Typically our team brings deeper experience than is available to a client internally, and that’s simply because we do more of this work more often across a wider spectrum of industries. Also important is that we’re vendor neutral, so we’re able to multi-spec the design to allow a number of vendors to bid on the installation. Generally, the result is a competitive bid situation that drives down the capital cost for our clients.”

Q: **What services will a Consulting Engineering (CE) firm provide if I choose a modular solution?**

A: A consulting engineering firm can provide valuable assistance throughout all three key stages of your modular data centre:
1) Planning,
2) Construction/ Commissioning and
3) Operations.

In the PLANNING stage, your CE firm can perform options analysis; perform risk assessment; develop a life cycle replacement strategy; and identify total cost of ownership.

During CONSTRUCTION/COMMISSIONING, your CE firm can oversee factory witness testing; ensure cost control; perform quality assessment; conduct performance testing; design IT supporting infrastructure; and conduct performance validation/Commissioning.

During OPERATIONS, the CE firm can provide training; perform retro-commissioning; conduct life cycle planning; suggest preventive maintenance; develop Standard Operating Procedures for seasonal variation; perform expansion planning (optimizing operational efficiency to reflect evolving floor plans); analyze energy use; and analyze growth and impact on building infrastructure.

Q: **Why isn’t everyone going modular?**

A: To paraphrase Mark Twain, reports of the death of the ‘snowflake’ – aka the custom designed data centre - are greatly exaggerated. Many people feel the modular solution poses constraints, that it won’t meet their needs or that it may force them into decisions they’re not ready to make. Where you are on the continuum between ‘total flexibility’ and ‘reduced cost’ also determines whether the modular approach will work for your data centre needs. The modular solution tends to work better for very large or very small facilities or those where time to market is extremely short and suitable locations are available. In any case, an experienced professional consulting engineering firm can help you determine your best solution.

DCIM – The Data Centre Infrastructure Management System – What is it and do you need one?

As the industry moves forward, data centre management will become more and more automated.

Depending on who you talk to, DCIM is variously an asset-management system, a monitoring system, a system for capacity planning and forecasting, or a complete data centre operating system. Andy Lawrence is Research Director and head of The 451 Group's Datacenter Technologies and Eco-Efficient IT research practices. His team examines how IT and business can monitor and reduce energy consumption and carbon emissions. Speaking at the 2012 Uptime Symposium, Lawrence defined DCIM as an information and control system for data centres, an admittedly broad definition: "A DCIM system collects and manages information about a data centre's assets, resource use and operational status. This information is distributed, integrated, analyzed and applied to help managers meet business and service oriented goals and optimize their data centre's performance."

Lawrence defines the main components of DCIM as:

- Monitoring
- Control
- Planning and non-real-time optimization,
- Data management, integration and reporting
- Asset configuration and change management

Lawrence gave an example of optimizing cooling and explained the benefits of DCIM over traditional systems: "People would say 'Well, we can take information from the data centre floor, temperature sensors, link it to the BMS and it can ultimately come up and adjust the cooling accordingly,' and that's been done for a long time. That isn't, in our view, DCIM. It's good to have, but it's not DCIM. DCIM would be able to do that but it would also collect all of that data in a real-time way; it would be able to report that in terms of metrics so you can see what the cooling's doing, you can see what the temperature's doing, you'd be able to do trend reporting, you'd be able to look back, you'd be able to look forward, and then you would also be able to pull that data up and do things like spot trends, exceptions, problems as they might be occurring, and you can embed some policy to instruct the system what to do. So, it's [examining] a lot more data and using a lot more intelligence. It's a lot more difficult to do but you're getting much, much more fine-grained control of the data centre as a result."

Lawrence asks 'so why invest?': "The big issue is over-provisioning or under-provisioning of power cooling, and the ability to monitor, track and model these is a really key reason for investing. Another really simple reason is just to help with the day-to-day management of where to place servers or new equipment, where to provision. You can instantly know what power is available, what cooling is available, reserve space, etc. So it helps day-to-day management, [which is] particularly useful in a colo environment. Then, accounting for energy, fairly obviously. This is really about getting much better control and professional management of the data centre in order to make better decisions and ultimately to avoid financial waste and improve availability."

Lawrence identifies cost as the overwhelming main reason companies do not invest in DCIM. And not because the systems are overpriced, but more likely because this has not been an area where data centres typically have spent invested, so deciding to spend here seems expensive at first. Other barriers to adoption include: the challenge of populating asset databases: integration with existing infrastructure; complexity; and unproven systems.

Photo Courtesy Uptime Institute

As the industry moves forward, data centre management will become more and more automated. Lawrence explains that humans don't need to be involved in the millions of micro-decisions made throughout the course of the year in a data centre. That's what technology is for: "So where is the technology going? It's clear that people want an integrated framework of technology to cover the whole data centre, so we're going to see far fewer point products and people coming up with much more integrated solutions."

Wrapping up his presentation, Lawrence highlighted these as issues to watch if you are considering the purchase of a DCIM system:

- Security – IP addressable devices in the data centre could create access for viruses and hackers
- Vendors/products subject to change – with rapid marketplace evolution, some systems may be quickly outdated and some vendors may merge
- Convergence – could change the financial picture
- The challenge of data volume - more hardware and more IT staff and expertise may be required
- Fragmentation – integration could be an issue as some systems may only be partially open.

Lawrence's research suggests that we can expect to see more discipline in the way that information is picked up from the data centre, how it's stored and tracked, and how change control is managed: "People adopting [DCIM] are going to have to go through some kind of cultural business-process reengineering, because they're going to have to learn to live by the DCIM system.

The Future of the Data Centre is DCIM

Using DCIM software, companies are optimizing hardware through an interface that lets them see at a glance what equipment is running, at what capacity and at what power consumption.

DCIM (Data Centre Infrastructure Management) software technology is changing how we view the data centre. Using DCIM software, companies are optimizing hardware through an interface that lets them see at a glance what equipment is running, at what capacity and at what power consumption. This in turn allows companies to make real-time decisions on their data centre management.

Kevin Malik, CIO of IO and GM of IO Labs, spoke at the 2012 Uptime Institute Symposium on the subject of viewing your data centre from a single interface. Malik oversees network operations, software development, governance, architecture, and IT strategy at IO, and is considered an industry expert in DCIM. He lamented the current lack of DCIM products providing true full spectrum management, noting: "One of the things we've seen with a lot of 'kind of' DCIM tools out there, is they specialize in one area or the other, but they're not covering the broad aspect…customers want us to deliver an entire solution to them. They want to be able to see everything on a single pane of glass. They want to be able to map things. They want to be able to see what's going on with their port and their traffic and how much power it's consuming. Why not show them how much power every virtual machine is consuming…why not show them how their applications are running?"

Malik is excited about the degree of flexibility that a DCIM systems offers. He keyed in on availability of multiple sources of information to enhance the data centre operator's ability to make informed decisions, such as designing modules and designing space that ties the IT workload to availability of services - "Why would I take Tier Three storage and put it in a Tier Four module? It's what we all do, right… say 'Here's some space'? That's how most data centres are [underutilized or overdesigned]. Why not say, 'This module is running Tier Three, let's run it at N'? And if we have to go to [a] generator, and that module's at N, and there's a lower quality of service, let's make sure that our highest priority of workload, maybe our web servers [or] our banking system, continues to stay up." He cites this flexibility as a key advantage of DCIM systems.

Malik agreed with many other presenters at the Symposium on the need to bring Operations and IT teams together in order to drill down to all of the core information required to design the best data centre for your needs. His company's DCIM system is capable of monitoring a quarter of a million points every second, and that creates a lot of data for analytics and simulations.

Looking ahead, Malik asks: "In the next five years, who's going to be using a lot of these tools? It's going to be kids coming out of college; it's going to be people who grew up with smart phones, with tablets, who don't carry laptops. We can't just scatter drawings on a work station; [the system] should be running on your phone. We have all these [devices] available today to make it completely pervasive anywhere on the cellular network through secure means.

Malik adds: "Everything from your workload to fibre channel utilization, to how much fuel you have in your generator, all that's delivered today completely IT-ready. We can't have this stand alone…we should have web services; we should have the IT capability to understand 'what's my average KW per rack across the organization?' Not just because I'm monitoring the grand circuit, but I should be able to slice and dice that data in any way. You should make it completely platform agnostic, portable, SQL, postscript, all of these things should be available through a DCIM tool."

Computational Fluid Dynamics (CFD) Solving your Data Centre cooling problems on your computer

CFD, also known as air flow modelling or air flow simulation, traces its roots back to the 1930s and the Navier-Stokes equations.

Wherever a group of data centre managers gather, air flow and cooling issues are sure to be a hot topic. The 2012 Uptime Symposium was no different. In a breakout session, Dr. Amir Radmehr, PhD, and a member of the technical staff at Innovative Research Inc., spoke on how CFD is increasingly being applied to improve the cooling performance of data centres.

CFD, also known as air flow modelling or air flow simulation, traces its roots back to the 1930s and the Navier-Stokes equations. And if you want to geek out for a moment, these arise from applying Newton's second law to the motion of fluid flow. But getting back to data centres, using CFD you can create a computer model of your data centre, representing all its relevant elements, such as cooling units, server racks, , obstructions, etc. CFD software then calculates air velocity, pressure, and temperature and provides you with a dynamic image of the movement of air in your data centre, identifying hot and cold spots and highlighting critical situations that require attention.

Armed with this information, facilities managers can more easily identify the cause of a cooling problem and ensure that servers and equipment in the data centre receive appropriate cooling. But, according to Dr. Radmehr, "at the same time you want to save on energy consumption and operating costs of the data centre. Using [CFD] technology you can explore different ideas and solutions without risking the operation of the data centre, and find the best solution for your data centre problems."

CFD modelling works well for both designing new data centres and finding remedies for existing data centres. In the case of data centre expansion, CFD provides the analysis to help you make decisions about how best to integrate new equipment.

Tom Doyle, Senior Mechanical Designer at HH Angus and Associates, finds CFD is a particularly helpful tool for IT managers to model projected IT loading on the floor to determine if the proposed configuration will create any heating or cooling issues. "We've been using CFD with several of our major financial institution clients to analyze energy saving options, such as cold aisle containment, higher supply air temperatures, and variable volume air flow.

"We are now able to use motion simulation to illustrate to their non-technical decision makers what problems and risks exist, in both still images and air flow animation. That brings tremendous value to the IT Manager or Facilities Manager, who can use our analysis to demonstrate projected improvement in performance and energy savings. For example, we recently completed a CFD project in which we were able to slow the air flow and raise the supply air temperature, while still maintaining optimal temperature at the top of the cabinet. It would have been difficult to convince the client of the benefits of our solution without CFD modelling of the cold aisle containment and the value of the energy savings; because they could see it, they totally supported our design solution."

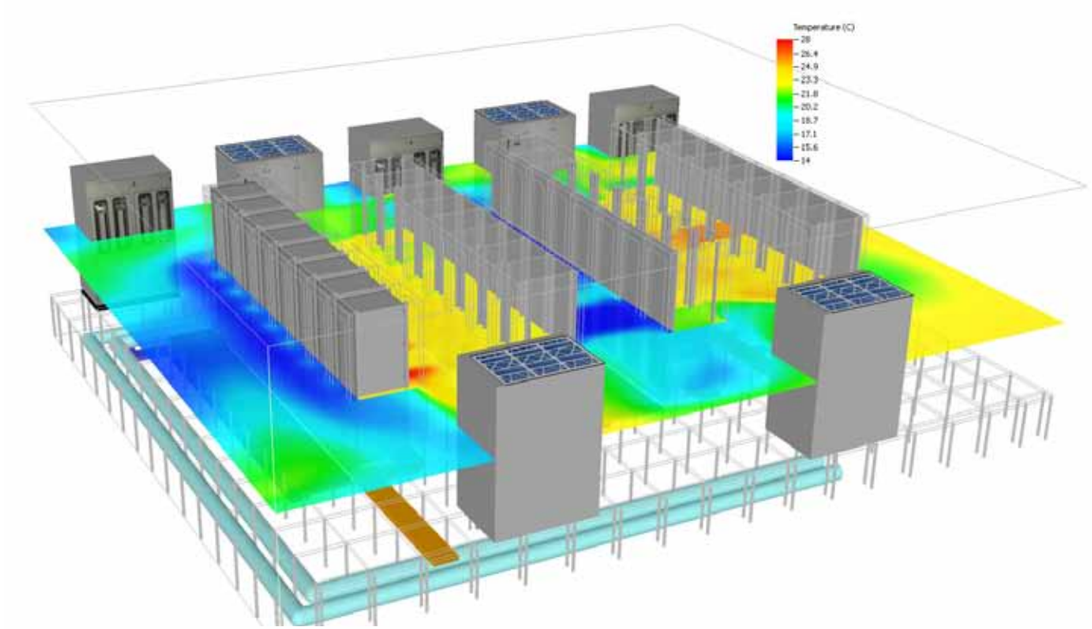

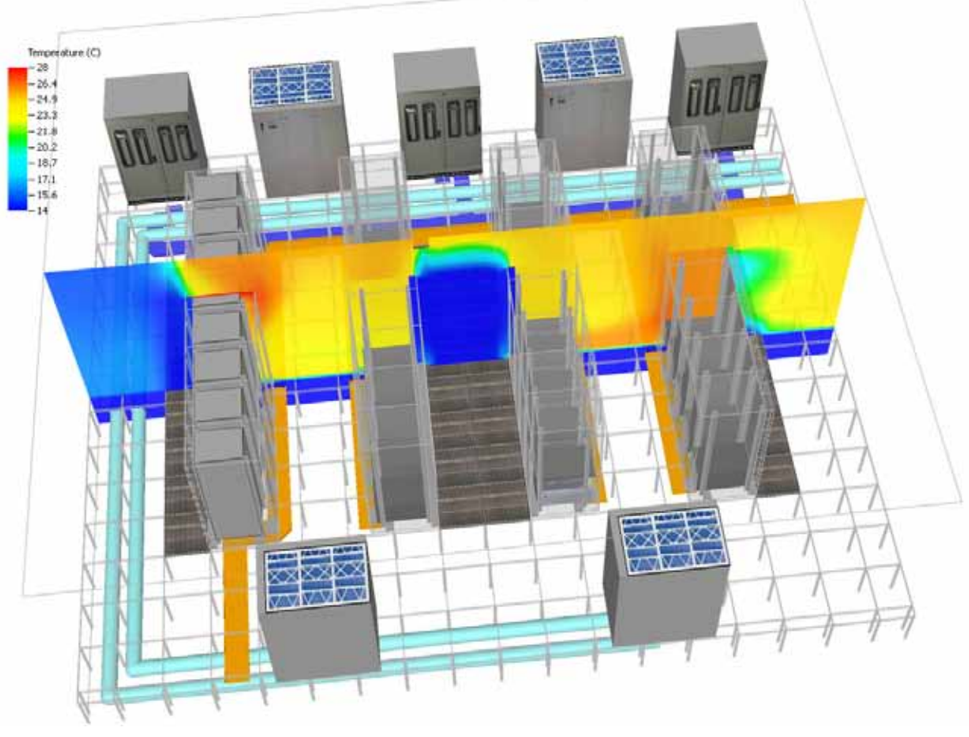

How Much is Lost Capacity Costing You?

Lost capacity exists for every data centre. The key is to know what that number is and to plan to mitigate it as you move ahead from one deployment to the other.

The goal of any data centre owner/operator in building a new data centre is to plan for ideal deployment of space, power, cooling and networking resources. According to Akhil Dokka, Senior Engineer at Future Facilities, we don't live in some fantasy world where such a thing could happen, and the harsh reality is that data centre operation always deviates from intent. In other words, the planned resources and the MEP to support them for Day One often depart dramatically from what your IT department ends up executing. According to Dokka, the elephant in the room is lost capacity. "Put simply, the assumption for any data centre, either new or operational, is that there's *x* amount of power that's being used and *y* amount of power that's available for deployment. But in reality, lost capacity is like a new car you just drove off the lot. You always lose money on it. Lost capacity exists for every data centre. The key is to know what that number is and to plan to mitigate it as you move ahead from one deployment to the other."

Dokka offered a telling example of why it's so critical to understand the parameters of your lost capacity: "Let's assume that you have three data centres at 60% capacity. They're all running at that level and you have to increase the capacity of your operations. So option one is to say, "Let's go build a new data centre. But, if you don't know what your lost capacity is and you don't change your business practices, then you end up again at 60% utilization, which means that you just added a new data centre that works as inefficiently as the three you had before." Ouch.

Dokka homed in on other aspects of lost capacity – power fragmentation and cooling fragmentation. As you begin to deploy more equipment, you may not have enough available power on the preferred rack, so you deploy to another one. "This is very typical fragmentation from a power standpoint. So now you have to plan for a new cabinet that can provide 1,400 watts or greater, and suddenly you have lost capacity in space and also lost capacity from the power standpoint. So this is a typical example of what happens in an everyday data centre."

Cooling fragmentation comes into play because the assumption is that if power, space and network are available, then cooling is also automatically available which, according to Dokka, is a common fallacy: "Most of the IT equipment is not designed the same way. They have different air flow requirements; they have different power requirements. What results is a state where the initial design was fine, but as IT loads up their remaining cabinets, it starts to fragment the cooling capacity."

Part of the challenge of managing capacity is that IT departments are implementing equipment today that didn't exist as little as five years ago, so nimble asset management is key to having a consolidated data centre, for moving servers around and redistributing your loads in terms of power, cooling and networking. Dokka advises looking for the lesser of the evils: "The main take away is to look at short-term deployment. Obviously IT has a lot of say [about] where the deployment will go, but there's a way to say that if you do it one way, the lost capacity is going to be 30%, versus 10% lost capacity if you do it another way…it goes back to design, but each of them has to be simulated before you're in deployment."

Uptime Institute's 2012 Data Centre Industry Survey

Despite the current stresses on the economy, companies are still spending on data centres.

Results from the Uptime Institute's 2012 Industry Survey were presented at the Symposium by Matt Stansberry, the Institute's Director of Content and Publications. Over 1100 data centre end users responded to the survey and respondent demographics indicated 45% Facilities Managers, 36% IT Managers and 19% senior or C-level executives from around the world. Fifty percent of respondents were from North America. More than 75% manage more than one data centre, and represent large financial organizations, technology service providers, government, or manufacturing. The survey revealed some interesting information for data centre users and providers:

- One-third of respondents are experiencing a 10%+ increase in budgets. Despite the current stresses on the economy, companies are still spending on data centres.

  More than 80% of respondents have built a new data centre or upgraded their data centre within the last five years. According to Stansberry, "Despite continued growth, some people are still running out of data centre capacity, but slightly less than shown in last year's survey results; in 2011, 35% said they were going to run out of capacity. Approximately one third of respondents said they're running out of power, cooling, or space in their existing data centres."

- Private cloud use is increasing. Nearly half of respondents are already deploying what that they term 'private cloud'. Public cloud use is also increasing. Data from July 2010 indicates cloud computing adoption was about 11%; by April 2012, it was reported at almost 30%.

  Stansberry raised a salient point: "Why are people going to the cloud? That's a good question. Apparently, to save money. And I want to be clear: the Uptime Institute isn't saying that cloud computing will necessarily save you any money. We're just saying that you think it will. Whether or not it's less expensive on your total cost to put a workload in the cloud is one of the fundamental reasons that we started [our] FORCSS program…we're trying to come up with a methodology to help people have a standardized, repeatable format to compare internal IT costs to external providers."

  The Institute's research shows that there are still important concerns that are keeping companies from deploying to the cloud, and these include security concerns, compliance and regulatory concerns and lack of credible case studies.

- Green IT/Data centre efficiency. More than 50% of respondents identified energy efficiency as a priority. Stansberry: "No one has their final financial incentives optimized to effect real change. Data centre facilities management executives really led the first wave into the data centre efficiency question because they paid the power bill." Seventy-one per cent of respondents reported that their Facilities/corporate real estate department pays the power utility bill, with only 20% reporting that IT foots the bill.

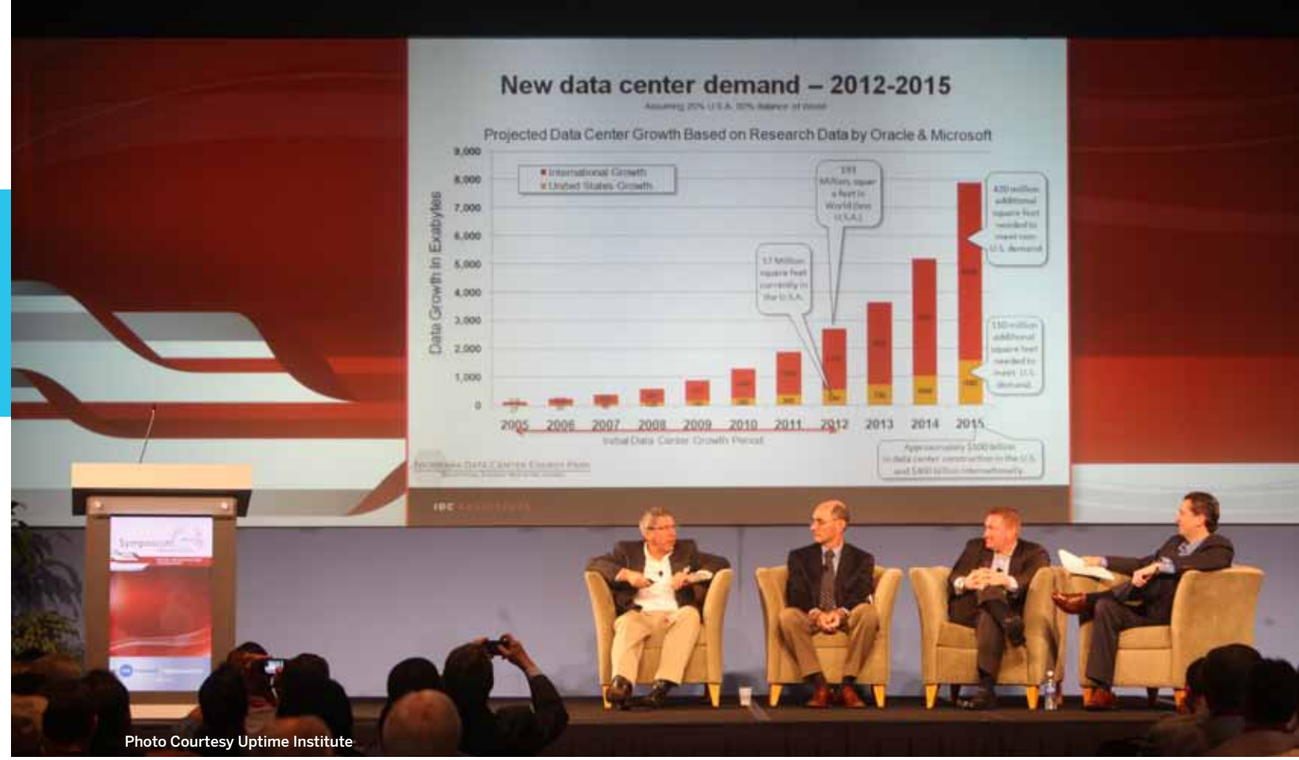

Photo Courtesy Uptime Institute

Stansberry: "Future improvements and data centre efficiency are really going to depend on incentivizing IT practitioners to take the next steps. So we need IT organizations to take a systematic approach at the application layer, consolidating apps, consolidating service, de-duplicating data, removing comatose servers, building the redundancy into the IT architecture rather than into physical systems." Approximately half of survey respondents reported pursuing some green or environmental certification in their data centres, such as LEED and ENERGY STAR.

- PUE measurement is getting more precise. Survey results show a large percentage of organizations no longer in the PUE Category 0, the simplest, least exact method, and more moving into Categories 1, 2 & 3, with the largest jumps in Categories 2 and 3. Stansberry reports that "a large percentage of respondents kept inlet air temps below 65 degrees. This year it's down 6 percent. And a much higher percentage are managing inlet air temps at 71 to 75 degrees, but you can't raise inlet air temp in a vacuum. This isn't something you can just go in and turn up the thermostat. There are multiple interlocked strategies for improving your data centre efficiency. In order to raise your inlet air temp, you need to have a detailed power monitoring [and] cooling monitoring benchmarking system. You need to have multiple ways of tracking to make sure that you're doing the right thing, [that] you're not putting your IT systems at risk."

- Data Centre Infrastructure Management (DCIM). Many organizations are acquiring or have implemented tools such as real time environmental monitoring and alarming; real time power monitoring at device or circuit level; trending analysis of operational data; capacity planning; and tracking asset inventory. According to Stansberry, data centre infrastructure management software is about making better decisions about future capacity planning requirements: "This market is the Wild West. [While] vendors are coming up with new bells and whistles, the end users are scratching their heads, and [asking] 'Do I need that product? Do I have the analysts to keep this thing up to date? What exactly do I need?' And they wind up going back to Post-it notes and Excel spreadsheets, because it's not just a superficial purchase. These things are expensive and cost is one of the primary barriers. You need to make the right decision that you're not buying shelf ware. And the difficulty of implementation is another significant barrier. These tools are going to be the decision engines of your organization, and it's important that you make the right choice."

You can see the Uptime Institute's full survey results and analysis at: http://uptimeinstitute.com/2012-survey-results

HHAngus

1127 Leslie Street
Toronto, Ontario M3C 2J6
Canada
T 416 443 8200
F 416 443 8290
Toll Free 866 955 8201

507 Place D'Armes, 3e étage,
Montréal, Quebec, H2Y 2W8
Canada
T 514 394 1440
F 514 205 6441

Suite 600, 3333 Lee Parkway
Dallas, Texas 75219
USA
T 214 526 0481
F 214 526 5660

Suite 806, 405 North Wabash Avenue
Chicago, Illinois 60611
USA
T 312 527 5552
F 312 527 9061

If you have any feedback on the content of this booklet or would like to comment generally, we'd enjoy hearing from you at info@hhangus.com.

hhangus.com