ANIXTER®

GLOBAL TECHNOLOGY BRIEFING

BUILDING BLOCKS FOR DATA CENTRE INTEROPERABILITY BEST PRACTICES

RISK MANAGEMENT

NETWORK MIGRATION

POWER OPTIMISATION

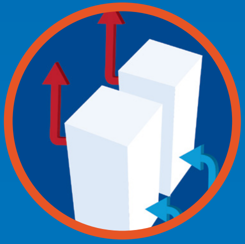

THERMAL EFFICIENCY

DCIM ENABLEMENT

Infrastructure as a Platform
by Anixter

How do you define best practices in your data centre?

Introducing Infrastructure as a Platform by Anixter for an agile, flexible and scalable data centre.

Data centres are continually evolving to keep up with the growing needs for capacity, performance and uptime. At the same time, managers of data centres are also anticipating future needs and the impact those needs have on budgets. An agile data centre allows you to plan and build with interoperable technologies that fulfill current demands, as well as scale as needed to meet future requirements.

Infrastructure as a Platform addresses the key building blocks for data centre interoperability that can provide agility for budgets, scalability for demand and flexibility for technology choices.

This approach not only addresses the five key technology areas, but it also integrates innovative solutions to meet your assessment and deployment needs.

Anixter's site-specific deployment solutions allow you to more accurately plan projects and improve scheduling, reducing non-productive labour and on-site assembly challenges.

Data Centre Infrastructure Challenges

RISK MANAGEMENT: Creating a data centre security solution isn't a one-size-fits-all proposition. Anixter has defined a **six layer approach to physical security** based on the current TIA-942 standard that addresses how to mitigate potential risks that can occur from the site perimeter to the cabinet.

NETWORK MIGRATION: Interoperability is the foundation of a converged infrastructure platform that allows for better return on IT investments through upgrade flexibility and enhanced system longevity. To maximise interoperability, **high performance structured cabling** is critical to support a migration path for 10/40/100 gigabit network demands.

POWER OPTIMISATION: When evaluating power distribution configurations, multiple factors such as efficiency, reliability, equipment availability, safety and cost must be considered. Developing an **intelligent power chain** that starts at the grid and flows through the IT cabinet will make sure those considerations are met.

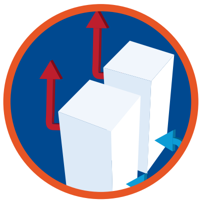

THERMAL EFFICIENCY: Selecting the appropriate cooling system for your environment is essential to manage airflow to and from the IT equipment. Proactive monitoring will help to balance cooling with the demands of the IT load, improving efficiency, reducing costs and moving toward a state of **conditional environmental control.**

DCIM ENABLEMENT: Intelligence gathered from data centre infrastructure management (DCIM), along with a process to act on that knowledge, is vital to understanding the who, what, where, when and how of your data centre. **The five senses of DCIM** addresses common data centre challenges to help achieve a faster return on investment.

To learn more about Anixter's approach to solve these challenges, visit **anixter.com/datacentre.**

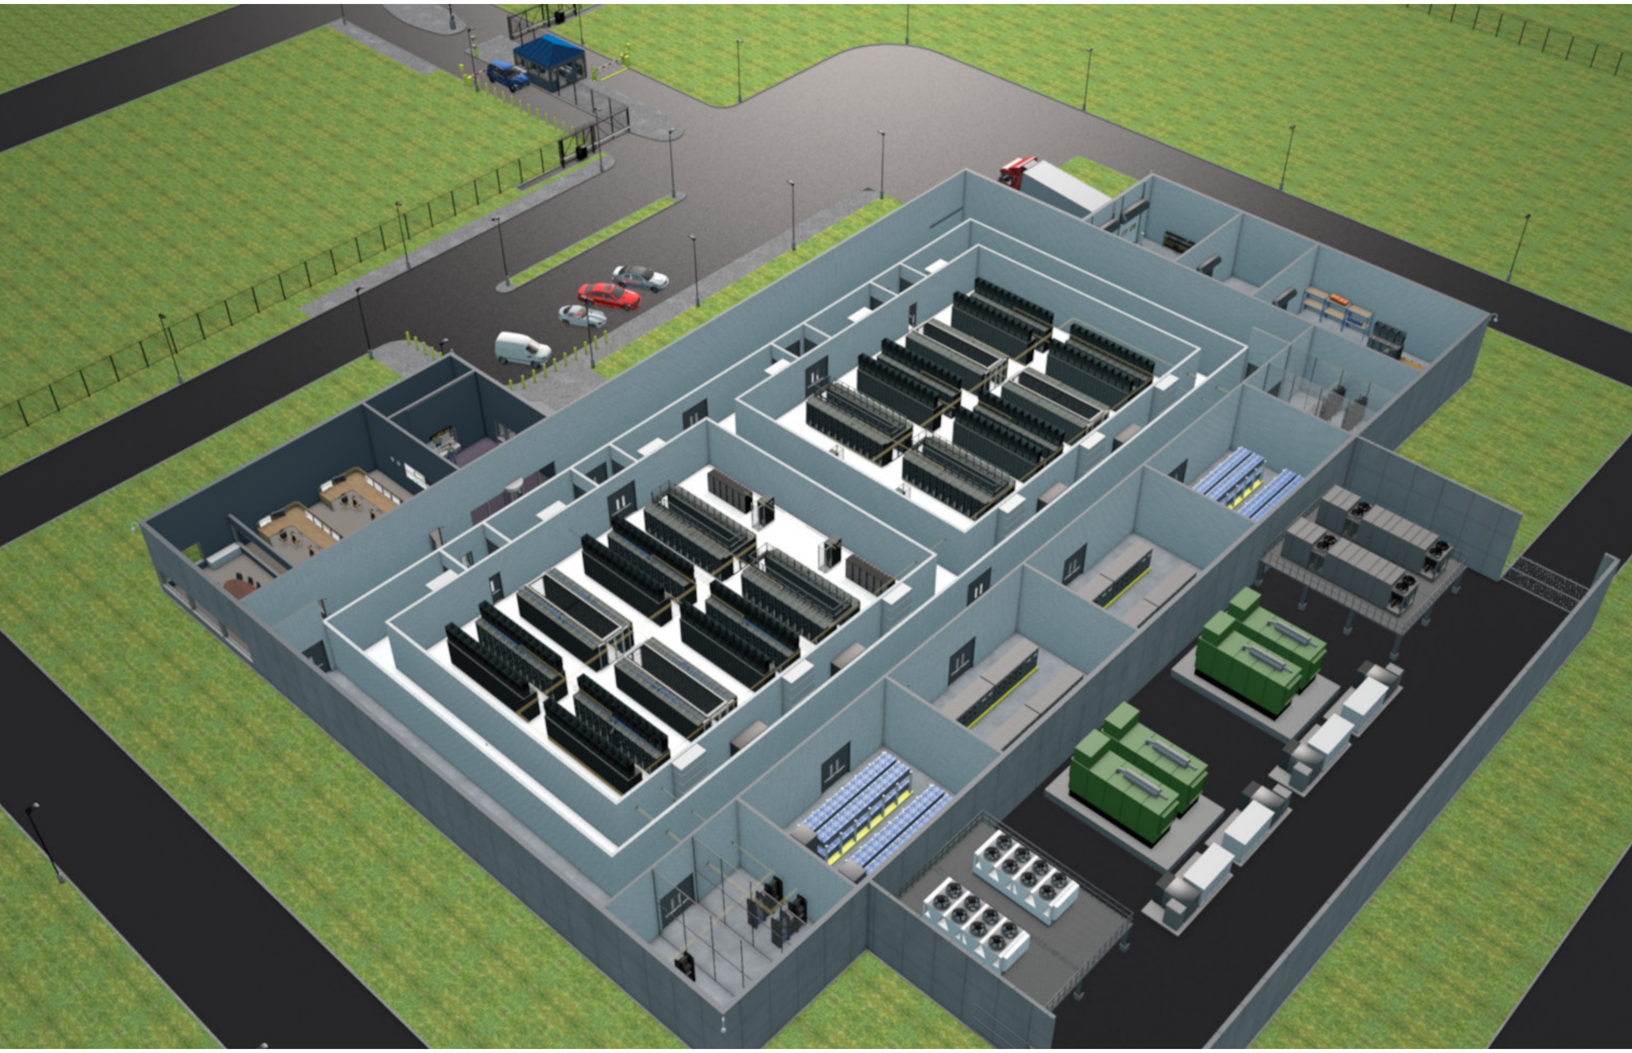

CONTENTS

GLOBAL TECHNOLOGY BRIEFING

RISK MANAGEMENT BEST PRACTICES

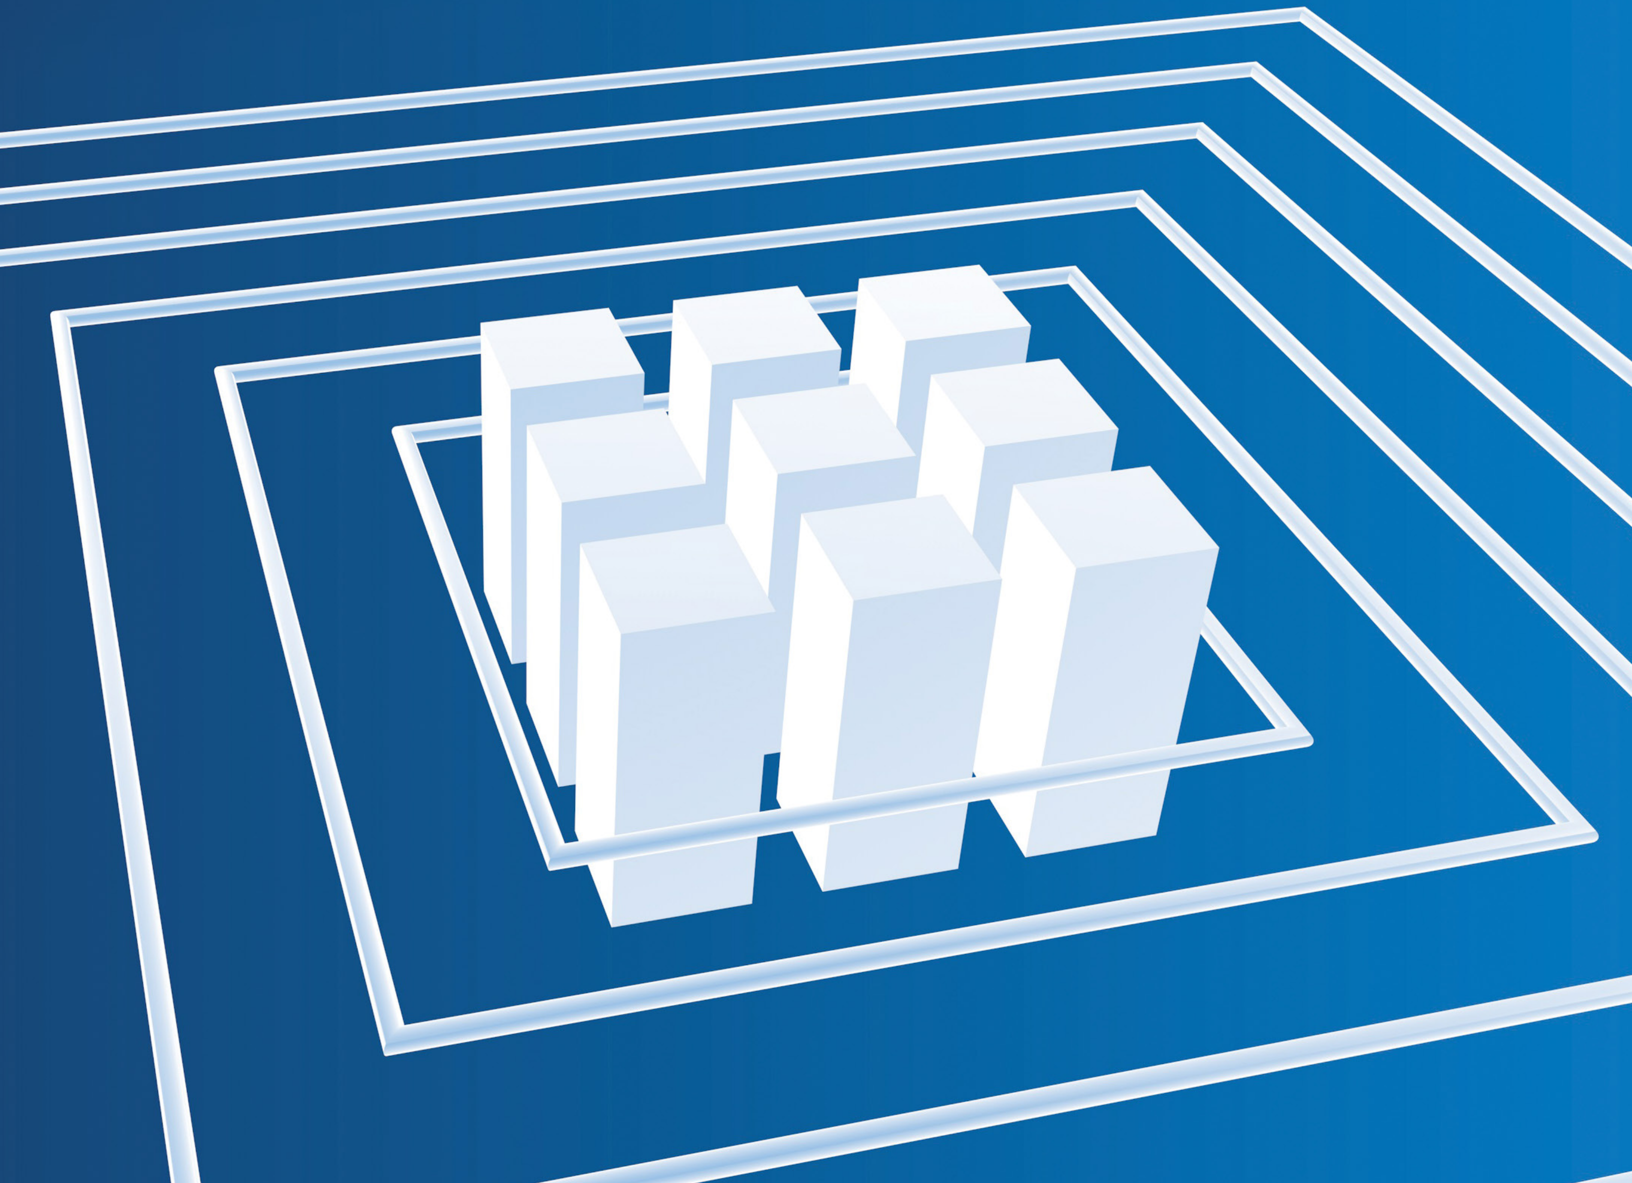

EXECUTIVE SUMMARY

This report identifies how the rapid rise of the data centre is leading to an increased risk of theft, sabotage and corporate misconduct. However, many data centre operators focus solely on logical security, ignoring physical security and its potential threats. This type of thinking leads to vulnerabilities in a data centre's security and increases its exposure to unnecessary risk.

Balancing physical security with logical security is the only way to protect a data centre. This paper analyses the cost vs. risk of data centre security, the importance of physically protecting assets and the value of declaring security budget ownership. Driven by industry regulations, design standards and best practices, physical security in the data centre is as important of an investment as the latest firewalls and cyber security protocols.

Throughout the chapters, the paper lays out the steps to create an effective multifaceted physical security strategy that focuses on implementing a six layered physical security approach.

Finally, this paper examines the need for data centre physical security and the risks to organisations that don't properly protect its critical data. It also sets forth a series of standards and best practices for protecting data centres at six different layers.

The information presented in this paper is for informational purposes only and shall not be considered technical advice. Readers should seek appropriate independent professional advice prior to relying on, or entering into, any commitment based on material published here, which material is purely published for reference purposes alone.

INTRODUCTION

Whether it's for business or education, entertainment or shopping, nearly every financial transaction, phone call or text, movie download or Internet search either now takes place or is recorded in a data centre. With nearly everyone's professional and personal lives dependent upon a healthy ecosystem of data centres, it is only natural that data centres are now targeted by thieves, spies and others maliciously seeking to cause damage or to steal the information contained within a data centre.

Data Breaches are Universal

As data centres have increased in importance, data breaches have become nearly universal. A data breach has happened at each of the 314 companies surveyed in a Ponemon Institute study (*"2014 Cost of Data Breach Study: Global Analysis"*) with those breaches ranging from over 2,000 records to "mega breaches" affecting well over 100,000 records.

Never before has a data centre breach been more damaging to the future of an organisation. As the risk has risen, a growing array of sophisticated threats continues to emerge designed to penetrate data centre defences.

SECURITY BREACH FACTS

- The average cost of a data breach is $3.5 million. Not developing or adhering to proper protocols has proved very damaging, as 39 percent of all breaches are a result of negligence.
- Malicious attacks are on the rise and threaten to supersede negligence as the most common cause of data centre breaches.
- Despite these costs, data centre security professionals report only being allocated, on average, 50 percent of the budget that they believe their organisations require.

Source: Ponemon Institute 2014 Cost of Data Breach Study: Global Analysis

Physical Security: The True Bedrock of Data Security

Data centre security is typically thought of as a logical security issue – the use of firewalls, intrusion detection, the hypervisor security protocols in a cloud environment, and other strategies to deter and defend against hackers and other online breaches. However, controlling physical access to the data centre is the true bedrock of data security and guards against:

- Hardware theft
- A malicious attack on a network port through a virus physically attached
- Corporate espionage
- Unauthorised access to sensitive computing and telecom equipment
- Disgruntled employees and contractors

Data Centre Security

When reading through this report, ask the following questions:

- Whom in the organisation is ultimately responsible for the physical security in the data centre?
- What is the physical security strategy?
- Does the focus on and budget for physical data centre security match the risks?
- Top to bottom, is the organisation committed to data centre physical security, or is it complacent?
- What is the weakest link? Is it the cabinet, grey space, visitor tracking or something else?
- When was the organisation's physical security policy and procedures last updated or reviewed?
- Does the organisation regularly test and enforce physical security procedures?
- What physical security measures are in place to protect the organisation's data centre?

We've shown some of the generic threats to the physical security of the data centre in Figure 1.

Don't Wait for a Catastrophe

Many organisations are not focused on physical security in the data centre, and in the past, have ignored the issues we raise in this report. However, this can be catastrophic for an organisation – as well as for the careers of the executives that didn't develop, implement and execute a suitable data centre physical security strategy.

Figure 1: Data centre security threats

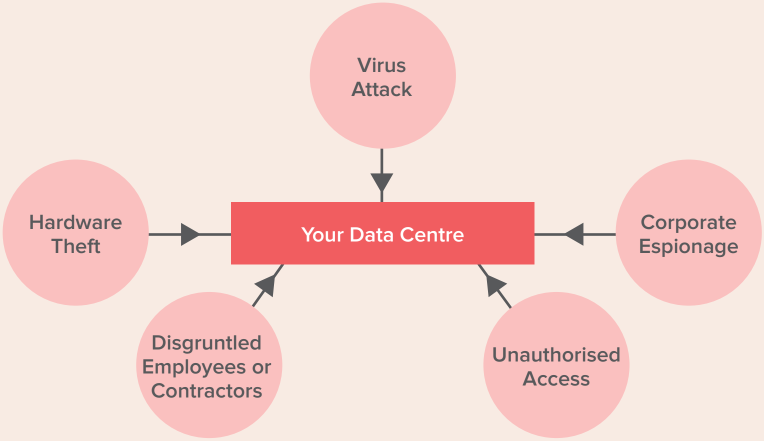

THE RAPID RISE OF THE DATA CENTRE

Importance and Critical Nature

In a way that would have been unimaginable just 20 years ago, nearly all the key business functions and communications of enterprises, businesses and government agencies are now occurring in data centres, which host the information and applications that an organisation uses as the backbone to serve its customers. It frequently contains all the organisation's communications, emails, records, financial activities and customer lists used to serve customers, employees, partners and other stakeholders. DCD's global data centre market overview and forecasts 2014-2020 shows a continuous rise in the requirement for data centre space (see Figure 2).

Figure 2: Data centre space (m^2 billion) requirement forecast – 2012-2020

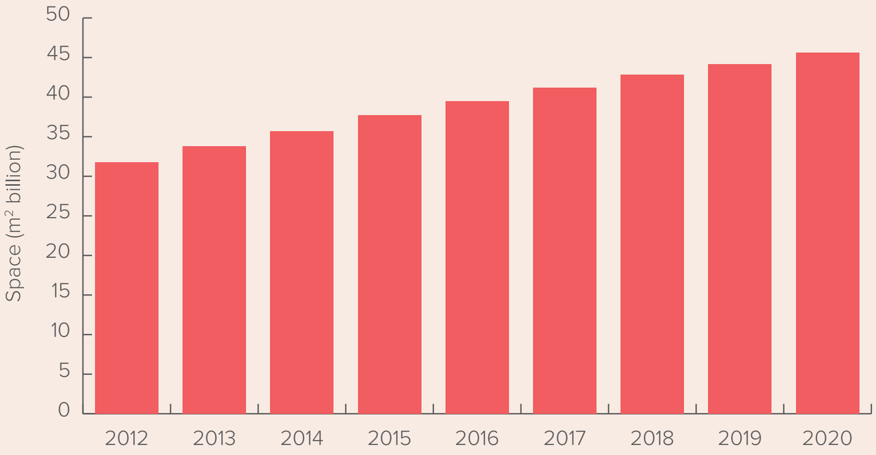

1 meter = 3.28 feet

Source: DCD Global data centre market overview and forecasts 2015-2020

Impact of Failure

This increased leveraging of technology allows for far greater productivity, efficiency, scalability and speed than has ever been enjoyed before. The positive attributes of the rise of the data centre have a counterpoint: the disturbing spectre of potential failure. A security breach could have a devastating impact to an organisation if something knocks it offline. Such a breach could have consequences beyond just the organisation that is directly affected. Imagine the economic impact if a major financial institution suddenly could not process customers' financial transactions: local economies, even national and global economies, could be affected.

Mission Critical

Nearly everything an organisation needs to accomplish is contingent upon efficient, secure data centre operations. The importance of the data centre in the achievement of an organisation's mission and business goals cannot be overstated.

Big Data, IoT, Mobility, Cloud: All Growth Trends Continue

Nearly all trends in technology are fuelling data centre growth (see Figure 3). Big data analytics, the Internet of Things (IoT), the continued explosion of device-driven mobility and cloud computing are just four trends driving growth in the data centre. As companies and individuals continue to rely more on technology in nearly every aspect of their lives, it is hard to imagine a future where data centres do not continue to play a central and growing role.

Figure 3: Data centre growth drivers

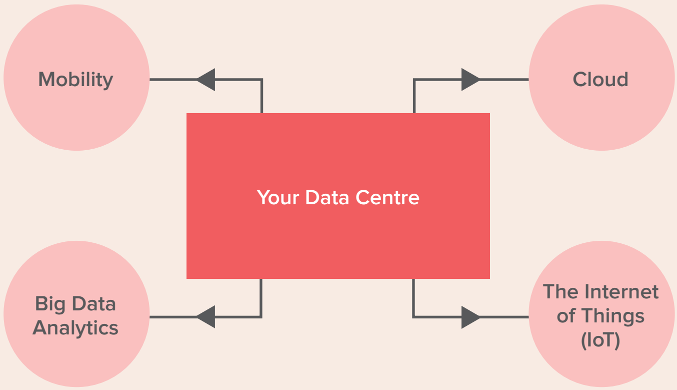

Outsourcing Fuels Data Centre Growth

Many companies, especially large enterprises, still prefer the control and security of ownership of their own data centres, and they will continue to do so. The continuing trend of outsourcing non-core operations has led to explosive growth of cloud computing and multitenant data centres. Figure 4 shows DCD's forecast for space requirements by area between 2012 and 2020: it demonstrates the greater need for co-location/outsourced sites than in-house ones.

Figure 4: Data centre space (m^2) growth forecast – in-house, co-location/outsourced and total – 2013-2020

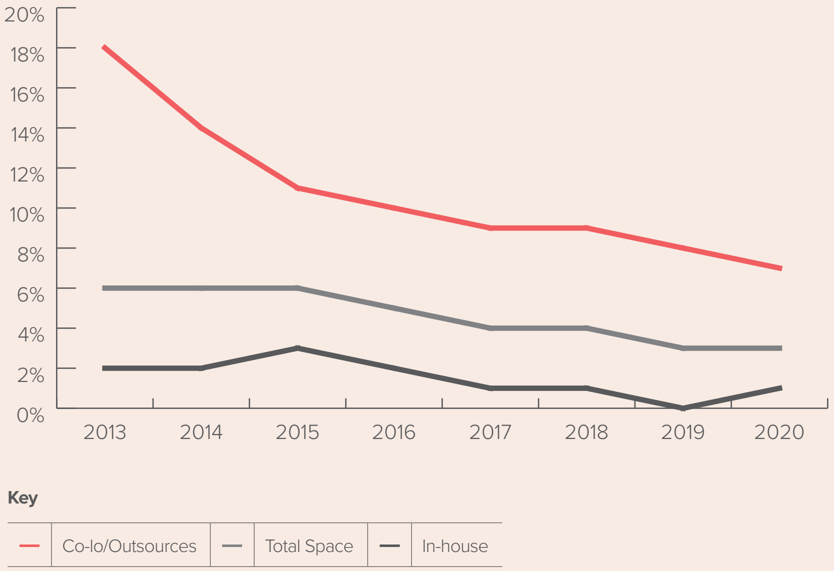

Source: DCD Global data centre market overview and forecasts 2015-2020

Cloud Computing Hubs

The two major beneficiaries of the trend toward outsourcing of IT infrastructure have been cloud computing and multitenant data centres. Cloud platforms are hosted in data centres. The most successful cloud providers, Amazon Web Services and Microsoft's Azure and Office 365 SaaS platforms, are building out data centre space worldwide at a rapid pace.

Multitenant Data Centres

Many organisations prefer multitenant data centres where their compute resources are not pooled, but the other aspects of infrastructure and operations (electrical infrastructure, generators, fibre links, security, etc.) are handled by the data centre owner on behalf of its tenants.

RISK IN THE DATA CENTRE ENVIRONMENT

There are all kinds of risks in the data centre (see Figure 5): fire, contamination (which is why no cardboard is allowed), water damage, power interruption and thermal. As a risk factor, physical security is frequently overlooked, despite its critical importance. As the data centre becomes increasingly integral to an organisation, the risk of an IT security breach increases and can cause far greater harm than ever before.

Figure 5: Physical threats to the data centre

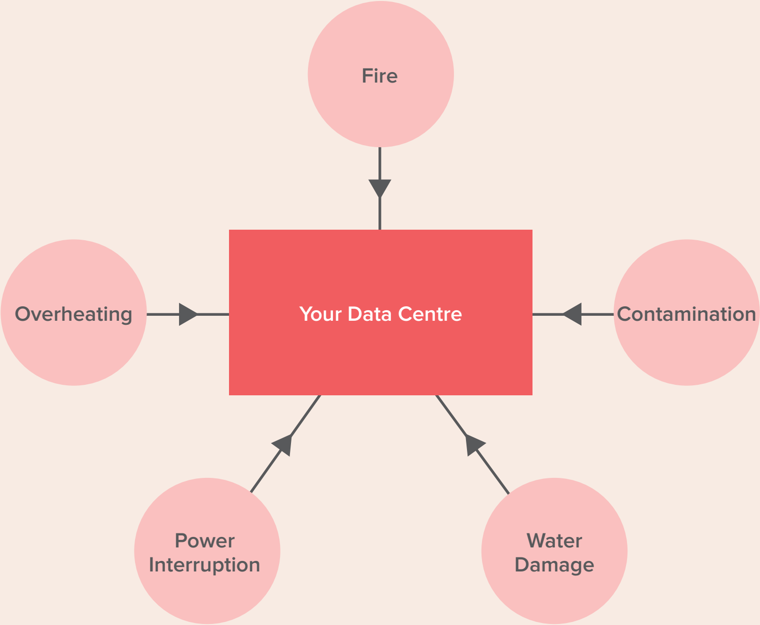

RISK MANAGEMENT BEST PRACTICES

Risk Management

An organisation must determine what the risk is of a data breach and use that to determine its investment in risk management. A key component of a risk management plan is the proper tools to establish a solid baseline of defence against potential threats, but equally important is the organisational buy-in needed to managing that risk.

DID YOU KNOW?

Despite the spotlight on data breaches, security experts agree that only a small fraction of these breaches even become known to the public. Most are handled quietly, particularly in B2B environments where both sides of a business relationship understandably choose to keep the incidents quiet to avoid embarrassment and the appearance of vulnerability, incompetence and mismanagement.

Security is Everyone’s Responsibility

Key stakeholders throughout the company should understand the damage that can be done and demand focus on and adherence to best practices in protecting against potential breaches can cause.

Even though the Chief Information Officer (CIO)/ Chief Information Security Officer (CISO) has frequently been the one to face the firing squad in the case of a data breach, Chief Executive Officer (CEOs) can also lose their jobs if they fail to foster an atmosphere of strict compliance with the best practices of data protection and security. A data centre security event can easily become a ‘resume-generating’ event not just in the executive suite, but among frontline staffers as well. Technicians, security guards, cleaning staff, operations professionals and others in the data centre need to make physical security a team effort.

Avoiding Data Breaches is Everyone’s Job

It is important for everyone in an organisation to understand that data breaches are not restricted to hacker incursions into a system. Just a single individual letting down his or her guard in a data centre and permitting an unauthorised person access to the IT infrastructure is all it takes for a malicious individual to get access to a server and upload a virus, cut some cords, rip out a server and walk out. Human error, mistakes and failing to follow protocols (assuming best practices are in place) can lead to tremendous damage.

THE COST OF SECURITY VS. THE RISK

The costs of a data breach are significant: a 2014 study by Ponemon Institute found that the average cost of a data breach was $3.5 million. This damage can take place whether the incursion is via a hacker or via a stolen server. In fact, lost or stolen devices are the number one factor in increasing the per capita cost of a data breach. Logical and physical security must be taken seriously and work in conjunction with one another.

WHAT WE HEAR

Challenges from the various data centre stakeholders are:

- Protecting the company's image
- Achieving regulatory compliance (HIPAA, PCI-DSS, SOX, GLB)
- Balancing investment in both physical and logical security
- Staying one step ahead of adversaries with technology
- Executing policies and procedures specific to the data centre

Investment vs. Risk: Striking a Balance

Physical security deployments are significant investments with the majority of costs incurred up-front. Once installed these security systems are often considered adequate, receiving little to no maintenance post installation. However, as we have discussed risk is continually evolving as equipment nears end-of-life, processes become outdated and the those with malicious intent get more sophisticated.

By building a scalable interoperable security solution, pro-active updates and upgrades are simpler, quicker and more cost effective than complete system revisions, helping you effectively balance the costs of maintaining a robust physical security system with reduced risk of attack and breach (see Figure 6).

DON'T BE A STATISTIC

Average organisational cost of data breach was $3.5M. The number one factor that influenced the cost was lost or stolen devices.

Source: Ponemon Institute 2014 Cost of Data Breach Study: Global Analysis

Figure 6: Investment vs. Risk

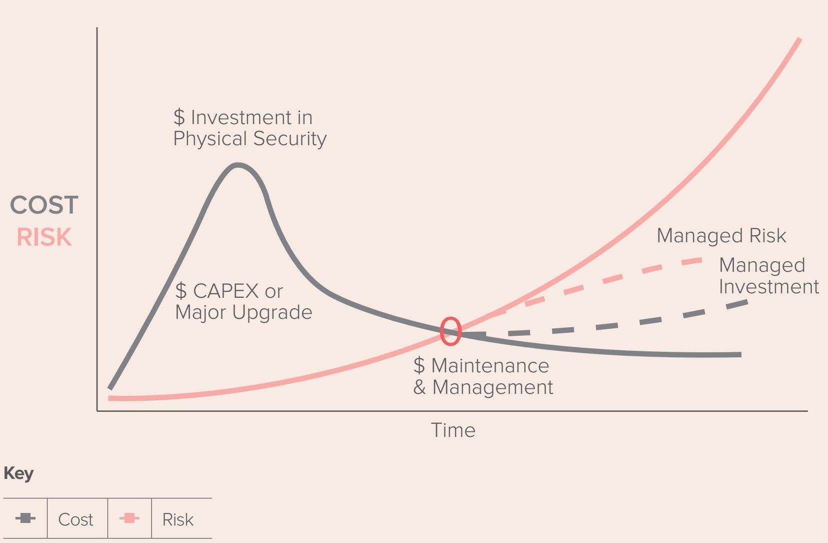

Defending Against the Array of Constantly Evolving Logical and Physical Security Threats

Any data centre operator, as well as any tenant in a co-location environment, has a need to stay on top of rapidly evolving and growing threats to its security.

Serious, Smart and Skilled Adversaries

The malicious thieves, corporate spies and revenge-minded vandals executing attacks of all kinds are numerous, relentless and ruthless. The perpetrators may be engaging in corporate espionage, they may be seeking to financially harm an organisation or they may simply be maliciously seeking to damage an organisation for whatever reason (including just the thrill of it). They may be backed by foreign governments, or they may be competitors, or they may just be common thieves with uncommon skills.

Whoever they are, they can do lasting damage to customers, an organisation and careers. An organisation must invest in stopping them.

Establish a Culture of Preventive and Active Security

Security professionals hear of new tactics on nearly a weekly basis, and truthfully, there is reason to be impressed by the creativity and resourcefulness on the other side of the battle. The competition is tough. An organisation must take the threat seriously and must commit to the fight.

Steps You Must Take

Create a Battle Plan
What are the plans and procedures to defend against threats that are increasing in sophistication and complexity?

Invest the Time and Be Informed
Being properly plugged in with the data centre security world should be a priority for an organisation. Join some of the active, vibrant communities of security professionals who share the latest threats and ways to overcome them. Attend industry conferences (BICSI, ASIS and DatacenterDynamics) and join, participate, listen, learn and share.

Complete Organisational Commitment: Including the Budget
From the boardroom to the front lines, an organisation must have complete buy-in and commitment to the level of security it requires (see Figure 7).

To stay ahead of the risk factors, the strategy must be to exceed the requirements, not to simply accomplish the minimum necessary to hit the marks.

Figure 7: The rights and wrongs of corporate commitment

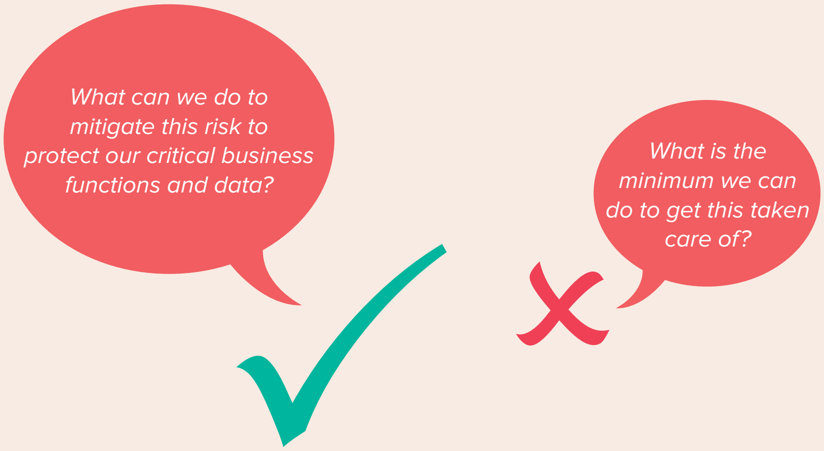
What can we do to mitigate this risk to protect our critical business functions and data?
✓
What is the minimum we can do to get this taken care of?
✗

DECLARE BUDGET OWNERSHIP

Having a substantive budget is not the only issue. The ownership of that budget and the responsibility for it to be wisely used are critical factors. In most organisations, logical security is part of the IT budget. The physical security is not so clear-cut.

Figure 8: Budgeting for physical security – meeting different executive requirements

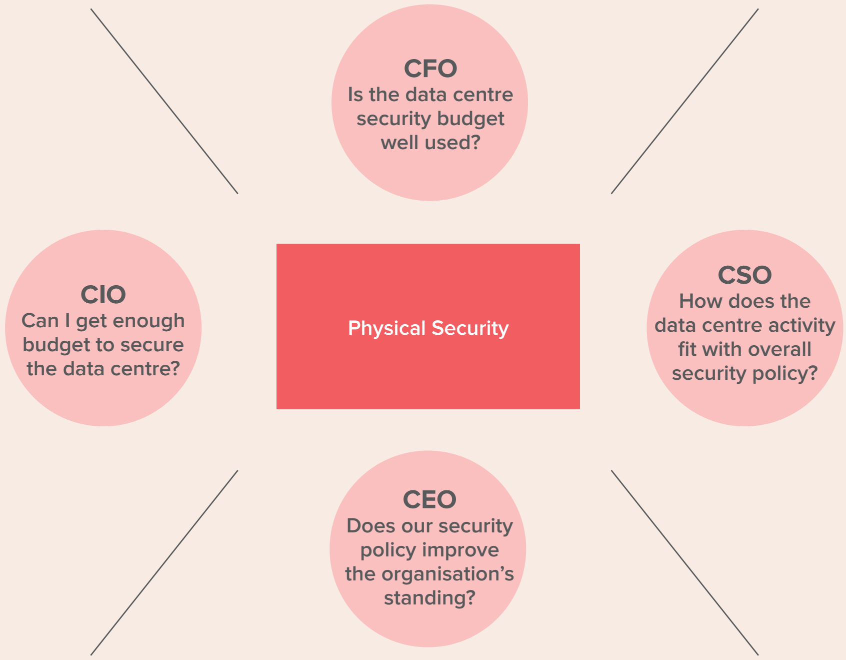
CFO
Is the data centre security budget well used?
CIO
Can I get enough budget to secure the data centre?
Physical Security
CSO
How does the data centre activity fit with overall security policy?
CEO
Does our security policy improve the organisation's standing?

Cost Analysis to Determine the Budget

Unfortunately, adopting the proper defence posture is not an inexpensive proposition. Surveys of data security professionals show that they believe their budget should be 50 to 100 percent higher. When considering the alternative, allocate the proper resources to protect the data centre. The costs of not doing so are just too high of a risk to accept.

PHYSICAL SECURITY BUDGET CHECK LIST

- Where does budgetary responsibility lie for the physical security of the data centre?
- What is the annual budget for physical security of the data centre?
- What is the decision making process for this budget? Who is involved?
- Is the budget focused just on the data centre white space or on all six layers of the data centre, which are defined on page 28?
- Is the budget proactive or is it instead more of a patching of weaknesses?
- How much input is solicited and received throughout the organisation?
- What is the strategy to make sure that the money is used well?
- Can the physical data centre security be scaled as needed?
- Is there the proper interoperability to keep adding pieces or do systems that don't communicate cause a need to "rip and replace"?

We've captured some of the different views of typical executives in Figure 8.

PROTECTING INFORMATION AND PHYSICAL ASSETS

The Environment Determines the Protection

The strategy and methodology for protecting the IT infrastructure depends on the business taking place in the organisation.

- A service provider that hosts its entire business platform in the data centre in a busy city should rightfully have a very high level of security.
- An enterprise data centre with just a few employees in a rural area where it is easy to monitor who comes and goes may not need a large budget to maintain security.
- A federal contractor doing work for a high-security agency needs the highest levels of security.
- A multitenant data centre with 80 customers that has employees and contractors coming and going has very specific security needs, particularly in tracking where visitors go within the data centre and making sure they cannot access the infrastructure of other tenants.
- A large company that has a disaster recovery data centre that backs up email and other records likely would not require as high an investment.

What Tier is it?

Uptime Institute and Telecommunications Industry Association (TIA) rate data centres by tier level, from Tier 1 to Tier 4. There are only a handful of Tier 4 data centres in North America. The more mission critical a data centre is (and the more importance achieving and maintaining a high tier level is) the higher the protection requirements. Tier levels focus on uptime and securing the primary and backup infrastructure that maximises uptime is a critical part of achieving a desired tier level.

Even though the Tier Classification System does not prescribe security provisions, physical security risk factors should be addressed in the owner's operational sustainability requirements. The level of security is largely determined by the industry, the criticality of the IT function performed on the site and the owner's policies. For example, a financial institution will typically invest in a level of security exceeding that of an academic institution.

Within the IT infrastructure there are numerous 'best practice' security measures to be taken, regardless of Tier. These include biometric readers, bollards, guard houses, hiding in plain sight, mantraps, re-enforced perimeter fencing, video surveillance, etc. Best practices are critical in reducing the risk exposure of curiosity, mischief, crimes, and accidents. However, best practices should not be confused with Tier requirements.

Physical Security Plan

A well thought-out, strictly defined physical security plan is necessary when developing and maintaining a strong security stance. Such a plan should cover every aspect of data centre security, from designing a system, to having everyday policies on operations, handling visitors, and emergency and disaster response policies.

Security doesn't just happen; it is a result of a process. A well-designed physical security plan is a key part of achieving the desired results.

Maintaining Sound Policies and Procedures

No matter what the focus of a data centre, it is critical to maintain and follow sound policies and procedures. This is part of an overall security plan that balances best practices with a willingness to evolve to properly defend against new threats.

- Sound policies and best practices support security in a data centre.
- A policy document is a living, breathing thing. Policies are not static. Policies should be frequently reviewed and updated.
- The logical security strategy should have a physical component to it.
- Security protocols should be understood by all and followed closely.
- Complacency must be avoided throughout the organisation.
- Logical security and physical security policies are interrelated and support one another.

DRIVEN BY COMPLIANCE

Security is one of the two key components of compliance, as is uptime demonstrated by the use of redundant systems and best practices for maximum uptime. Many organisations have worked hard to achieve compliance only to lose those hard-earned certifications necessary for operating within their industry. A weakening of an organisation's compliance stance can be devastating to its future prospects.

Addressing physical security is complying with regulations and standards. Those who set the standards compliance in various industries have an increasing understanding of how controlling access to a facility helps threats from interfering with the power and telecommunications systems integral to the operation of a data centre, along with the data itself.

Compliance standards are unique to individual industries, but share many commonalities. For example, Payment Card Industry Data Security Standard (PCI DSS) certification, which covers the payment card industry, requires a compliant organisation to:

- Protect stored cardholder data
- Restrict access to cardholder data by business need to know
- Identify and authenticate access to system components
- Restrict physical access to cardholder data
- Track and monitor all access to network resources and cardholder data
- Regularly test security systems and processes

These security requirements can be met by physical security protocols and support the need for multiple layers of security as defined in this report.

A layered approach that meets the specific needs of the payment card industry to achieve PCI DSS compliance would include steps welcomed by an auditor evaluating the security stance of organisations seeking compliance with the standards in other industries as well. With the same goal of data protection, it is no surprise that different industries have commonalities in what is required for compliance.

There are many other certifications in a wide range of industries (see Figure 9 for how some apply to different industry sectors). Even though there is significant standards overlap, each individual certification is unique. There are several examples of organisations that lost their critical certifications and have hamstrung their ability to operate in their industry or serve customers in a particular industry. An organisation's security stance is an integral part of maintaining compliance and needs to provide sufficient focus and budget to deliver true physical security in the data centre.

Figure 9: Notable industry certifications with security standards

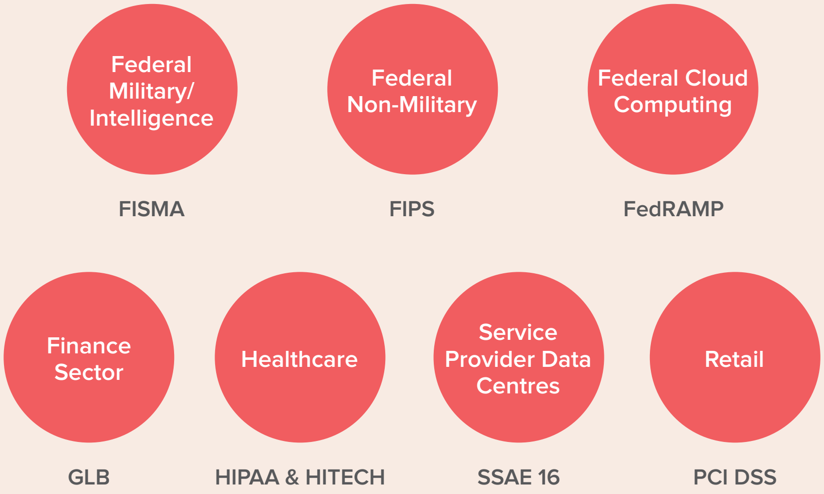

The Importance of Auditors

Independent auditors determine if a data centre meets proper standards, and that auditing report will usually be universally accepted (see Figure 10 for the ideal conversation a CIO might have). Auditors see different things and interpretation can be a challenge. Having clarity on plans, procedures and protocols minimises the risk of leaving an organisation open to an interpretation issue.

Figure 10: The Ideal Conversation with an Auditor

"Here's our physical security plan. We classify our data centre as a critical facility. We have these different rings of security to provide defence in depth. Here are our protocols and procedures. Here's what we do. We're committed to doing things the right way."

DESIGN STANDARDS AND BEST PRACTICES

Whenever possible, implement key data centre design standards. A critical infrastructure approach to physical security for data centres should protect from the facility perimeter to the data centre cabinet by leveraging industry standards and best practices.

Proper design and layout of a data centre enables the effectiveness of these and the many other systems that provide the attributes and services required of a modern data centre. The standards are guidelines that tend to be general, broad and baseline, creating another instance where a comprehensive, strategic approach is more likely to earn the certifications desired.

Data Centre Design and Security Standards

- **ANSI BICSI 002-1014:** Best practices for data centre design
- **TIA-942:** Sets requirements for telecommunications infrastructure within data centres
- **Uptime Institute:** Focused on improving the performance, efficiency and reliability of critical business infrastructure
- **ASIS:** Seeks to advance security standards worldwide
- **SANS:** Information security best practices and standards
- **ONVIF:** Focused on network-based physical security product interoperability

A layered approach maximises the achievement of all necessary certifications. It leverages the insight from all major standards and builds a holistic approach.

Awareness and understanding of these standards is a must for anyone building a data centre, whether it be a greenfield build, a conversion of a facility or an upgrade of an existing data centre.

Integral to the overall standards of data centre design is the understanding that each component of the physical security system meets the standards set by subject matter experts in each area of data centre design. These include technological standards for networking infrastructure, physical infrastructure and physical security.

A MULTIFACETED STRATEGY FOR PHYSICAL SECURITY

Given the increasing reliance upon the data centre, the harm that can be caused by a security breach and risk inherent in the loss of (or inability to earn) necessary certifications and compliance standards can be crippling. As a result, DatacenterDynamics and Anixter have partnered on a project to properly define the expanded best practices required for secure, compliant data centre operations in the current environment.

This technology report defines the six key layers of a holistic data centre's "defence in depth" physical security strategy.

1. Data centre cabinet
2. Data centre room and white space
3. Hallways, escorted areas and grey space
4. Facility façade and reception area
5. Clear zone
6. Perimeter defence

Some data centres are not designed to support all six layers and some layers may be combined. For instance, a smaller data centre may not have a significant clear zone, with the perimeter defence leading immediately to the facility façade and reception area. Additionally, a multitenant data centre with significant foot traffic may have different security needs than an enterprise data centre, but the proper physical security protocols and best practices should be implemented wherever possible.

MACRO-SEGMENTATION STRATEGY

Logical security frequently relies upon the concept of micro-segmentation, which creates barriers within a virtual environment to keep someone who has penetrated the system contained. There are barriers placed throughout the system to keep them from moving from virtual machine to virtual machine. When properly designed and implemented, it keeps a hacker contained while also creating alerts to inform those monitoring the system of any potential incursion.

The six layers of physical security defined here support a macro-segmentation strategy for physically defending a data centre. Like a virtual micro-segmentation strategy, a physical macro-segmentation strategy should strive to not limit the performance of the data centre, in this instance, by making a data centre overly difficult and time-consuming to visit or navigate.

Six Layers Supporting 5 D's: A physical macro-segmentation strategy seeks to limit the damage of a threat by supporting the 5 D's of perimeter security.

1. Deter
2. Detect
3. Delay
4. Defend
5. Deny

A layered physical security approach provides a strategic series of obstacles to protect against a potential physical incursion of a data centre, making it increasingly difficult to gain access to the mission-critical data.

The Building Blocks

Key industry protocols leveraged as baseline standards for the six layers include BICSI, TIA, ASIS and ONVIF. These and other standards were critical in the development of a holistic physical security approach designed to protect the data centre.

The Necessity of Interoperability

The entire ecosystem that serves the data centre physical security market must do an increasingly better job of providing interoperable solutions that support standards-based open architectures. A multifaceted, layered approach will have different components provided by best-of-breed manufacturers. Systems that don't integrate with one another are an impediment to the mission of keeping data centres secure.

Network-based solutions are clearly the future of data centre security and disparate systems that do not talk to one another will be left behind. Decision-makers must not tolerate inefficiencies in systems management. Proprietary manufacturing can lead to a dead end for the end-user seeking a comprehensive solution from multiple providers. Open-architecture solutions enable a scalable, flexible, long-term security solution and put the end-user in control.

Physical security systems have reached a significant point in their evolution. Formerly, manufacturers built proprietary systems without regard to integration with others. Limited interoperability then emerged as manufacturers began building their own ecosystem of partners; however, with volatility in the physical security market, these solutions can put end-users' security platforms at risk. The industry has now grown to the point that a true open architecture has emerged to support standardised communication between network-based physical security products, regardless of manufacturer.

True Scalability

True, standards-based open systems and open architectures allow end-users and their integration partners to be in control of the system. They can confidently scale and build in forward compatibility as they grow into their future needs. With the addition of physical security standards, end-users can now build end-to-end standards-based physical security solutions to protect and scale with their critical data centre environment.

The graphic represented below shows some of the industry standards that support a true open architecture.

Networking Architecture Standards	IEEE IEC ISO International Organization for Standardization PoE Power over Ethernet Wi Fi™
Infrastructure Standards	CENELEC IEC ISO International Organization for Standardization TIA ADVANCING GLOBAL COMMUNICATIONS
Physical Security Standards	SIA SECURITY INDUSTRY ASSOCIATION UL ITU ONVIF Driving network video through global standardization SMPTE

An Open Systems Challenge: Manufacturers must be truly involved in open standards development, and not just give the appearance of compliance in order to optimise sales. As an end-user, the flexibility, efficiency and scalability of a security systems depends upon a true commitment to open architecture from the technology partners. Technology partners should be asked how they are investing time, energy and resources to continue the development of open physical security industry standards. This is a good indicator of the type of technology partner being selected as part of a data centre physical security strategy.

An effective, layered approach requires all systems to work in a cohesive manner. Interoperability has become an absolute requirement in the design and operation of a data centre security strategy and manufacturers, integrators and other suppliers that do not focus on the need for interoperability will be left behind.

THE SIX LAYERS SUPPORTING THE PHYSICAL SECURITY OF A DATA CENTRE

LAYER 1: PERIMETER DEFENCE

Recall the six layers of data centre physical security should be based on the 5 D's of deter, detect, delay, defend and deny. The first layer - perimeter defence - controls authorised and unauthorised access to the data centre's property. When properly implemented, the perimeter defence layer can reduce the overall cost of a data centre facility's security system and improve the effectiveness of the security plan.

Figure 11: Layer 1: Perimeter Defence

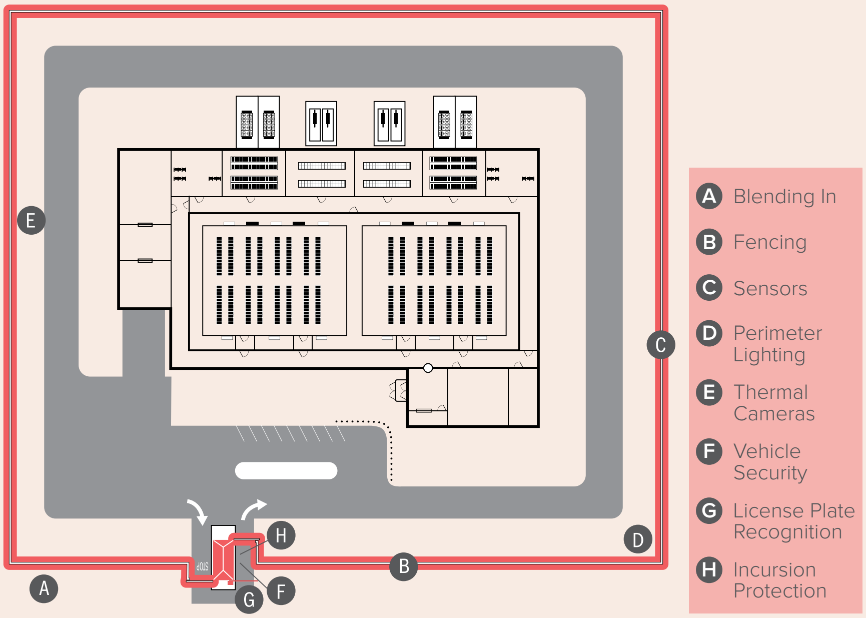

Crime Prevention Through Environmental Design

The concept of crime prevention through environmental design (CPTED) is a philosophy that encourages the use of architectural design to support security.

CPTED is a philosophy of territorial control, where natural surveillance and access control are used. Think of berms, big boulders, clear views without trees or brush in the way for natural surveillance and other methods of natural access control. This not only allows for better monitoring, but also serves as a deterrent as well. It supports the perimeter defence and frequently does so at a low cost.

Blending In

Most data centres do not have signs that indicate the nature of the business that takes place. Enterprises have no reason to attract attention to the data centre and multitenant data centres' clients frequently feel the same way.

Fencing: The first line of facility protection

For many data centres, a relatively simple fence that separates a facility from the environment may be suitable. However, for high-security data centre installations, numerous fencing options are worthy of consideration.

Sensors for Security

Sensors for perimeter intrusion systems use special cable media (copper and fibre optic) and electronics to sense vibrations and disturbances to identify if there is an intruder attempting to compromise the perimeter defence. These solutions can be deployed as zones in fences, walls and underground applications. They can also seamlessly integrate with various physical security subsystems to provide complete situational awareness.

Perimeter Lighting

Lighting is an important part of a data centre's perimeter security intrusion detection system. Not only does it provide improved visual coverage for security personnel and video surveillance cameras, but it also improves the effectiveness of the overall security system. Most perimeter lighting applications have turned to infrared (IR) and white-light LED lighting solutions in place of traditional outdoor lighting such as low-pressure sodium, halogen and metal halide.

IR LED lighting solutions can be used in covert perimeter applications covering long distances. White-light LED has multiple functions, such as providing light for video surveillance cameras to identify intruders, as a part of a CPTED strategy as deterrent lighting, or for general area illumination for the safety of site personnel.

Thermal Cameras

The perimeter of the facility can also be monitored with thermal cameras. In many cases, only the main entry point will be well-lit. A thermal camera allows confirmation of someone outside the facility even in complete darkness by detecting heat signatures emitted by objects. Thermal cameras can also be used with video analytics (perimeter tripwire) and other intrusion detection sensors to create a strong perimeter defence.

Vehicle Security Station

The key entry point for many data centres is a manned security booth at a vehicle entry point and/or a camera and audio system that allows an internal security desk to communicate with those in vehicles as they go through a checkpoint. Only authorised visitors are allowed and tools such as mirrors and bomb-sniffing dogs are frequently used at higher-security facilities. A motorised operating gate is an additional option for both deterrence and protection.

As a vehicle approaches a data centre's guard station, long-range badge readers can determine the identity and level of clearance of someone driving to a gate and assist the security personnel in identifying those occupying the vehicle.

License Plate Recognition (LPR)

Specialised infrared-sensitive LPR cameras are capable of capturing the license plates of even fast moving vehicles in ambient light levels from bright sunlight to complete darkness. The license plate numbers are then converted into a computer-readable format and compared with a database of vehicle registration numbers. This allows an assurance that an authorised user is in an authorised car, allowing a guard to quickly validate this person or to discover why a non-authorised person is approaching the facility.

Incursion Protection

Wedge barriers that rise up out of the ground to create an obstruction to potential intruders and crash-proof fencing may be used to guard against vehicles engaged in an incursion. The outer perimeter is another area where policy, video, deterrence, protection and active monitoring create a secure environment.

LAYER 2: CLEAR ZONE

The clear zone is extremely important. After getting through the perimeter, the threat may have access to critical electrical and mechanical areas, such as the primary power plant and power wires running into the facility, or other backup areas, such as generators and fuel tanks. The clear zone also contains equipment loading docks and secondary entry points into the facility.

Video Surveillance

Video monitoring is a key part of data centre access control. It can serve as a deterrent, as a monitoring tool and as a way to review an incident. An active video monitoring system can guard against "tailgating," and video monitoring will help enforce security policy.

Video has uses in every layer of a data centre's defence, and defending the clear zone is no different. In the clear zone, the primary goal is to identify individuals and monitor restricted areas. Video surveillance is key to keeping the clear zone secure.

DID YOU KNOW?

To enhance the performance of the video surveillance system, a best practice is to use LED lighting in the clear zone in place of low-pressure sodium or mercury vapor light. LED lighting offers numerous benefits from energy savings to reduced maintenance. LED lighting also provides superior lighting conditions allowing cameras to deliver high quality nighttime images. Here are some of the major benefits of using LED lighting with video surveillance:

- › Increases high-resolution camera performance in low-light
- › Reduces bandwidth consuming video noise during low-light scenes
- › Provides better colour rendering vs. traditional lighting (e.g. low-pressure lights produce a monochromatic light that impact colour rendering)
- › LED lighting is offered in both infrared (IR) and white light options

Figure 12: Layer 2: Clear Zone

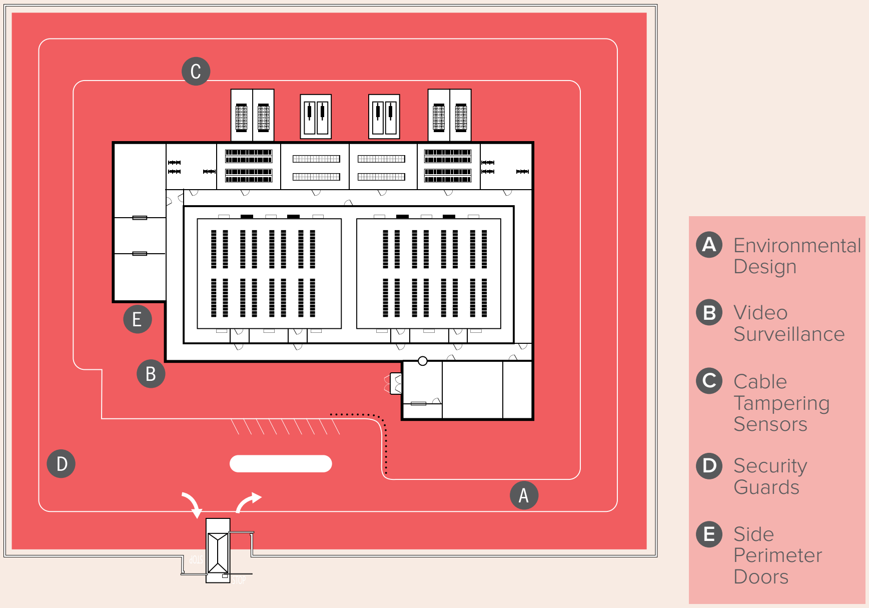

Video Strategy Supporting Security Protocols

There are several strategies to support security protocols with video including the following:

- A focus on power sources. Who is in proximity of the power infrastructure, fuel tanks, loading areas and generators within the clear zone?
- Use high-resolution 180° degree cameras. Pan, tilt and zoom cameras don't provide constant coverage.
- Detection and identification of people within the clear zone. 20 to 80 pixels-per-foot cameras are required to see who is approaching the facility.
- Get complete coverage of the clear zone. Cameras may be required from the building looking out and the perimeter looking in to ensure there is complete coverage and no dead zones.

Determining the Optimal Camera Resolution

Paying attention to optimal camera resolution makes sure there is proper image detail (pixel density) in everyday and worst case scenarios. Many times, it seems easy to pick a high-resolution camera and feel comfortable that it will do the job. That's not necessarily the case. Too high of a resolution camera can be overkill for the application and create unnecessary stress on the network by creating excessive bandwidth and storage requirements and costs. On the other hand, too low of a resolution camera can be the difference in being able to properly identify a potential threat or suspicious activity. So, what is the best way to optimise camera resolution? Here are some simple steps to follow to help optimise camera resolutions for data centre video applications:

- Determine operation requirement. What is the 'everyday' use scenario? What is the 'worst case' use scenario?
- Work with technology partners to help determine the pixel density requirement to meet the operational requirements.
- Be aware that complex issues such as lighting, optics, compression and others can impact image quality.
- Leverage the pixel counting tools available in most network cameras to validate pixel density once a camera is installed.

When developing a video strategy, it is important to remember that video is only as good as the people who are watching it. How vigilant are they going to be after they've been on the job a while? How alert and focused is the overnight security personnel at 4 a.m.? Those realities are fuelling the growth of video analytics.

Video analytics can be used at different layers. Uses for video analytics in the clear zone include the following:

- Alerts in case individuals are loitering in monitored areas based on suspicious movement.
- Object classification provides alerts based on the type of object (person, vehicle, animal, etc.) detected.
- Direction flow provides alerts based on the direction of a moving object.
- Object added/removed alerts if an object is added or removed from a predefined area.
- Motion tracking follows a moving object across the cameras view.

A SECURITY GUARD REALITY CHECK

Security guards may have varying degrees of experience and training. If the security staff is outsourced, that is another level of control over the guards that isn't available. Given the varying quality of security guards, company policies and procedures are paramount, as is the buy-in of everyone in the facility to support the security protocols of the data centre.

Cable Tampering Sensors

In high-security data centre applications, it is important to deploy physical intrusion detection sensors that protect critical telecommunications and power transmission cabling infrastructure throughout the various layers of the data centre's physical security strategy. Sensors should be deployed in parallel within a network conduit, embedded in a carrier or used to physically protect optical networks from an individual attempting to compromise the infrastructure. This is done by providing a security alarm to detect a physical intrusion or the tapping of information. These systems can also aid in identifying if an intrusion in the network is a physical or logical attack.

Side Perimeter Doors

Another element of clear zone security is protecting non-main doors from attack. As a physical barrier, a door with a single point of latching into the frame is easy to defeat. Higher security, multipoint locking systems are available. They add latch points on the top, bottom, middle and hinge sides, making the door invulnerable to physical attack from outside.

Use of Controlled Key Systems

In a restricted, high-security key system, key blanks are under control of the manufacturer. Under no circumstance can they be copied anywhere else and the distribution and management is strictly controlled.

LAYER 3: FACILITY FAÇADE AND RECEPTION AREA

Commit to Visitor Management

Visitor control starts at the reception area and sets the tone for the data centre. Visitor control is becoming a larger part of achieving compliance. PCI DSS, a key financial industry certification, is among the standards requiring strict visitor controls. In a PCI DSS compliant facility, a visitor needs to be clearly identified via a visitor badge.

A visitor badge is not merely a generic badge; rather, it should be created with a photo badging system onsite for visitors. The visitor badge should additionally have an expiration date and time. This badge clearly indicates who the person is and how long they are permitted to be in the facility. A good practice is to fully integrate the visitor management software with the access control system. This gives security personnel one place to manage both visitor and employees.

Complacency is the Enemy of Security

The security staff in the reception area must understand that just because a contractor or visitor has been through the process before, it doesn't mean that the process can be relaxed in any way. Complacency should never set in. Complacency is a breeding ground for a security breach, allowing even the most sophisticated security apparatus to be defeated.

UNDERSTANDING THE ENVIRONMENT

Frequently, sunlight permeates a glassy visitor area. As a result, there may be a need for wide dynamic range (WDR) cameras. A standard camera without WDR cannot compensate for the dark inside of a building and bright sunlight light coming into reception area; WDR cameras adjust and compensate for both bright light and dark areas.

Figure 13: Layer 3: Facility Façade and Reception Area

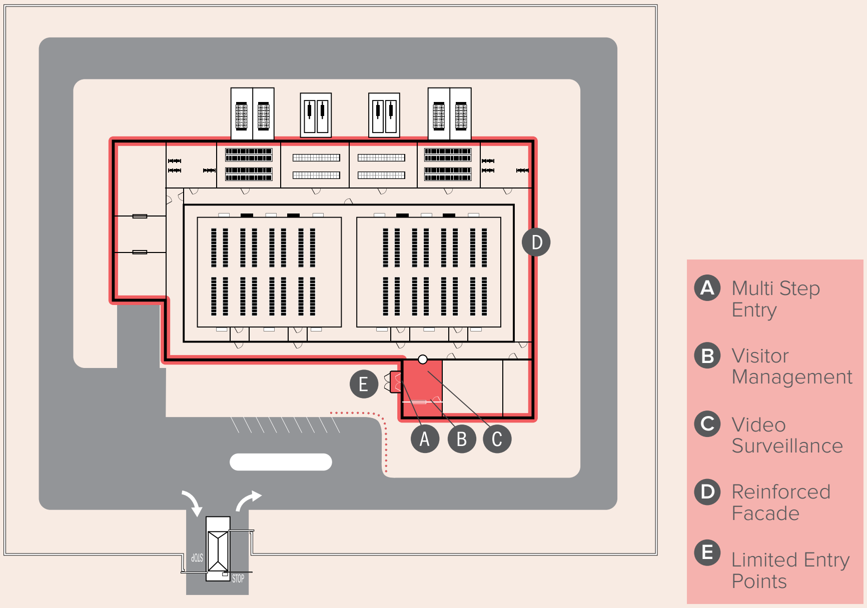

Interoperability Keys Guest Verification

An access control system can integrate and sync up with HR management systems such as Oracle PeopleSoft in order to verify the to-the-minute employment or contracting status of the visitor. This avoids the potential actions of a disgruntled employee immediately following termination.

Multistep Entry

Ideally, a visitor would be badged in before the reception area. A video camera near a card reader can see who that person is and where he or she is coming into the facility. An intercom system is used to facilitate communication and verify that someone has business with the facility. The intercom system can utilise VoIP technology and become a part of the communications infrastructure.

Video Surveillance

Like the previous areas of the data centre, video is critical as visitors approach and enter the building. High-resolution cameras are recommended to provide clear identification of who the person is coming into the facility. Also, if the video and access control systems are integrated, a positive identification and verifiable record of entry is provided for each instance.

Everyone’s Responsibility

An adherence to policies and procedures is extremely important when allowing admittance to the data centre. It starts with the security staff and receptionist, but following best practices as laid out by the organisational security policy by all employees in the data centre is required at this critical stage of the security process.

LAYER 4: HALLWAYS, ESCORTED AREAS AND GREY SPACE

Most data centres have significant focus on the white space of the data centre room, but the grey space, hallways and escorted areas that lead to the data centre floor are frequently an area where proper security measures are overlooked. It is helpful to think of the grey space as the perimeter to the data centre and secure it appropriately.

Who is in the Grey Space?

Visitor control is critical. Without proper visitor tracking if a group is touring a facility, it is not difficult for someone to slip away to use a restroom and not be noticed if they fail to return. There is now an unescorted, unaccounted for visitor one open door away from the data centre white space.

Mantrap In and/or Out

Allowing only a single person in is obviously more secure than allowing a group. Users of a mantrap, and truly all secure areas of a data centre, must be aware of potential “tailgaters” who follow an authorised individual closely to bypass the proper security procedures. Even if the individual is authorised, tailgating is a breakdown in the ability to track those in the data centre and must not be permitted. A strict policy should be in place that does not permit such unauthorised access.

A mantrap is frequently used either into the grey space from the reception area or out of the grey space into the data centre halls. This is not mandatory, but it is helpful for higher security installations.

Figure 14: Layer 4: Hallways, Escorted Areas and Grey Space

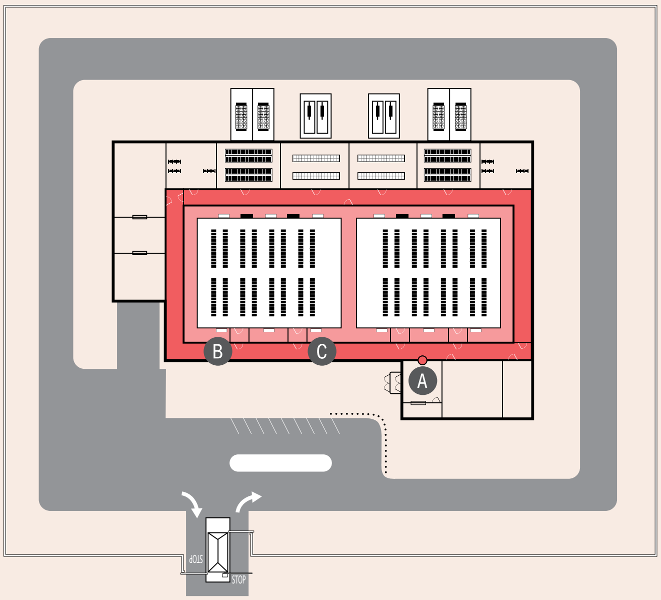

A Mantrap

B Secure Facility Entry

C Video Surveillance

Building Infrastructure Access Control

Once in the grey space, it is much more attainable for someone to penetrate the power and telecommunication systems that provide the backbone of a data centre, along with their backup systems. The security of a building's infrastructure is a paramount concern for any data centre and multiple security methods should be in place to protect it. A mechanical key cylinder is not difficult for a motivated and skilled malicious individual to defeat. A catastrophic incident can occur if an unauthorised person accesses important building infrastructure from the grey space.

Every door in a facility that does not have active access control is a possible problem point. Many of these doors can be accessed via the grey space by someone who passes through the main security entrance of a data centre if the grey space is not properly monitored.

Secure Doors and Locks

Doors that lead from the grey space into infrastructure or other areas should be considered for high-security doors and lock sets. Locks that engage not just at the door handle, but at the top and hinge side provide a greater level of protection should an opening of the cylinder lock occur.

Video

Video is an important tool in securing the grey space, and it can provide good, clear coverage of the hallways. A video strategy can include:

- High-resolution cameras ranging from HD 720p and 1080p to multi megapixel
- 360º degree or 180º degree cameras that can cover a significant area
- Cameras with speakers and microphones to enhance the ability of security personnel to judge potential situations and converse with subjects
- Motion detection that can alert to movement within sensitive areas or network cameras that come with intelligent motion detection that allows multiple specific points of interest to be masked off for motion detection versus the entire field of view
- Identification of who went into the chiller plant or the battery room and confirm their access.

Proper Emergency Plans

An additional area of concern is fire exits from all points of the facility. What happens if somebody pulls a fire alarm? Would a visitor suddenly have access to areas that would be off limits without the emergency? Would data centre employees know what to do?

LAYER 5: DATA CENTRE ROOM AND WHITE SPACE

As the importance of each layer of protection is explored, it becomes obvious that each individual layer of protection is extremely important. A threat can do great damage once inside a data centre cabinet, and it speaks to the importance of controlling access to the data centre room and white space. Someone intent upon server theft or a virus upload cannot achieve that if access to the floor is not granted.

If there is an overall strength in data centre security, it is likely the focus that is placed on preventing unauthorised people from entering the white space. Regardless of this focus, many data centres still have holes in their plans, processes or systems that render the white space not as secure as it could be.

Mantraps

Part of the security feature in and out of many data centre rooms is a mantrap. The mantrap requires authorised access into a small room, where both sides of the door must close and then another authorisation procedure takes place to clear an individual for entry.

Figure 15: Layer 5: Data Centre Room and White Space

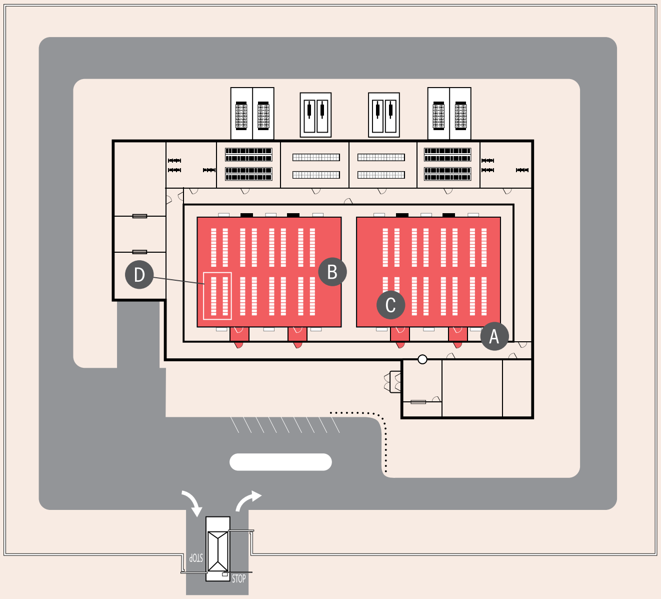

A Mantraps

B Video Surveillance

C Biometrics

Cages

Within the white space, a cage system can provide enhanced protection to separate and secure customer specific cabinets or network equipment. When using a cage system you should ensure your desired security measures for access control, intrusion detection and video surveillance can be supported.

Video

There are many uses of video throughout the different layers, but specific to the data centre room and white space, there are specific cameras that can be used. A 360° degree can be used in the middle of hot or cold aisle. Based on how it integrates with the cabinet access control, it digitally zooms into a preset position when a cabinet door is opened.

A pan-tilt zoom camera can be used at the end of a row as a fixed-position camera. It is stationary and looks down the row in its home position. Once a cabinet is opened, it will automatically optically zoom in and record the person going into a cabinet based on a preset position, and then zoom back into its fixed home position.

Low-light cameras can also be used in areas where data centre floor lights are dim when no motion is detected to save on energy costs. Advances in technology now allow these low-light cameras to provide colour images in dark environments better than the human eye can see and provide more lifelike colours in low-light conditions.

Thermal cameras can be used in the data centre to provide protection in the event a malicious individual intentionally turns out the lights in an equipment storage area or on the data centre floor. Thermal cameras operate in complete darkness by providing a heat signature that always radiates from any object or person. Thermal cameras can also be used to find hotspots caused by overheating servers, uninterruptible power supplies, cooling systems and power distribution units inside data centres.

VIDEO AS A BUSINESS SERVICE:

Cameras in the white space offer capabilities beyond just protecting people and assets. A high-resolution video camera can allow a user to visually identify equipment to see what's going wrong. This provides a forensics capability for reviewing what has taken place in the data centre. Cameras can assist a network operations centre by visually identifying and reading equipment displays and LED light patterns, allowing for more efficient monitoring and maintenance. Providing video surveillance as a business service is an added benefit of security video monitoring.

Biometrics' Triple-Factor Advantage over Dual-Factor Authentication

Access control into the data centre white space and between the white space and the data centre cabinet is of primary importance. Some data centres, even those with otherwise high security protocols, still rely upon badges that require dual authentication even in protecting the inner core of a data centre. Even though badges are still acceptable in the outer portions of the protective rings, best practices dictate the use of biometric identification closer to the data suite.

What are the Three Key Advantages of Biometrics?

Proof: Biometrics, whether via fingerprint, iris scan or otherwise, positively proves an individual's identity. Keycards with dual authentication provides some semblance of security, but as anyone who has ever used a friend's or spouse's debit card can attest, possession of a card and knowledge of a PIN code is simply not proof of an individual's identity.

Convenience: Biometric systems avoid situations where a card has been lost or misplaced, which requires time, energy, effort and expense to deal with. A fingerprint or iris doesn't need replacing, making a more efficient and less time-consuming solution.

Cost: Even though a keycard reader is initially less expensive than a biometric reader system, the management costs of a system add up quickly. The initial issuance of such keycards can become costly, as is the replacement of misplaced keycards or those left at home.

The cost advantage of biometrics is especially true in a facility that requires authorisation for a large number of people to enter. The cost of biometric systems should not be a barrier for implementation; in the long run, it is frequently less expensive, as there are few costs beyond that of the reader. Even though a biometric reader is more expensive than a keycard reader, there are not the initial and ongoing costs of the keycards themselves, which typically cost around $3 a piece. If over time the number of employees and visitors requiring keycards becomes a fairly large number, the keycard reader system will frequently be more expensive than the biometric system.

LAYER 6: DATA CENTRE CABINET

The core of the data centre is the IT infrastructure housed within a data centre cabinet. As a rule, these cabinets are remarkably insecure.

Cabinet access control is being implemented into more compliance standards. The ability to positively identify who is getting into those cabinets is now required for TIA Tier 3 or Tier 4 data centres, and PCI DSS is one of the many certifications that is continually raising the bar on standards. The vast majority of data centres lack proper cabinet access control. This is a significant issue.

Figure 16: Layer 6: Data Centre Cabinet

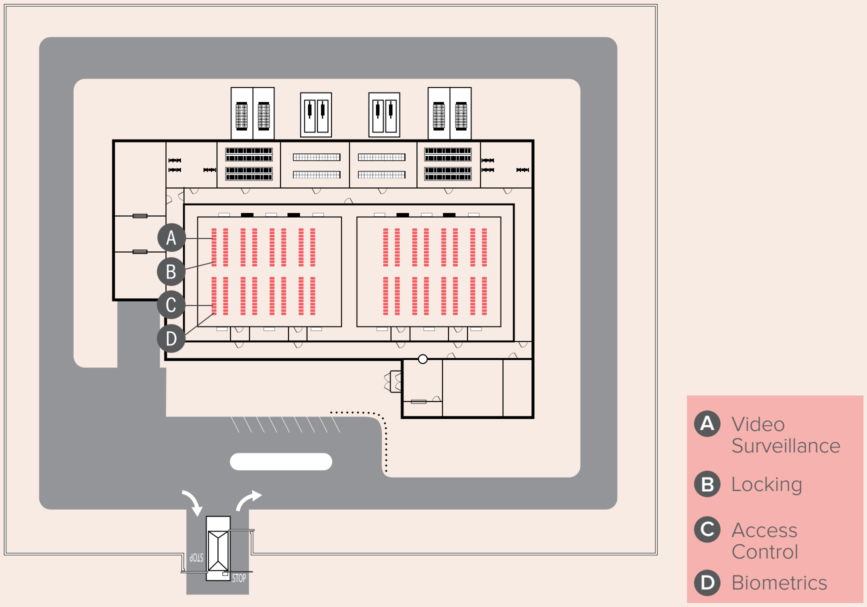

A Key Problem

A common problem is that many cabinets today are still being shipped with a single master key. Surprisingly, each override key can open all other cabinets of the same model in any data centre anywhere. As a result, the majority of cabinets in use by data centres today can be opened by anyone who possesses one of the override keys.

This situation creates a significant gap in the security of many data centres, where somebody with a manufacturers' override key can quickly and easily access a cabinet. The implications of this ready access to a data centre cabinet have significant consequences, including the following:

- **Server theft:** Server theft is surprisingly common, and according to DCD the average price of a data centre server in 2014 was $7k but the disruption caused is far more
- **Storage theft:** The average price of a storage system in 2014 was $12k according to DCD – but the loss of data could be very significant
- **Virus upload:** A similar concern is the ability to upload a virus when given access to the cabinet.
- **Interruption with power or connectivity:** No matter how redundant a data centre, if a server has wires clipped, there will be a period of downtime until the problem can be diagnosed or fixed.

This is obviously a significant issue for all data centres, but particularly in data centres with a lot of visitors and foot traffic through the facility. It should be a particular concern of multitenant data centres where a customer's representative (a technician or contractor) can be legitimately cleared to be inside one customer's cabinets in easy reach of other customers' cabinets. If they possess the proper manufacturer's override key, and are not properly monitored, they would have easy access to other customers' servers and infrastructure.

In many multitenant data centres, a customer's IT infrastructure can easily be accessed by the customer next door, but this isn't just a multitenant data centre problem. Few data centres do not at least occasionally have outside contractors and others doing work on the data centre floor. Whether it's a private data centre or a co-location environment, a contractor or employee with malicious intent can quickly and easily access a data centre cabinet and do significant and expensive damage in a very short period of time.

Co-location facilities also frequently have significant foot traffic with people on sales tours and many try to host events for technologists and other business leaders in an effort to increase awareness of the facilities. Those additional visitors create additional need for a cabinet access control solution.

Key Restrictions

If circumstances dictate the use of a cylindrical key access, some steps can be taken to reduce the risk.

- Understand that “do not duplicate” is meaningless in an era of self-serve key replication machines in major retailers.
- Some manufacturers have key blanks that are stamped with end-user ID numbers, which enable the manufacturer to identify the source of any unauthorised keys.
- Always be aware that a key system with no other cabinet access control is sub-optimal and is a significant security risk.

Best Practice Support

Basic security best practices can mitigate the risk somewhat. Policies that can reduce the risk include strict tracking of all those given access to the white space, requiring escorts for visitors, not allowing individuals to be alone on the data centre floor, and requiring two or more team members together whenever they are in the white space. While serving as deterrents, none of these methods is completely foolproof.

A plan for cabinet access control is also needed for private data centres as well. It is not rare for a data centre to have instances where staffers have reason to be in the data centre space regularly but not necessarily in the cabinets.

Cabinet Access Control Products

In recent years, a number of access control products that attempt to solve the issue of cabinet-level security have come to market. Typically, these products use dual (or greater) factor authentication or biometrics in an effort to make sure that only those with proper clearance have access to a cabinet.

An Audit Trail

Access control products must create complete access records to support a true level of security and compliance via a verifiable audit trail. They can also limit access to specific times for specific individuals. For example, a contractor that works on a server on Tuesday mornings can have access restricted to 8 a.m. to noon on Tuesdays.

Interoperability and Video

Cabinet access control is improved by including video integrated for visual confirmation as well as access control. Integrating video surveillance and access control creates a verifiable audit trail of those who are within a cabinet. Many customers are insisting upon interoperability from manufactures and security integrators.

End-of-Row or Integrated Cabinet Reader

Cabinet access can be controlled by either an integrated cabinet reader at the handle or by an access device at the end of a row. As a general rule, end-of-row solutions are more affordable but do not provide as detailed an audit trail on a cabinet-by-cabinet basis as an integrated cabinet reader provides. Every application will be unique. In addition to budgetary factors, the system size, future scalability, needs, security requirements and policy and procedures should all be factors in the decision.

Key Override

In case of an emergency situation or power loss, a person in authority may have need for a key that bypasses all other security measures. The presence of a manual override key obviously requires significant management, including the following best practices:

1. The key system must be restricted by a utility patent to prevent unauthorised key duplication.
2. Override keys must be tightly managed to make sure they are only used by authorised users.
3. Keys should be locked in a safe or emergency key cabinet that has restricted access.
4. Keys should be logged in and out – and only by authorised users.
5. Cabinet access control system should be able to detect and log any use of bypass keys.
6. Keys should only be used in emergency/power loss situations.

Proper Procedures and Cabinet Access Control are Keys

The risk of penetration to a data centre cabinet is obvious, and the large numbers of manufacturer keys in circulation increases the risk. A combination of a solid plan, proper security procedures, a dedicated adherence to those procedures, and new tools and technologies are necessary to properly secure a cabinet.

CONCLUSION

Creating a security strategy isn't a one-size-fits-all proposition. By using the "Defence in Depth" approach outlined in this report, data centres can create a secure environment by using multiple security strategies, policies and protocols. Whether it's a military-grade installation, a multitenant facility or even an on-site data centre of a local business, the six layers help create a Defence in Depth posture, tailored to each individual's need, to protect critical infrastructure, deter potential threats and promote the achievement and maintenance of necessary compliance objectives.

Data breaches are becoming nearly universal, and data centres of all sizes need to be aware of the risks and prevention strategies. As more data are created in all areas of personal and professional life, the information stored in the cabinets becomes increasingly valuable. Evaluating and budgeting for logical and physical security is the only rational step.

Never before has a data centre breach been more damaging to the future of an organisation. As the risk has risen, a growing array of sophisticated threats continues to emerge designed to penetrate data centre defences. It's important to take the first steps now to develop, implement and execute a suitable data centre physical security strategy.

GLOBAL TECHNOLOGY BRIEFING

NETWORK MIGRATION BEST PRACTICES

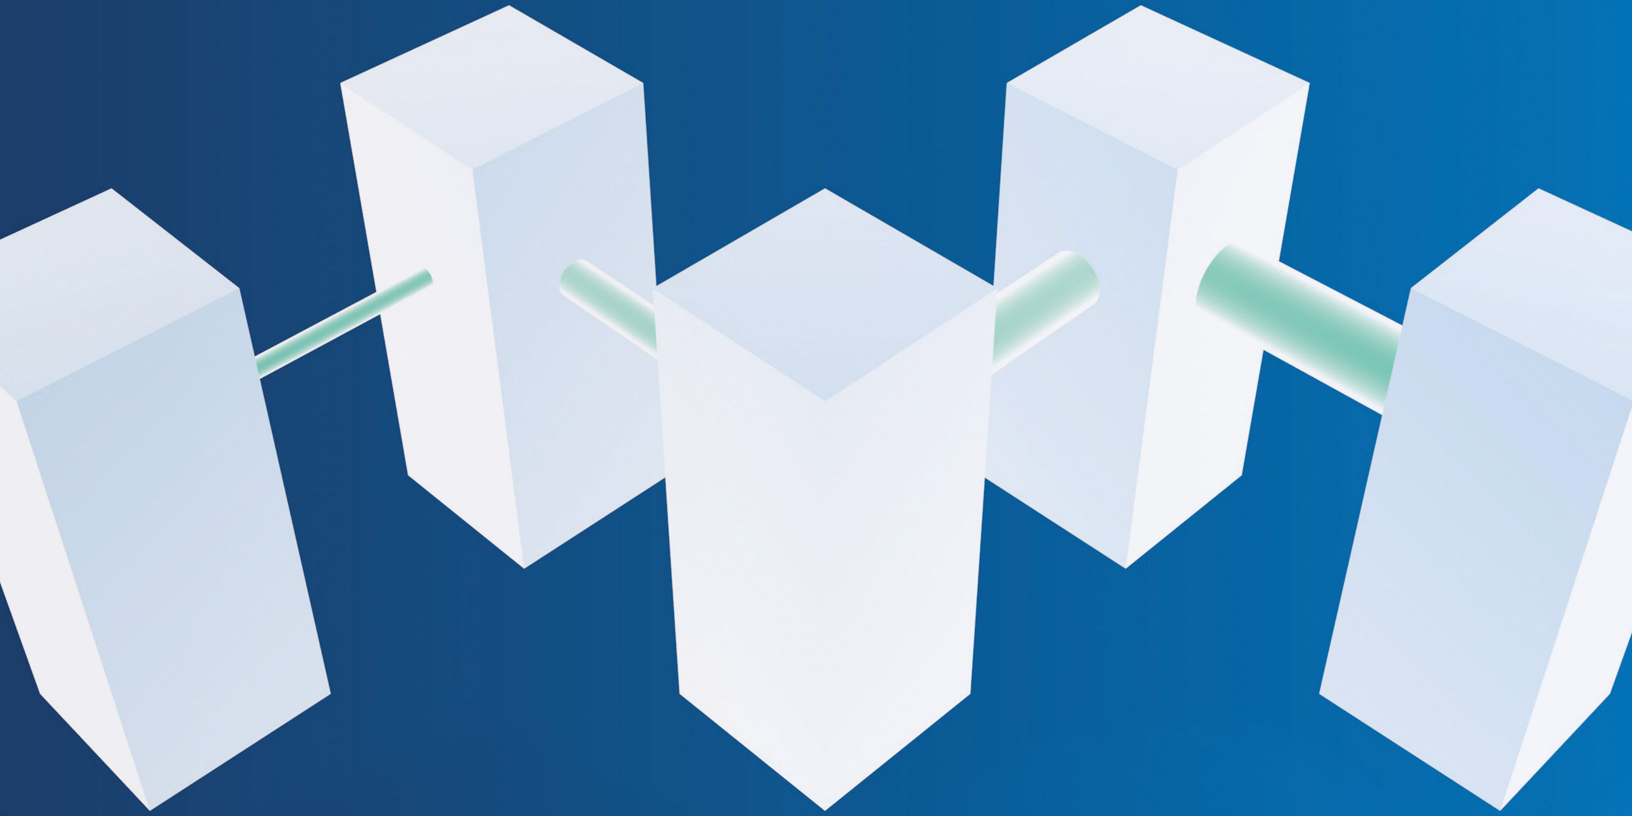

EXECUTIVE SUMMARY

With the ever-increasing demands of data-intensive applications and the immense storage requirements of modern data centres, navigating through the myriad of network-infrastructure choices needed to support high-speed data rates ranging from 10Gbps to 40Gbps can be a challenging undertaking. For this and several other practical reasons, it’s vital to have a well-planned network-migration strategy that will make sure a data centre can support LAN and SAN computing requirements now and in the future.

During the original Levels program in the 1990s, which led to the development of cabling performance standards such as current Category 5e and Category 6, Anixter recognised the challenges faced by designers of low-voltage communications cabling systems when choosing the type and grade of cabling that would best support the emerging Ethernet protocol being deployed across enterprise networks.

Today, the industry has evolved from 10Mbps systems to industry standards that support data rates up to 100Gbps in data centre environments. Yet, media selection, cabling architecture and cable management choices remain complex. As data centre planners try to control costs by maximising floor space, choosing the right high-density cabling architectures and computing equipment is mission-critical. Rack space is quickly migrating from 40RU to 48RU in server cabinets, while data speeds are moving from 10Gbps to 40Gbps in the server and eventually to 100Gbps in the data centre backbone.

NETWORK MIGRATION BEST PRACTICES

Faster, denser technology is driving costs, and the right high-performance cabling is needed to provide stability in the data centre. Unsuitable infrastructure can become an expensive problem, delaying necessary upgrades and creating other potential obstacles needed to stay competitive.

This report contains the four best practices needed to achieve a high-performance, future-ready structured cabling solution for a data centre.

Network Migration Best Practices:

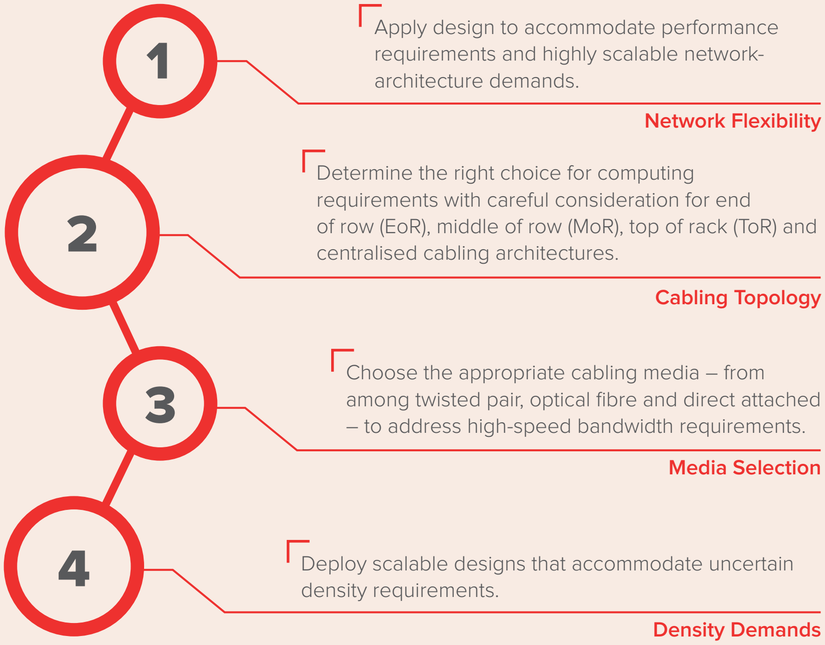

1 Apply design to accommodate performance requirements and highly scalable network-architecture demands.

Network Flexibility

2 Determine the right choice for computing requirements with careful consideration for end of row (EoR), middle of row (MoR), top of rack (ToR) and centralised cabling architectures.

Cabling Topology

3 Choose the appropriate cabling media – from among twisted pair, optical fibre and direct attached – to address high-speed bandwidth requirements.

Media Selection

4 Deploy scalable designs that accommodate uncertain density requirements.

Density Demands

INTRODUCTION

The Evolution of Data Centre Cabling

Over the past few years, data centre design and cabling infrastructure architecture have evolved as needs and technologies have changed, as shown if Figure 1. In the past, the data centre manager relied on trial and error or on solutions that had previously worked. Planning today's data centre, however, requires a more rigorous and methodical approach, because of ultra-rapid change.

The current state of data centre cabling is highly dependent on high-performance connectivity. Data centre cabling must be able to accommodate today's large number of diverse bandwidth-intensive devices, including virtualisation appliances, backup devices, bladed servers and clustered storage systems. Of course, all these devices are interconnected by a network infrastructure and rely heavily on high-performance cabling to meet the growing bandwidth and storage needs.

The future state of data centre cabling is difficult to predict, though it can be assumed that the current state of bandwidth growth will only continue. Data centre transformation and technology overall have moved much faster than anyone could have imagined just a decade ago, and advances promise to continue over the next few decades. There's no doubt that data centres will depend on improvements to cabling and higher performance standards to make sure there is an adequate flow of data at speeds that meet anticipated leaps in bandwidth.

NETWORK MIGRATION BEST PRACTICES

The prevailing wisdom is that a simple, cost-effective migration path can be achieved by installing a structured cabling system that can support future 40/100 Gigabit Ethernet networking demands. The challenge for data centre builders today is to make sure that the cabling infrastructure can keep pace with future needs.

Simplifying network migration hinges on planning that strikes a balance between capacity space, power consumption, flexibility, cost and speed. Although striking that balance requires highly technical considerations, the first step in reducing complexity should be to examine the challenges and concerns before exploring best practices.

Figure 1: Changes in Data Centre Traffic

Device	Traffic Multiplier
Tablet	1.1
64-bit laptop	1.9
Internet Enabled TV	2.9
Gaming Console	3.0
Internet 3D TV	3.2

Source: IEEE 802.3

Figure 2: The Server Roadmap

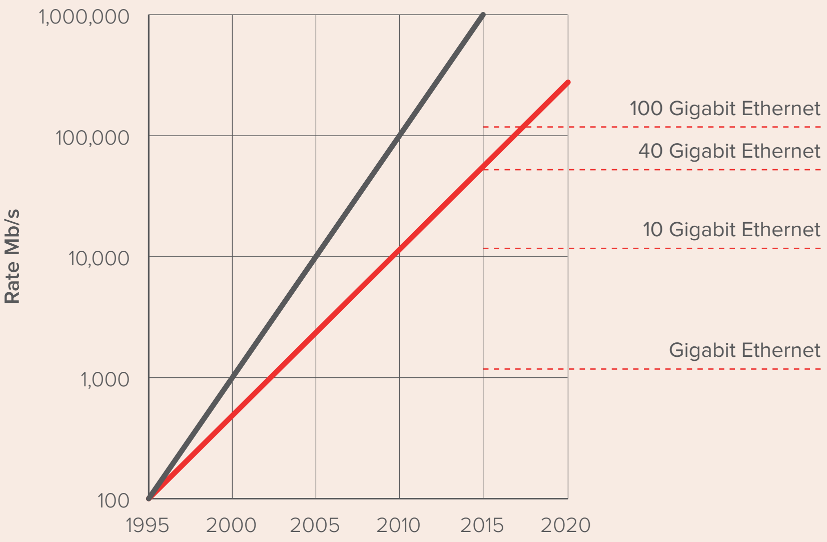

Key

—	Core Networking Doubling ≈18 mos	🟥	Server I/O Doubling ≈24 mos

Server Upgrade Path	• 2014: 40 GbE	• 2017: 100 GbE
Blade Servers	• 802.3ba: 10 GbE to 40 GbE	• 802.3bj: 40 GbE to 100 GbE
Other Future Server I/Os	• 40GBASE-T	• 100 GbE over MMF

Source: IEEE 802.3

TOP CHALLENGES AND CONCERNS

Network Migration Challenges

Amortisation of cabling investment

Research shows that data centre amortisation periods run on average for 10 to 15 years. Server amortisations are typically in the three-year range. Servers are purchased three to five times during the life of a data centre, so, when amortised properly, they continue to dominate the cost equation.

When looking at data centre cabling costs, one should consider not only the initial costs but long-term costs, too. Understanding the full life cycle and industry trends fits into the equation, too. It's also important to note that cabling represents only 2-3% of the initial network hardware investment.

Cabling is also expected to outlive most network components, but it might be the most difficult and cost-intensive component of a network to replace.

Data centre cabling systems that turn out to have a shortened life span requiring replacement sooner than expected might negatively affect an entire IT budget for years to come.

Infrastructure complexity

Data centre complexity is another significant obstacle to smooth migration. Even though virtual machines let numerous applications run on larger machines, many of these hungry applications have outgrown servers and now run across multitudes of nodes, on-premises or in the cloud. Adding to the intricacy is cloud integration manufacturers that deploy products that increase complexity by layering on more management components that need constant monitoring.

Bottom line, the data centre is coming under increasing pressure to perform from every direction. Infrastructure complexity is an increasing challenge to a well-planned migration that can help deliver the availability, capacity and efficiency that businesses demand of today's data centre.

Pace of innovation adoption

Although there is a constant drive for rapid innovation around the data centre, a lot of what is driving the tremendous advances is on the device side or the Internet of Things (IoT). More devices and people are coming online, which has led to innovation adoption in the data centre. Yet the slow pace of that adoption can be a challenge to migration.

According to Gartner, consumer applications will drive the number of connected things through 2020, while enterprises will account for most of the revenue (see figure 3). Gartner estimates that the consumer segment will account for 2.9 billion connected things by 2015 and surpass the 13 billion mark in 2020, while the automotive sector will show the highest growth rate.

This new approach to computing puts a tremendous emphasis on the back-end data centre services rather than the capability of the end-user's device. Many existing data centres' entire supporting infrastructure may be ill prepared to handle this dramatic shift. In addition, the server continues to be redefined, which greatly affects the hardware and associated cabling choice for network migration.

According to Cisco, global IP data centre traffic will grow 23% annually, to 8.7 zettabytes, by 2018. Over three-quarters of the traffic will be within the data centre, generated by exchanges between servers or between servers and storage. With this in mind, it is paramount that data centres leverage infrastructure technologies that will optimise data centre traffic.

Figure 3: Internet of things. Units (millions) Installed base by Category

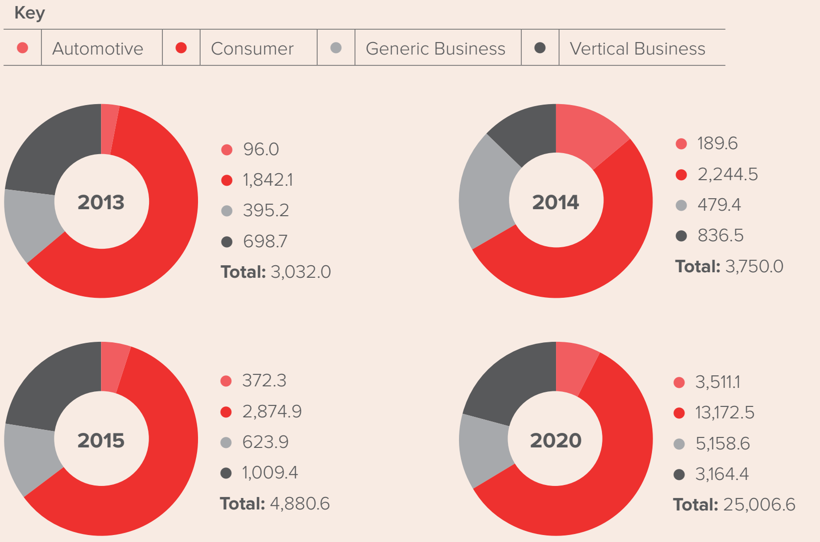

Source: Gartner (November 2014)

Figure 4: Global Data Centre Traffic by Destination

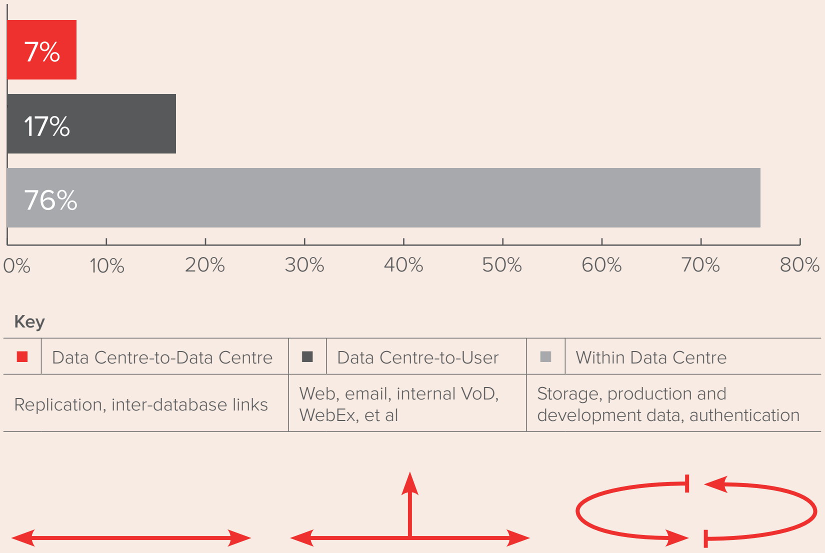

Source: Cisco

Speed of deployment

A significant challenge in network migration is the speed of deployment and the potential interruption or slowdown of business as usual. A data centre is a mission-critical delivery platform for business, so migrating a data centre can be a complex, large-scale effort with significant investment and risk unless carefully planned and managed.

Business continuity must be assured by maintaining application performance and availability. Operational service levels must also be delivered throughout physical and platform moves.

> There's little or no margin for error, so it's crucial to prevent costly, unplanned service outages or performance degradation with rapid deployment.

To keep up with the demands of critical business operations and urgent demand for performance, pre-terminated cabling solutions have been designed with features of upgradability, modularity and scalability. In addition to rapid-deployment product features, a business should consider the importance of project planning, controlled repeatability and key supply-chain partnerships.

Restrictions of legacy systems

Legacy systems can be an obstacle to a smooth network migration. Planning and decisions regarding exactly how legacy systems and their data will be migrated must take into consideration systems' inherent restrictions, especially as they relate to the functionality and innovations of the new data centre. In some cases, the retirement of legacy systems may be in order.

Concerns for data centre investment

In addition to these significant challenges, there are several common concerns regarding data centre investments:

- **Data centre connectivity –** all businesses rely on distributing massive amounts of data across multiple locations, so connectivity plays a major role in buying decisions.
- **Availability –** high availability and sustained continuous operation are vital if business continuity is to be provided.
- **Resilience –** confidence is needed that an entire data centre will continue to operate even when there has been an equipment failure, power outage or other disruption.
- **Control over the facility –** monitoring and control are critical elements of maintaining maximum availability for critical operations.
- **Access to the cloud –** determining how cloud computing figures into a new data centre will affect some of the investment decisions.

NETWORK MIGRATION BEST PRACTICES

In the 2014 DatacenterDynamics Industry Census, 20.7% of European data centre end-users expressed significant levels of concern about the potential impact of insufficient access to networking on their operations over the next 12-18 months. An additional 35.6% of end-users expressed moderate levels of concern. Concern about access to networking was almost has high as the concern expressed about reduced capital and operating budgets, as well as power cost and skills shortages.

Despite the range of challenges and concerns facing data centre operators, executing the smooth and seamless network migration of a data centre can be done, starting with its new vital arteries: high-performance cabling.

Figure 5: Key Data Centre Operator Concerns (proportion of DCD Census respondents expressing significant levels of concern about the impact of specific issues on their operations over the next 12-18 months)

	2012	2013	2014
Increasing labour costs	14.1%	12.9%	18.6%
Increasing energy costs	28.0%	23.9%	25.9%
Shortage of power-blackouts, brownouts	18.2%	18.6%	20.7%
Lack of suitable local outsourcing facilities	9.1%	9.4%	14.5%
Lack of suitable real estate for development	12.0%	11.5%	15.3%
Reduced availability of capital	25.8%	23.7%	28.0%
Reduced operating budgets	28.0%	27.4%	31.3%
Shortage of staff with necessary skills	19.8%	21.3%	31.2%
Insufficient networking access	-	-	23.1%

Source: DCD Intelligence

DEFINING THE FOUR BEST PRACTICES

Next-generation enterprise applications are emerging that require the support of Ethernet ecosystems that go beyond current speeds. That makes it even more important to have standards that allow the full range of manufacturers to be able to interoperate with the current infrastructure among the ever-increasing volumes of data traffic brought on by everything from wireless access points to hyperscale computing growth.

Structured cabling systems provide a standards-based approach for connecting and administering data centre computing hardware. By making sure there is manufacturer interoperability and backward compatibility with legacy systems, a structured cabling approach maximises both the utility and useful life of the data centre cabling infrastructure.

Standards

ANSI/TIA-942-A is the most recognisable data centre standard in North America today. It was revised and published in August 2012 and contains best-practice recommendations for all elements of data centre design.

Figure 6 from ANSI/TIA-942-A represents the functional elements of a generic data centre cabling system. It depicts the relationship between the elements and how they are configured to create the total system.

There are eight basic elements of a data centre cabling system structure, as shown in figure 6 below:

1. Horizontal cabling (Cabling Subsystem 1).
2. Backbone cabling (Cabling Subsystem 2 and Cabling Subsystem 3).
3. Cross-connect in the entrance room or main distribution area (Distributor C, Distributor B or Distributor A).
4. Main cross-connect (MC) in the main distribution area (Distributor C or could also be Distributor B or Distributor A).
5. Optional intermediate cross-connect (IC) in the intermediate distribution area or main distribution area.
6. Horizontal cross-connect (HC) in the telecommunications room, horizontal distribution area or main distribution area (Distributor A or could also be Distributor B or Distributor C).
7. Consolidation point in the zone distribution area (optional).
8. Equipment outlet (EO) located in the equipment distribution area or zone distribution area.

Figure 6: Data Centre Cabling System Infrastructure

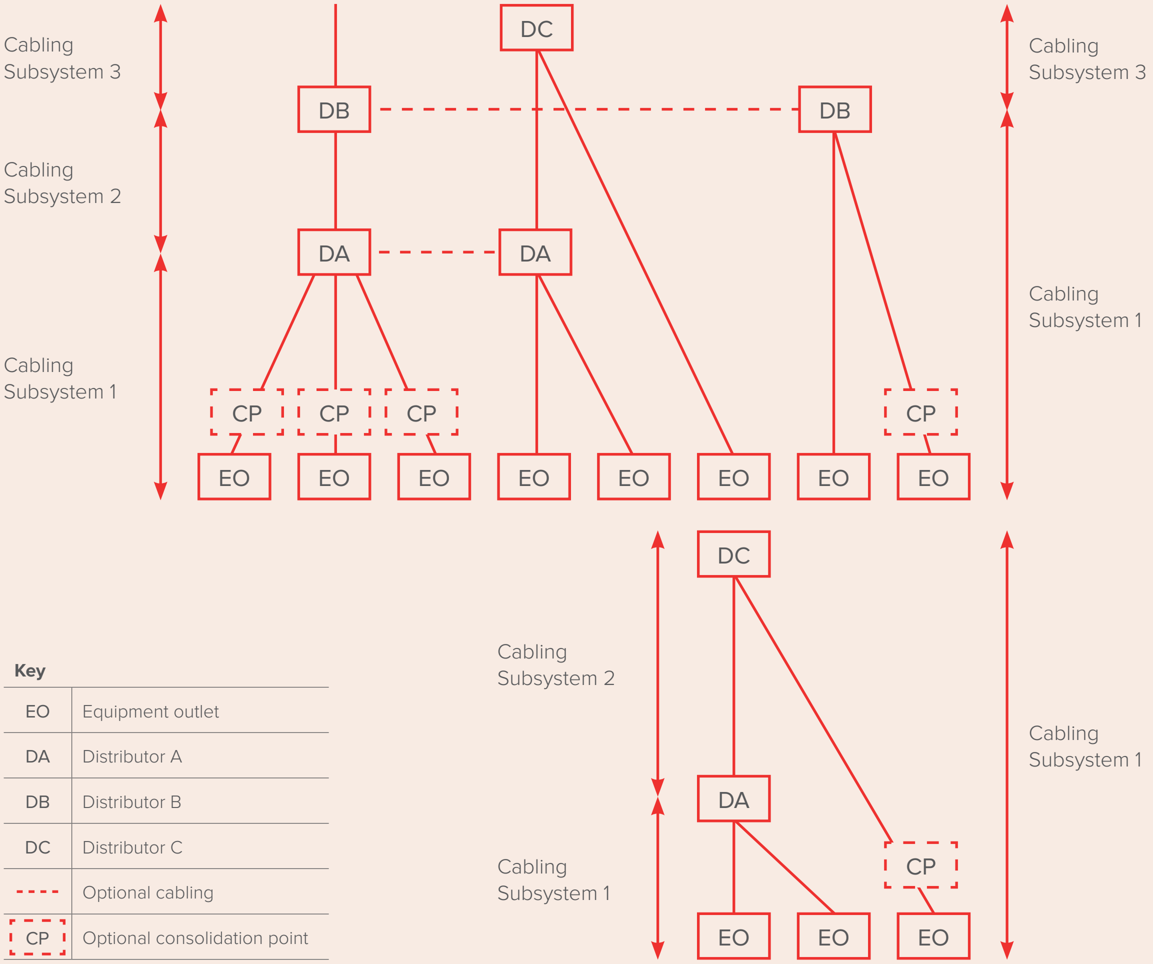

Key

EO	Equipment outlet
DA	Distributor A
DB	Distributor B
DC	Distributor C
- - - -	Optional cabling
CP	Optional consolidation point

Source: ANSI/TIA-942-A

The ISO/IEC FDIS 24764 standard only covers the cabling system within the data centre. This standard is widely referred to in European data centre design processes.

Generic cabling systems in data centres contain up to three cabling subsystems: network access cabling, main distribution cabling and zone distribution cabling. Where present within the premises, a distributor in accordance with ISO/IEC 11801 is connected to the generic cabling within the data centre using the network access cabling.

Data centre network needs
When designing your cabling infrastructure for migration, it's important to address three areas that are essential to meeting current and future data centre objectives.

- ***Scalability*** *– cabling choices must allow businesses to scale up to meet growth quickly and cost-effectively.*
- ***Reliability*** *– without reliability, availability is a moot point, so your data centre must be able to operate despite outages and other potential interruptions.*
- ***Flexibility*** *– A new data centre must have the ability to adapt over time to accommodate unforeseen changes without disruption or a significant reinvestment.*

Figure 7: ISO/IEC 24764 (2010) - General Structure

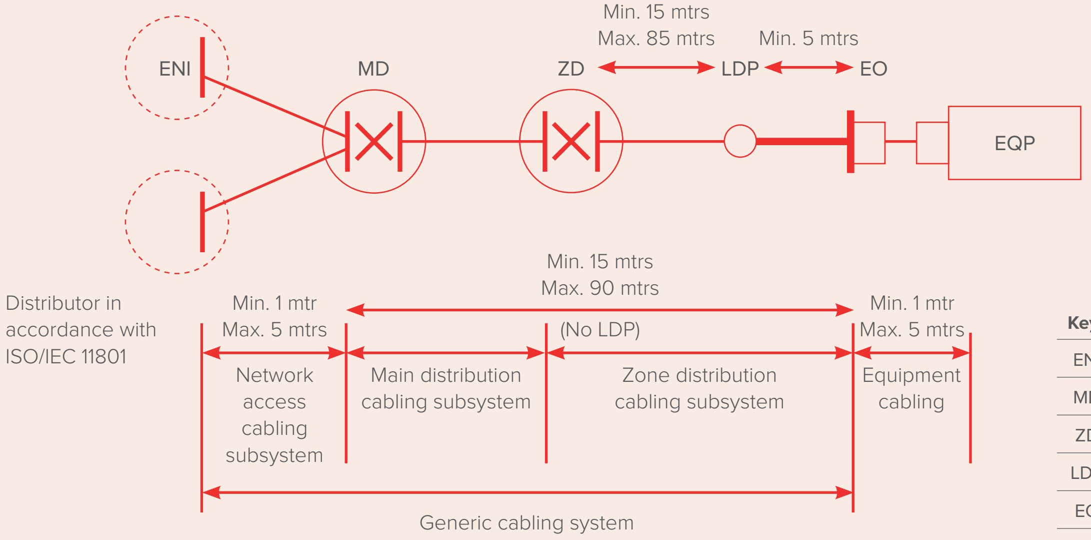

Key

ENI	External network interface
MD	Main distributor
ZD	Zone distributor
LDP	Local distribution point
EO	Equipment outlet

Source: ISO/IEC 24764 (2010)

BEST PRACTICE 1 OF 4

CREATE NETWORK FLEXIBILITY

1

Building a long-term cabling plan

As the worldwide pool of data grows, corporations are increasingly consolidating and centralising data centre operations to save on operation and maintenance costs.

With more than 3 trillion bits of data created every second, the world's appetite for more feature-rich information keeps growing, and with it comes the need for data centres to store and process that information.

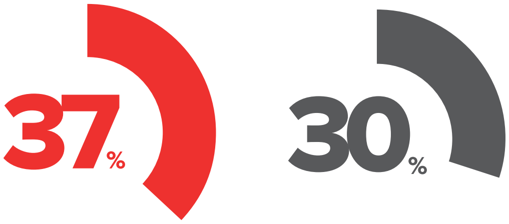

Thirty-seven percent of data centre operators responding to DCD Intelligence's 2014 Industry Census indicated that they were involved in data centre consolidation initiatives, while 30% indicated that they were involved in migration projects.

The subject of data centre consolidation initiatives include real estate, taxes, utilities and other physical support groups, such as security, electrical and mechanical systems. The new high-density data centres save on physical costs by reducing equipment and floor space for servers in remote offices, cutting software license and distribution fees, and reducing operating expenses.

NETWORK MIGRATION BEST PRACTICES

One consolidation strategy is server virtualisation; this harnesses the computing power of multiple servers into logical groupings, or virtual servers, running concurrently on the corporate network. The DatacenterDynamics Industry Census shows that virtualisation has consistently remained a key investment driver for European data centre operators.

Virtualisation essentially breaks the link of the physical server and the software applications that run on it. Given that software applications run on virtual machines, the result is greater computing and power efficiencies by maximising the use of the physical servers. Virtualisation provides network administrators with essential flexibility and agility in managing data centre environments while delivering rapid deployment, quick change adoption and flexible disaster recovery.

Virtualisation requires an improvement of network bandwidth and latency performance. High-bandwidth technologies, such as 10 Gigabit Ethernet server interfaces and 40/100 Gigabit Ethernet switches used in the data centre backbone, can alleviate potential bottlenecks from aggregating computing resources using virtualised servers and storage platforms.

Figure 8: Global Data Centre Operator Investment Avenues, 2014 -2015 (proportion of DCD Census respondents identifying specific investment avenues for a particular year)

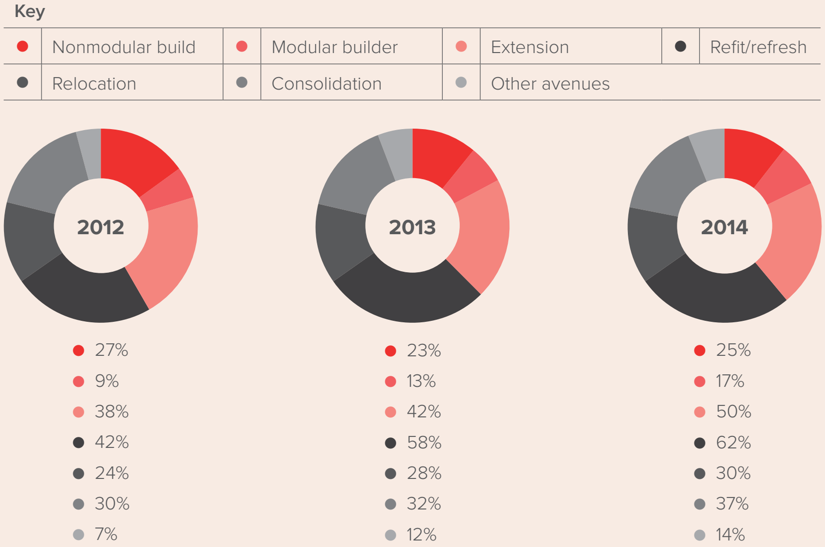

Source: DCD Intelligence

Figure 9: Global Data Centre Operator Investment Drivers, 2014 -2015 (proportion of DCD Census respondents identifying specific investment drivers for a particular year)

	2014	2015
Increased IT capacity requirements	39.2%	48.2%
To reduce operating costs	33.1%	43.5%
Improve network performance	31.9%	35.5%
End of facility life	31.9%	30.5%
Improve security	31.8%	34.7%
Enable virtualisation/cloud computing	30.7%	38.0%
Increase redundancy	30.6%	36.1%
Improve space use	29.4%	34.9%
Changing corporate and client requirements	25.6%	32.5%
Increase power into facility	24.7%	33.0%
Greener and more sustainable	20.3%	34.3%
Meet legislative or accreditation requirements	20.1%	24.5%
Support the requirements of big data	15.8%	27.9%
Other reasons	10.6%	12.5%
Increase competitive differentiation	8.0%	10.5%
Access software-defined utilities and "as a service" delivery models	7.3%	10.5%
Attract different client groups	7.3%	9.5%

Source: DCD Intelligence

Data rates

Within the data centre, there are distinct network subsystems and a cluster fabric, as shown in figure 10.

- **Data centre LAN –** connects application, Web and database servers to corporate LAN
- **SAN –** storage area network
- **NAS –** network attached storage
- **Cluster fabric –** enables high-performance computing (HPC)

In most cases within the data centre LAN, the network server connections are Ethernet-based. The server connections typically will use 1000BASE-T technology operating over a balanced twisted-pair cabling system. However, with the advent of server virtualisation, 10 Gigabit Ethernet technologies are becoming more commonplace.

The good news is that faster connection speeds are being developed. These include 100 Gigabit Ethernet and 40 Gigabit Ethernet networking technologies for transmitting Ethernet frames at rates of 100Gbps and 40Gbps, respectively. The technology was first defined by the IEEE 802.3ba-2010 standard.

Figure 10: Data Centre Subsystems and Cluster Fabric

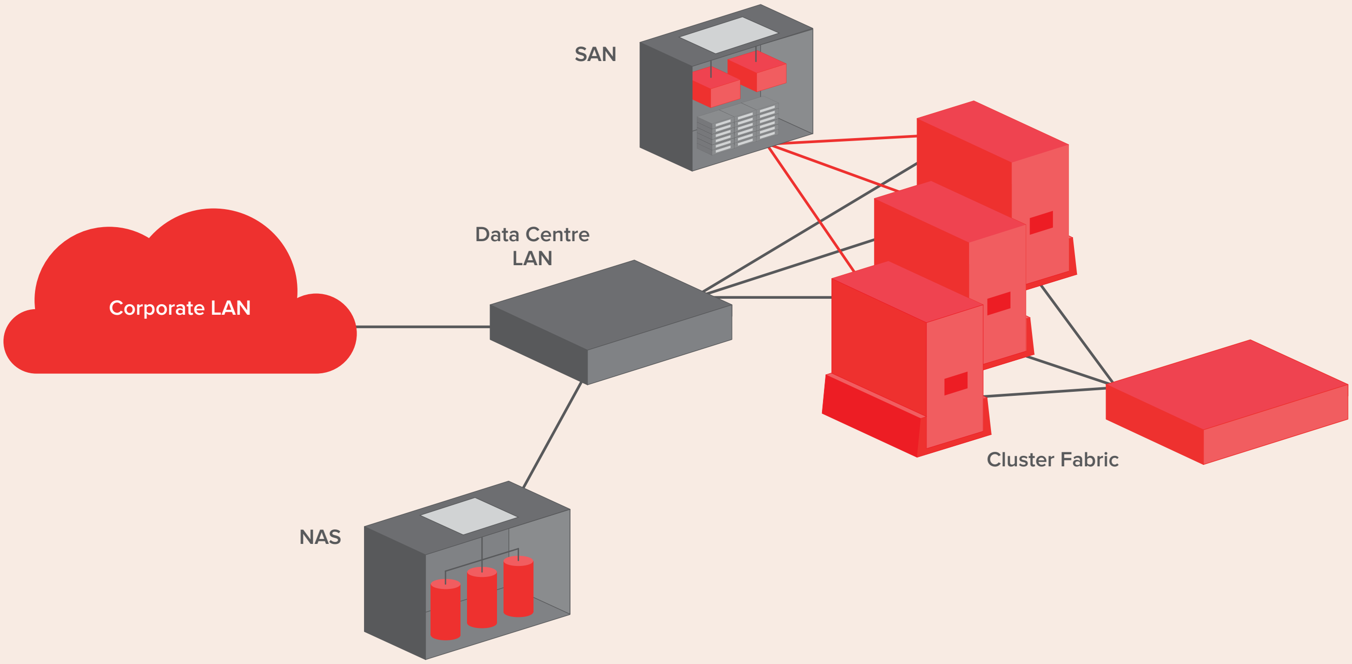

The data centre SAN uses the Fibre Channel, or FC, which is a high-speed network technology (commonly running at 2Gbps, 4Gbps, 8Gbps and 16Gbps rates) primarily used to connect computer data storage. Fibre Channel is standardised in the T11 Technical Committee of the International Committee for Information Technology Standards (INCITS), an American National Standards Institute (ANSI)-accredited standards committee. Despite its name, Fibre Channel signaling can operate over copper twisted-pair or coaxial cables, in addition to fibre-optic cables.

Figure 11: Fibre Channel Variants

NAME	1GFC	2GFC	4GFC	8GFC
Line-rate (gigabaud)	1.0625	2.125	4.25	8.5
Line coding	8b10b	8b10b	8b10b	8b10b
Nominal Throughput full duplex; MB/s	200	400	800	1,600
Net throughput per direction; MB/s	99.6	199	398	797
Efficiency	78.6%	78.6%	78.6%	78.6%
Availability	1997	2001	2004	2005

NAME	10GFC	16GFC	32GFC	128GFC
Line-rate (gigabaud)	10.52	14.025	28.05	4x28.05
Line coding	64b66b	64b66b	-	-
Nominal Throughput full duplex; MB/s	2,400	3,200	6,400	25,600
Net throughput per direction; MB/s	1,195	1,593	-	-
Efficiency	95.3%	95.3%	-	-
Availability	2008	2011	2016 (projected)	2016 (projected)

BEST PRACTICE 2 OF 4

USE A FUTURE-READY CABLING TOPOLOGY

2

ANSI/TIA-942-A cabling system

Building upon the earlier definition of the generic cabling system definition of the ANSI/TIA-942-A standard, the basic cabling topology should be laid out as follows.

Principle spaces for the data centre network

Figure 12 illustrates the hierarchical nature of data centre infrastructure design. It does not represent the topological design. It is recommended that the infrastructure distribution emanate from the centre of the room, rather than following a typical legacy topology of locating the “network row” on the perimeter of the computer space.

ENTRANCE ROOM (ER) – carrier circuits and demarcation equipment located here; typically separate from the computer room for security reasons.

MAIN DISTRIBUTION AREA (MDA) – where core layer equipment, such as routers, LAN/SAN switches, PBXs and MUXs, are located.

HORIZONTAL DISTRIBUTION AREA (HDA) – houses aggregation layer equipment, such as LAN/SAN/KVM switches.

EQUIPMENT DISTRIBUTION AREA (EDA) – where access layer equipment, such as LAN/SAN/KVM switches and servers, are located.

ZONE DISTRIBUTION AREA (ZDA) – a consolidation point or other intermediate connection point.

Figure 12: Hierarchy of Data Centre Infrastructure Design

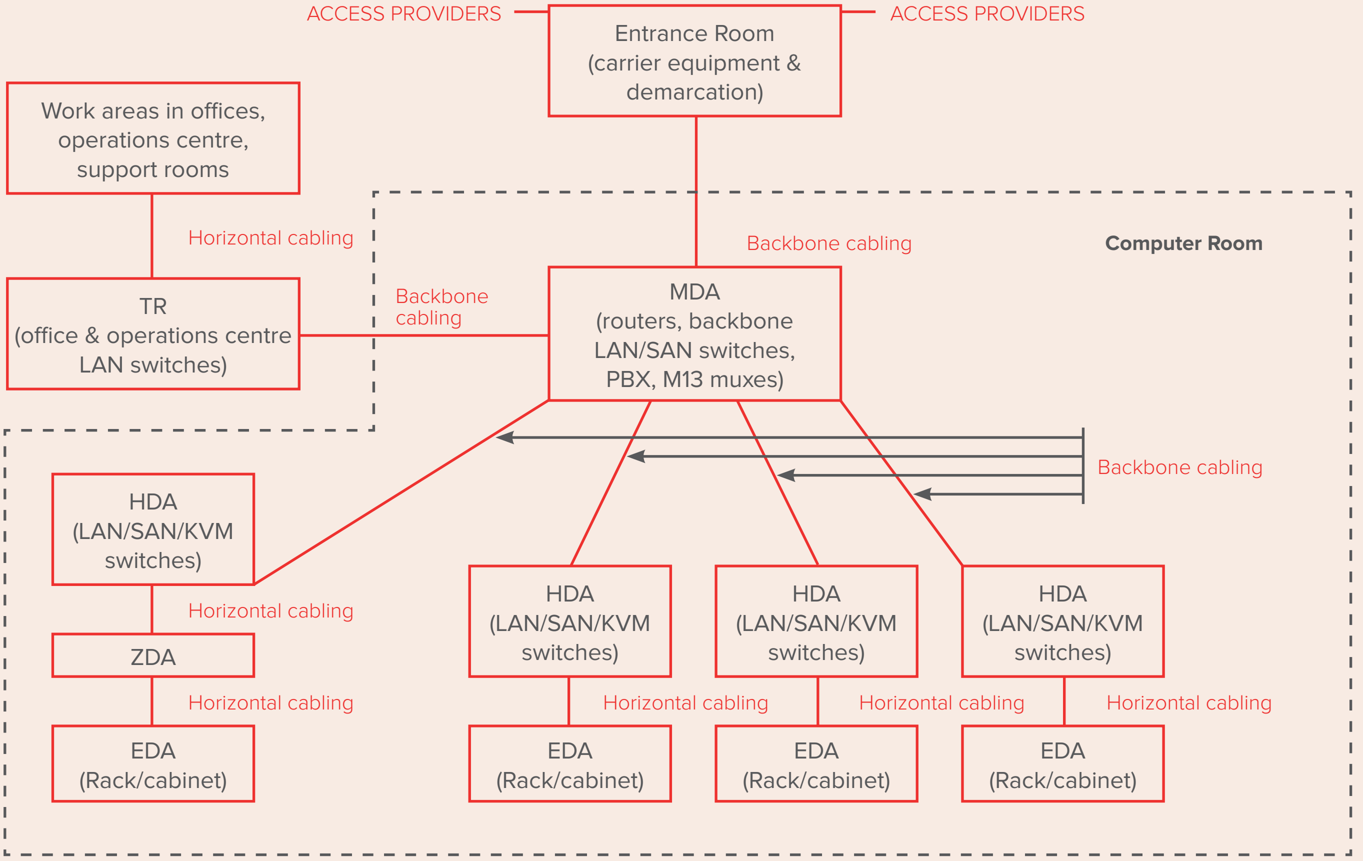
ACCESS PROVIDERS
Entrance Room (carrier equipment & demarcation)
ACCESS PROVIDERS
Work areas in offices, operations centre, support rooms
Horizontal cabling
Backbone cabling
Computer Room
TR (office & operations centre LAN switches)
Backbone cabling
MDA (routers, backbone LAN/SAN switches, PBX, M13 muxes)
Backbone cabling
HDA (LAN/SAN/KVM switches)
Horizontal cabling
ZDA
Horizontal cabling
EDA (Rack/cabinet)
HDA (LAN/SAN/KVM switches)
HDA (LAN/SAN/KVM switches)
HDA (LAN/SAN/KVM switches)
Horizontal cabling
Horizontal cabling
Horizontal cabling
EDA (Rack/cabinet)
EDA (Rack/cabinet)
EDA (Rack/cabinet)

Source: ANSI/TIA-942-A

Balancing Cost vs. Performance

The selection of the network cabling architecture is driven by both the technical and financial factors of the data centre design. In general, top-of-rack (ToR) architectures are better suited for environments requiring low latency and high-performance server connections, whereas end of row (EoR) and middle of row (MoR) architectures look to optimise cost and flexibility.

Figure 13, 14, 15, 16: Cabling Topologies

Key

🟥	MDA	🟩	EDA	🟦	Fibre Cable	🟨	Copper Cable	🟦	Switch	🟪	HDA

Top of Rack (ToR)	Application	Pros	Cons
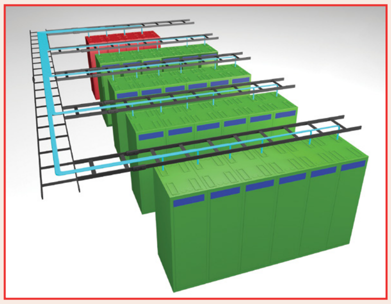	Cabling connects ToR switch to server within the EDA cabinet or rack. Cable types include direct-attached cabling (DAC), OM3 multimode optical fibre jumpers and Category 6A rated patch cords. OM3/OM4 multimode or single-mode cabling is used for backbone connection between EDA and MDA. Cabling architecture is deployed widely in hyperscale and cloud data centres (>100K ft^2/9.3K m^2) where performance is at a premium over cost	• Efficient use of floor space • Excellent scalability • Easy cable management	• More switches to manage • More server-to-server traffic in aggregation layer • Higher network equipment costs (redundancy) • Creation of hotspots due to higher density power footprint
Middle of Row (MoR)	**Application**	**Pros**	**Cons**
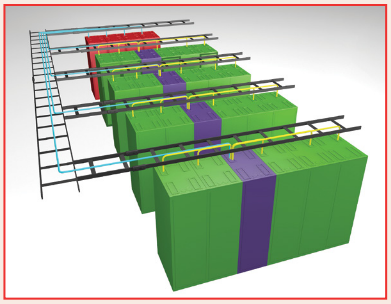 **End of Row (EoR)** 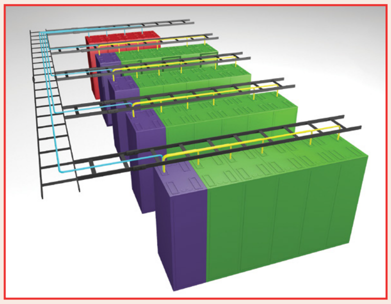	Utilises a traditional MDA-HDA-EDA cabling topology, but the HDA is physically located at the end or middle of EDA cabinet row. OM3/OM4 multimode cabling is used for backbone connections and Category 6A cabling between HDA and EDA. Typically deployed in small to medium sized data centres (<20k ft^2/1.8 m^2)	• Fewer number of cables than direct-connect architectures • Good scalability • Cost effective compared to top of rack (ToR)	• Increased management overhead • Network stability risks due to potential Layer 2 loops that cause broadcast storms
Centralised Cable	**Application**	**Pros**	**Cons**
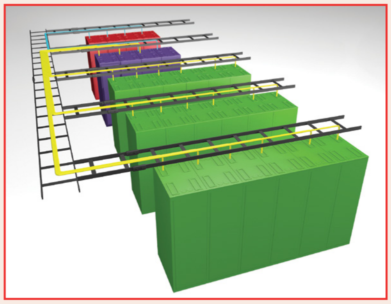	Traditional cabling architecture where network switching is centralised in a row of HDA cabinets and racks. Due to shorter reach requirements, OM3/OM4 multimode cabling is used for backbone connections between MDA and HDA and Category 6A cabling between HDA and EDA. Typically deployed in small to medium sized data centres (<20k ft^2/1.8 m^2)	• Simple to design, implement and maintain • Minimised network bottleneck • Good port utilisation • Easy device management	• Large number of cables • Cable overlaps • Difficulties in cable pathway design • Lack of scalability

BEST PRACTICE 3 OF 4

MAKE THE RIGHT MEDIA SELECTION

3

Copper and fibre considerations

A data centre cabling system may require a balance of both copper and fibre to cost-effectively meet today's needs and support the high-bandwidth applications of the future. Rapidly evolving applications and technologies are drastically increasing the speed and volume of traffic on data centre networks. Making sure that a cabling solution is designed to accommodate the higher transmission rates of these evolving bandwidth-intensive applications is critical.

Growth in 10 Gigabit Copper Solutions

10 Gigabit Ethernet technology continues to evolve and improve, causing wider adoption. Blade servers, networked enterprise switches, video servers and other applications can benefit now from 10 Gigabit speeds in storage, system backup, teleconferencing and surveillance. Technical advances have enabled the higher density, reduced power and improved cost-effectiveness needed to attract all of the major system developers.

For example, the development of 10GBASE-T defined in the IEEE 802.3an standard is expected to accelerate its market acceptance, as 10 Gigabit Ethernet can be delivered over a twisted-pair cabling using the ubiquitous RJ45 physical interface.

Twisted-Pair Technology

There are three strong reasons for the broad acceptance and rapid growth of twisted-pair as the cabling media of choice for switch-to-server interconnect:

- Low initial cost.
- Ability to deliver higher-data-rate LAN services.
- Flexibility to use one medium for all services.

Figure 17-19: Twisted-Pair Cable Types

Cable Type	Main Features	Diagram
Category 6A/Class EA Unshielded Shielded Twisted Pair TIA I ISO Designation: (UTP) I (U/UTP)	· Originally was designed to support 500 MHz operation · This media is used the most for 10 GBASE-T applications · The diameter was originally about 0.354in, but most manufacturers have been able to reduce the diameter to ~0.28in, which is negligibly larger than Cat 6 / Class E · Conductor gauge size is typically 23 AWG · Patch cords are available with solid or stranded conductors	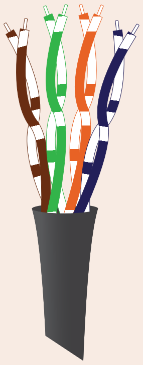
Category 6A/Class EA Foiled Twisted Pair TIA I ISO Designation: (ScTP) I (F/UTP)	· Originally was designed to support 500 MHz operation · The diameter is 0.290in · Conductor gauge size is typically 23 AWG · F/UTP construction has an overall foil shield and unshielded pairs · In order for a shielded cable to work properly and not introduce additional noise to the channel, it must be paired with shielded terminations (jacks & path panels) along with shielded patch cords, and complete telecommunications grounding (earthing) system must be in place · Patch cords are stranded conductor	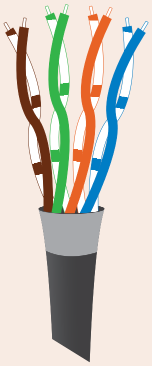
Category 7/Class F Pairs in Metal Foil (PIMF) TIA I ISO Designation: (SSTP) I (S/FTP)	· It is designed to support 600 MHz · Diameter is smaller than unshielded but larger than shielded F/UTP · Conductor gauge size is typically 22 AWG · The diameter is 0.33in · Category 7 relates to components only · Classes ISO relates to channel performance · Classically called "Individual Shielded Pairs with Overall Braid Shield," similar to the original IBM Cabling System Type I Design but in a 100-ohm 4-pair cable · Patch cords are stranded conductor	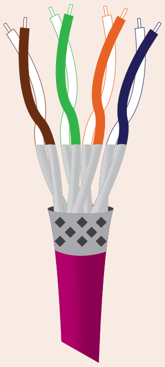

Fibre-optic technology

Fibre can provide numerous advantages over twisted-pair copper in the context of new data centre architectures. These advantages include the following:

› Greater bandwidth and error-free transmission over longer distances.

› Simpler testing.

› Immunity to EMI/RFI.

Because space is always at a premium, high-density fibre-optic systems may be preferred for maximising valuable square footage. Fiber's small size and weight require less space in cable trays, raised floors and equipment racks. As a result, smaller optical networking provides better underfloor cooling and gives precious real estate back to the data centre.

Supporting Higher Data Rates
Fibre has the ability to support higher data rates, taking advantage of existing applications and emerging high-speed network interfaces and protocols.

Multimode fibre optics support:

› 10Mbps/100Mbps/1Gbps/10Gbps/40Gbps/100Gbps Ethernet.

› 1/2 /4/8/16Gbps Fibre Channel.

Laser-optimised 50-micron OM3 and OM4 fibre are generally recommended by storage area network manufacturers because of their higher bandwidth capabilities.

Figure 20: Fibre Optic Cable Types

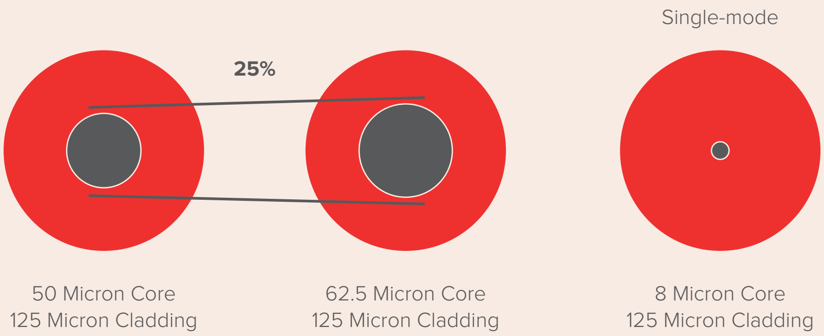

Cable Type	Main Features
OM3 = 50/125 micron multimode fibre	• Laser optimised for 10 GbE up to 300 meters at 850 nm • Aqua jacket on indoor cable
OM4 = 50/125 micron multimode fibre	• Laser optimised for 10 GbE up to 550 meters at 850 nm • Aqua jacket on indoor cable
OS2 = 8.3/125 micro single-mode fibre	• Enhanced performance - Qualified for use across 1310 nm to 1550 nm spectrum - Improved manufacturing processes have eliminated the "Water Peak" attenuation spike • Introduced in 2001 - currently the standard performance in single-mode fibre cable • Yellow jacket on indoor cable

BEST PRACTICE 4 OF 4

ANTICIPATE AND PLAN FOR DENSITY DEMANDS

4

From a physical-hardware perspective, the use of high-density blade server technology facilitates server virtualisation. By containing multiple servers in a single chassis-based enclosure, blade servers maximise CPU processing power per watt of power consumed.

However, this higher-density platform approach changes the data centre design paradigm. As computing resources consolidate into smaller physical footprints, the kilowatt usage per square foot increases, as does the number of port connections. Managing the increased port counts becomes a challenge when trying to make sure there is operational efficiency and flexibility of the structured cabling system.

Minimise space, maximise efficiency

Pre-terminated cable assemblies and modules should be considered when trying to manage dense port counts associated with high-density LAN and SAN equipment and servers within the data centre. Assemblies that use 12-fiber MPO (multi-fibre push on) connectivity provide for both space savings and operational efficiency when compared with traditional multifibre cabling options.

These assemblies provide the following benefits:

- Up to 60 percent smaller optical networking solution frees raised floor and racking space.
- Pre-terminated connectors in 12-fiber increments allow for quick commissioning of network services and a data centre layout that can adapt to increased network computing requirements.
- Preassembled solutions allow for greater cost control, significantly reducing installation variability costs in the field.
- An optical fibre link can be completely assembled and tested prior to leaving the factory (but does not completely remove the need for field testing after installation for link certification).

Migration Path to 40 and 100 Gigabit Ethernet

Two-fibre (duplex) transmission is a common optical infrastructure technology used in today's data centre environments. Two-fibre implementations can use either multimode or single-mode optical fibre media. Multimode infrastructure has the advantage of lower-cost optics that drive a lower overall system cost, compared with the more expensive single-mode optics that offer significantly longer-distance support (in the order of kilometers).

In many instances, the shorter reach associated with multimode optical fibre is sufficient to support the higher-speed 40 Gigabit Ethernet and 100 Gigabit Ethernet systems. However, careful planning is vital because these systems use parallel optical technology with multiple fibre strands to send and receive data.

In order to understand the migration from two-fibre transmission to multifibre parallel optic transmission, the figure below helps outline the optical and connector configurations in each scenario.

It is important to note that the IEEE is pursuing the definition and operation of a "four lane" 100 Gigabit Ethernet interface that will simplify the migration from 40 Gigabit Ethernet to 100 Gigabit Ethernet over multimode fibre in the data centre.

Figure 21: Optical and Fibre Connector Configurations

Data Rate	10GB/s	40Gb/s	100Gb/s
Laser Type	VCSEL	VCSEL Array	VCSEL Array
Fibre Type	OM3	OM3/OM4	OM3/OM4
Connector	LC X 2	12-fiber IMPO	2x 12f MPO or 1x 24f MPO
# of Fibres in Cable	2	12	24
Schematic	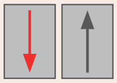	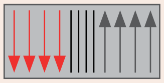	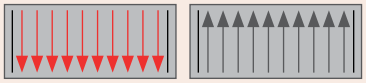

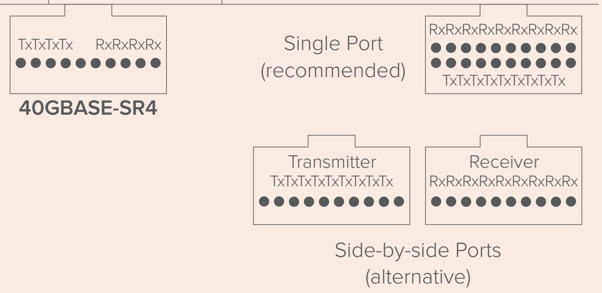

Note: IEEE is working on 100GBASE-SR4 standard using 8 fibres (4 x Tx, 4 x Rx)

Source: Ethernet Alliance

TAKE THE PRACTICAL APPROACH TO SUPPLY CHAIN SERVICES

Another thing to consider for a successful data centre migration is the approach to supply chain services. The foundation for an efficient data centre deployment project is having a distribution network that can be leveraged for product inventory and coordinating deliveries with installation resources. Data centres require materials from multiple manufacturers to arrive at the right place and to be installed on time and within budget.

Fundamental distribution services should include:

- The ability to view and allocate inventory in any warehouse globally.
- A significant investment in a diverse breadth of inventory.
- IT systems that provide customers with real-time information.
- Predictable delivery times and processes to help plan projects.

When planning a complex data centre network migration, it's important to partner with a provider that has a proven distribution network that fulfills these requirements.

Three Supply Chain Principles
Supply chain management has several core principles, including the following:

- **Minimise cost.**
- **Maximise customer service.**
- **Optimise operational effectiveness.**

FINAL RECOMMENDATIONS AND CONCLUSIONS

Anixter understands the challenges facing today's data centre professionals and embraces standards such as ANSI/TIA-942-A, whose goal is to provide a comprehensive road map for success. Anixter supports data centre customers in achieving their goals of availability, reliability and operational efficiency.

Besides the many technical considerations and decisions that contribute to a trouble-free and seamless network, there are a few other points to note. In the area of communications infrastructure, Anixter has some specific recommendations:

› Even though the ANSI/TIA-942-A standard for data centres specifies a copper minimum of Category 6A cabling, cabling should support 10 Gigabit Ethernet in new or expanding facilities.

› At minimum, a Category 6A construction is needed to provide bandwidth in the 500MHz range with proven control of alien crosstalk.

› Fibre cabling should be in all backbones, including the data centre.

› Laser-optimised, 50-micron is the fibre of choice – again to make sure the data centre is ready whenever and wherever 10 Gigabit Ethernet becomes a requirement.

Contact Anixter to learn how network migration works with the other building blocks of Anixter's Infrastructure as a Platform solution. Infrastructure as a Platform focuses on helping you create an agile and scalable data centre by addressing the five key building blocks for data centre interoperability:

› Risk management
› Network migration
› Power optimisation
› Thermal efficiency
› DCIM enablement

For more information on how the four best practices of high-performance structured cabling can improve your network's flexibility, determine the proper topology for your applications, aid in the proper media selection and accommodate uncertain density requirements in your data centre, visit anixter.com/datacenterdcd or contact your local Anixter representative.

GLOBAL TECHNOLOGY BRIEFING

POWER OPTIMISATION BEST PRACTICES

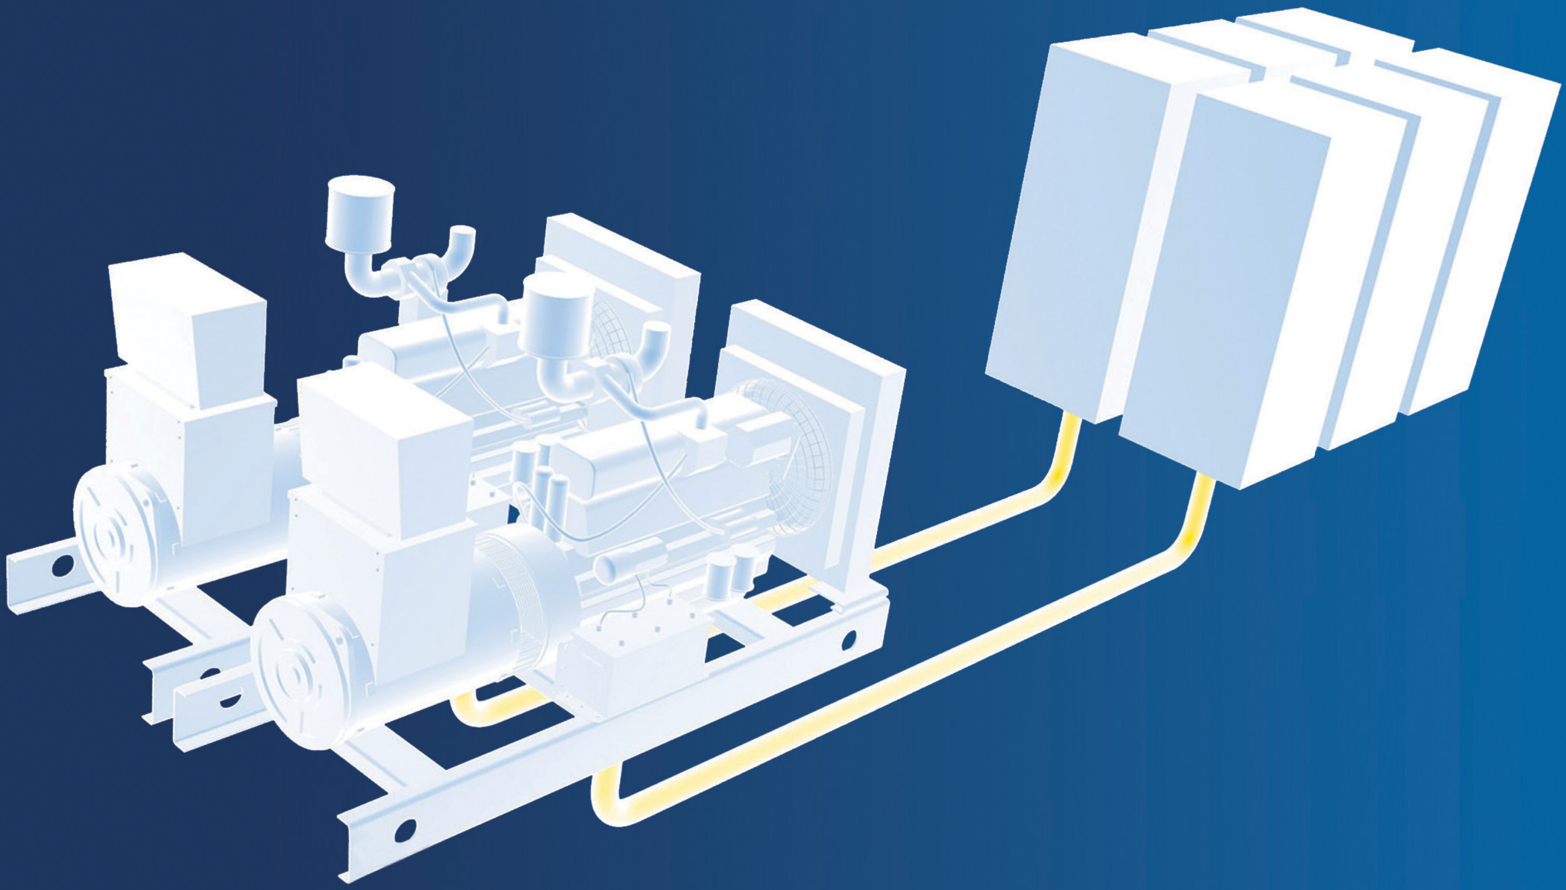

EXECUTIVE SUMMARY

The progression of the data centre mirrors how power is provided, measured, and used for maximum efficiency. The importance of data centre power delivery is paramount. In the past, the data centre's power chain was designed to support where the business was predicted to be in 10 to 15 years.

Now, pressures to operate efficiently to reduce costs, increase IT capacity, and move to a cloud model have forced a shift in thinking on how power is being generated, designed, and managed in the data centre. The data centre itself is becoming a source of revenue rather than just a cost. This means that operating a facility that can scale real-time based on business demands requires smarter capital investment and higher operating margins. This evolution gives more weight to efforts to decrease operating expenses (OPEX) and increase capacity through more efficient power distribution and avoidance of costly outages.

Yet, there exists a gap between efficiency monitoring and consumption with many data centre managers. According to DCD Intelligence, 46 per cent of data centre managers continuously monitor efficiency, yet 68 per cent watch consumption closely as a key energy parameter. The takeaway is simple — if you can't measure it, you can't manage it as management guru Peter Drucker once said.

This report details current data centre approaches, challenges, and standards. It also covers Anixter's concept of the intelligent power chain, which includes the entrance feed, UPS systems, room and cabinet distribution, and the IT equipment.

The report concludes by exploring the future of the data centre including power, industry trends, and energy technologies and what they mean for owners, managers, and operators.

The same basic principles apply to all different approaches to a data centre, including on-premise enterprise facilities, co-location, or cloud. What is different is ownership and responsibility of the equipment. For example, a multi-tenant data centre will likely own the UPS and generator functions, whereas the enterprise customer co-located there is responsible for what is inside the server cabinet.

The majority of this paper will focus on the data hall and rack-level distribution, because it is easier to modify in a retrofit environment, representing the majority of today's data centre activity. In addition, roughly 60 per cent of the total power generated for the facility is consumed by IT equipment, according to DCD Intelligence. Bottom line, you can make a big impact — even with a limited budget — when you focus on optimising power within the data hall and cabinet. So, while all components of a data centre is important, this paper is weighted towards those areas that will provide the maximum benefits from applying an intelligent power chain approach.

According to DCD Intelligence, 46% of data centre managers are continuously monitoring **efficiency**, yet 68% continuously monitor **consumption**.

INTRODUCTION

THE EVOLUTION OF DATA CENTRE REQUIREMENTS

Increasing requirements for continuous accessibility to the Internet — more processing and more storage, for example — have driven the need to increase IT capacity. Other drivers include:

- High-performance computing requirements
- Increased use of Internet communication and streaming entertainment
- Retail, banking, and financial institutions increased online availability and transactions
- Digitisation of healthcare and government records
- Migration toward online business models
- Heightened information and national security requirements
- Video surveillance now moving to IP-based video storage.

To keep up with increasing computing demands, IT is constantly adding servers. However, because space comes at a premium, facility operators face severe constraints in housing additional IT equipment. As a result, servers have been designed to incorporate higher processing power in less space. This type of server consumes three to five times more power than previous generations in the same footprint, so power densities per rack are on the rise.

Higher power demands have necessitated investment in the IT support infrastructure, namely physical cabling, network hardware, and server and storage systems. However, in today's economy the business is pressuring facilities and IT managers to reduce operating expenses, improve space utilisation, enable virtualisation and cloud technologies, and improve sustainability. How can a data centre make strides in improving operating efficiency through new capital investment? The answer lies in the power chain.

According to DCD Intelligence, 27 per cent of data centre operating costs are spent on facility power. By increasing the efficiency of the hardware throughout the power chain, operating expenses will theoretically be reduced. This optimisation can be done by using scalable, modular systems that will help limit data centre capital investment and will allow for growth as IT needs dictate. Measuring at different points along the power's path will provide the data needed to tune the system so it is running as efficiently as possible, and it will also help make sure space is being used effectively.

Figure 1: Data Centre Global Density Per Rack (KW / rack, estimated, IT & non-IT)

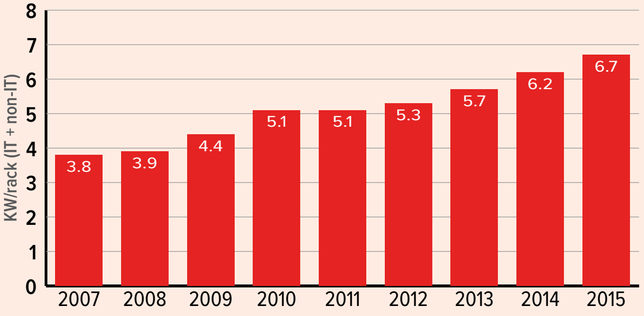

Source: DCD Intelligence

LEGACY POWER DESIGN

The modern data centre consumes more power per square foot than ever, yet many of the same methods and technologies for power distribution to IT equipment has remained the same for decades.

Legacy power designs that are over the 10-year mark are rigid and have limited ability to scale. These systems required high upfront capital investment and were designed to support more than 10 years of IT growth with limited software integration capabilities.

Social and economic drivers have led to the construction of data centres that are designed with a fixed amount of power despite rapidly changing IT needs—namely unpredictable IT requirements that lead to capacity fluctuations. Having a view into the total power chain is critical to ensure capacity is being managed effectively and efficiency goals are met. Having this complete view also prevents unnecessary outages that can lead to downtime, ultimately resulting in a potential cost in the hundreds of thousands of dollars.

In addition, today's applications require more processing power from servers though smaller footprints are increasingly the norm. Couple that with the virtualisation trend, and one can see why power densities per cabinet are rising. So how does your data centre address these issues and support your business's application-focused needs? Whether your data centre is a co-location facility or enterprise, you need to budget and have a contingency plan that addresses all of these challenges.

Designing an intelligent power chain requires careful planning, weighing performance versus costs, and scalable hardware and useful analytics to capture insightful data that will help companies quickly adapt to changing IT requirements.

Figure 2: Annual Data Centre Operating Costs (%)

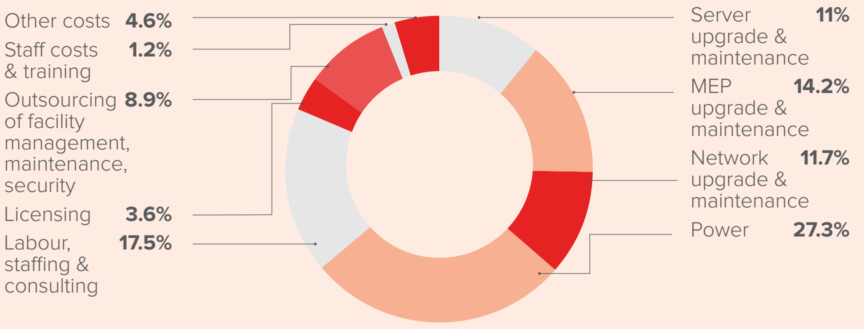

Source: DCD Intelligence

COST OF OUTAGES AND DOWNTIME

Every IT team understands that downtime has the potential to significantly impact the profitability of its business. In extreme cases, an outage can seriously threaten the viability of an enterprise.

Take for example the Singapore Stock Exchange (SGX) outage that closed one of the world's biggest exchanges for almost three hours on November 5, 2014, after its systems failed to cope with a voltage fluctuation caused by a lightning strike. Ultimately, it was discovered that a flaw in the design of the electrical infrastructure failed to prevent the outage, which was attributed to human error on the part of third-party contractors. This was followed by another three-hour outage due to a software error one month later.

The financial fallout and damage to the exchange's reputation was nothing short of devastating. Besides the millions of dollars in revenues lost during the outage, SGX spent 15 million USD to address the flaw and another one million USD on an education fund to help restore investor confidence. The country's regulatory authority also imposed a moratorium on exchange fee increases due to the outage, until an overseeing committee could guarantee the issues were addressed.

In a survey conducted by the Ponemon Institute, the results clearly illustrate the substantial economic impact of data centre downtime to the bottom line. The analysis was based on 67 independent data centres located in the United States, querying five key categories of functional leaders.

Some of the key takeaways from the study were related to the financial impacts around outages and how they are increasing. The cost now ranges from 45 to 95 USD a square foot or a minimum of 74,000 to 1.7 million USD per organisation in the study.

Also, 83 per cent of survey respondents knew the cause of the outages. The most frequently cited root causes were:

- › UPS battery failure – 55 per cent
- › UPS capacity exceeded – 46 per cent
- › UPS equipment failure – 27 per cent
- › PDU/circuit breaker failure – 26 per cent

Clearly, just this small sampling of outage examples demonstrates what downtime can do to a business, with the damage multiplying exponentially according to the size of the business. Outages are inevitable. They are going to happen. However, there are steps that can be taken to reduce the likelihood and duration of outages in the future.

Source: Cost of Data Centre Outage: Ponemon Institute, December 2013

SURVEY OF 67 INDEPENDENT DATA CENTRES		
Facility manager	CIO	DC manager
CTO	IT compliance	Security officer
The cost per square foot of data centre outages now ranges from 45 to 95 USD	Minimum cost of 74,223 USD to a maximum of 1,734,433 USD	

MAJOR CAUSES OF OUTAGE		
UPS system failure	Accidental error	Weather
CRAC failure	Generator failure	IT equipment failure
• UPS battery failure • UPS capacity exceeded • PDU/circuit breaker failure • UPS equipment failure		

TOP CHALLENGES AND CONCERNS

Many data centre power decisions and product choices aim to address and solve the following challenges and concerns.

Preventing accidental outages

Recent research shows that 73 per cent of data centre downtime is caused by human error, according to Rick Schuknecht, Uptime Institute vice president. These types of outages can be caused by poor performance in the areas of training, maintenance, and operational governance. The majority of outages are accidental. You can have the most robust system money can buy, but if someone accidentally unplugs something he isn't supposed to, it will still cause an outage.

Understanding capacity needs

Measurement of capacity throughout the entire power chain is vital to ensure facilities can effectively support future IT needs. Poor capacity planning can lead to outages and delays in deploying new business applications. On the other hand, proactive capacity planning helps to make sure the power chain is optimised and power is there when it is needed.

In some data centres, an opportunity may exist for better collaboration in gauging capacity needs.

Power chain visibility

Having visibility into the entire power chain is important in order to "tune" the system to make sure it is being run as efficiently as possible. According to the Green Grid Association, to have full insight into an infrastructure's energy efficiency, multiple components from the utility entrance through the IT equipment should be monitored. In addition, having this information gives facilities and IT the full picture to better make deployment decisions on new equipment.

Budgeting of the unpredictable

IT hardware requirements are constantly in motion. They are being driven based off the business applications that they need to support. Facilities managers are being pressured to reduce operational expenses by running as lean and efficiently as possible. It becomes a challenge to do that when the infrastructure that needs to be supported is a moving target.

Performance versus investment costs

Weighing design decisions with available budget is always an issue. What is the right balance? How can you design a system that is resilient at a reasonable cost? The reality is more complicated than it used to be. According to DCD Intelligence, the major investment drivers growing year on year have to do with the data centre becoming a strategic rather than just an operational resource. Does it make more sense for the business to move the data centre to the cloud? Have needs shifted due to a recent acquisition or significant market change? Can the current data centre still support the business?

**Cannot gauge
CAPACITY needs**

**Lack of TOTAL POWER
CHAIN visibility**

PERFORMANCE vs COST

**Budgeting of unpredictable
IT REQUIREMENTS**

**PREVENTING
accidental outages**

DATA CENTRE POWER STANDARDS

Standards have enabled our industry to effectively get on the same page, empowering faster advances in technology. In the early days of data centres, many facilities were designed in the absence of established standards, especially as it related to power distribution. Even today, many network administrators face the challenge of making power distribution choices and deciphering how to properly implement them without the benefit of well-researched, documented, and methodical standards.

The following is a selective offering of power distribution standards.

ANSI/TIA-942-A: Telecommunications Infrastructure Standard for Data Centres
The purpose of the ANSI/TIA-942 standard is to provide requirements and guidelines for the design and installation of a data centre or computer room. It is intended for use by designers early in the building development process and covers site space and layout, cabling infrastructure, tiered reliability, and environmental considerations.

As it relates to power, the ANSI/TIA-942-A standard offers an overview on the following categories:

- Utility service entrance and primary distribution
- Standby generation
- Uninterruptible power supply (UPS)
- Computer power distribution
- Data centre grounding infrastructure
- Building management system

ANSI/BICSI 002-2014: Data Centre Design and Implementation Best Practices
The BICSI standard is primarily a data centre design standard with installation requirements and guidelines related to implementing a design.

BICSI offers insight into everything from data centre site selection, electrical systems, mechanical systems, fire protection, data centre management, and building systems and telecommunications cabling. As it relates to power, the BICSI standard offers a comprehensive guide into the following areas:

- Utility service
- Distribution
- Mechanical equipment support
- Uninterruptible power supply (UPS) systems
- Standby and emergency power systems
- Automation and control
- Lighting
- Bonding, grounding, lightning

Figure 3: Data Centre Investment Drivers: Operational Efficiency and Outage Risk Mitigation

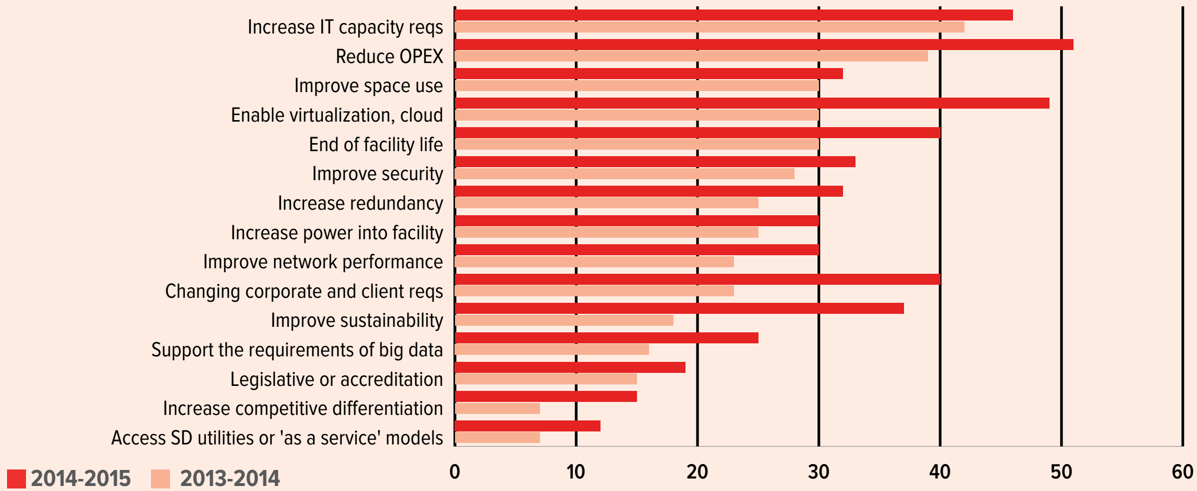

Source: DCD Intelligence

DATA CENTRE POWER STANDARDS

- protection, and surge suppression
- Labelling and signage
- Testing and quality assurance
- Ongoing operations
- Electrical systems matrix

European Committee for Electrotechnical Standardisation (CENELEC)
EN 50600: Information technology: Data centre facilities and infrastructures
CENELEC has developed a series of European standards for data centres. This document specifies recommendations and requirements to support the various parties involved in the design, planning, procurement, integration, installation, operation, and maintenance of facilities and infrastructures within data centres. The EN 50600-2-2 standard focuses specifically on power distribution and specifies requirements and recommendations for the following:

- Power supply and distribution within data centres
- Dimensioning of power distribution systems
- Availability
- Power supply
- Power distribution
- Incorporation of LVDC distribution
- Emergency power off (EPO)
- Energy efficiency enablement and power distribution

The EN 50600-2-2 standard also addresses using intelligence to improve

Figure 4: Possible Power Chain Measurement Points

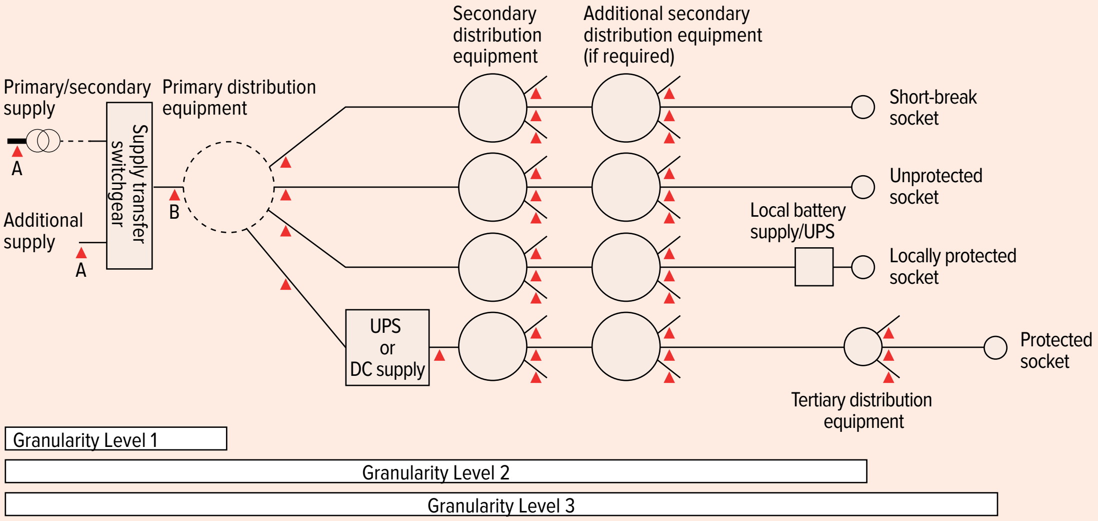

Source: CENELEC EN 50600-2-2 Section 8
Energy Efficiency Enablement and Power Distribution

DATA CENTRE POWER STANDARDS

data centre power efficiencies, outlining three levels of granularity:

LEVEL 1
Provides simple global information for the data centre as a whole

LEVEL 2
Provides detailed information for specific facilities and infrastructures within the data centre

LEVEL 3
Provides granular data for elements within the spaces of the data centre

Uptime Institute data centre site infrastructure tier classification system

Even though it's not necessarily a standard, the Uptime Institute is recognised for the creation and administration of tier classifications and certifications that enable data centres to achieve their mission while mitigating risk. The classification system establishes four distinctive definitions of data centre site infrastructure tier classifications and the performance confirmation tests for determining compliance to the definitions. The tiers can be defined as follows:

- Tier I: basic site infrastructure
- Tier II: redundant site infrastructure capacity components
- Tier III: concurrently maintainable site infrastructure
- Tier IV: fault tolerant site infrastructure

Table 1: Uptime Institute Tier Requirements Summary

	Tier I	Tier II	Tier III	Tier IV
Activity capacity components to support the IT load	N	N+1	N+1	N after any failure
Distribution paths	1	1	1 active and 1 alternate	2 simultaneously active
Concurrently maintainable	No	No	Yes	Yes
Fault tolerance	No	No	No	Yes
Compartmentalisation	No	No	No	Yes
Continuous cooling	Load density dependent	Load density dependent	Load density dependent	Class A

DEFINING THE FIVE BEST PRACTICES OF AN INTELLIGENT POWER CHAIN

Anixter defines the intelligent power chain for data centres as a combination of the right hardware (intelligent hardware design) and the right analytics (software for data collection). Data centres that execute on the concept of the intelligent power chain can gain substantial benefits.

- Improved overall power system efficiencies – for more efficient power distribution to IT load
- Better capacity planning and management – design closer to capacity with option to scale as needs grow, better understand your capacity ahead of time, and free up stranded capacity to extend the life of the facility
- Reduce risk of outages – proactive monitoring (data/analytics) that helps you identify trends and understand potential problems before they create an outage; create a clear business process to mitigate and resolve outages quickly

The following sections will examine five key areas where the application of the intelligent power chain concept will make the most difference in the dependability and efficiency of data centre power distribution.

Anixter defines the intelligent power chain for data centres as the right hardware (intelligent hardware design) with the right analytics (software/data collection).

THE INTELLIGENT POWER CHAIN

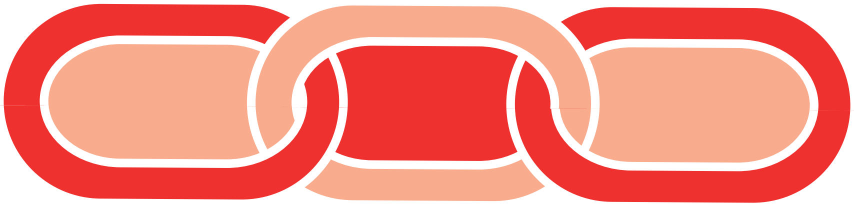

IMPROVED OVERALL
efficiency

BETTER CAPACITY
management

REDUCED RISK
of outages

BEST PRACTICE 1 OF 5

ENTRANCE FEED

1

The data centre entrance feed is the gateway that connects a facility with power utilities. Proper selection and sizing of the medium-voltage feeder cable from the utility to the service entrance can save a business a significant amount of investment capital. Paired with monitoring and measurement of the critical power applications, it will help to prevent downtime and improve efficiency.

Sizing

The ampacity of the cable should equal or exceed the maximum current the cable will be expected to carry during its service life. Conductors that are undersized can overheat, cause damage to the insulation or jacket of the cable, and potentially cause harm to equipment or people. It's better to install a cable with a higher ampacity rating versus one that is too small, allowing for future growth.

Jacket type

Data centres aren't industrial settings, so PVC tends to be a popular and suitable choice for the jacket. Most facilities in the U.S. use EPR insulation because of its excellent mechanical properties, high resistance, and durability.

Termination

As far as termination, medium-voltage cables require a special on-site termination. An experienced installer will ensure this is done correctly. Testing should be conducted to measure current and any discharge coming off the cabling after termination. This will help ensure that the cabling, once installed, is terminated correctly and performing properly.

Testing

Besides cabling installation, testing is another vital area for the entrance feed cabling system. Field tests can be broadly grouped into three categories:

- Acceptance
- Maintenance
- Fault location testing

Conducted on wire or cable after an installation but before placing it into service, an acceptance test detects installation or shipping damage that might affect cable performance. After the cable has been placed in service, maintenance tests detect in-service deterioration. Fault location tests pinpoint the exact failure site in a cable. Identifying the failure point of a cable permits the cable to be repaired or replaced as necessary.

Deployment

Medium-voltage cables for data centres are a heavier type of cabling installation. They contain larger conductors that may be difficult to work with and move. Installation can put a significant amount of stress on cabling, so it is important to:

- Allow for adequate clearance between conduit and cable
- Use suitable lubrication — there are a number of commercially available wire pulling compounds (many of which are UL Listed)
- Avoid sharp bending of the cable at the first pulley in overhead installations
- Use sufficient rollers to stop cable from dragging on the tray.

Monitoring the entrance feed

Consider monitoring at the point of utility service entry into the building. A power meter can be installed directly on the low- or medium-voltage switchgear that brings the incoming utility power into the building. The meter can be mounted on the wall or inside the switchgear cabinet, but it is typically mounted into a cutout in the front panel of the switchgear cabinet for easy viewing without opening the equipment. The meter has a wiring harness that connects current transformers (CTs) to the incoming three-phase power to deliver measurements of voltage, current, and frequency. Based on that data, the meter can calculate watts or kVA. Consider meters that use Modbus or Ethernet so the data can easily be pulled into a data centre infrastructure management (DCIM) or building management system (BMS) software without the use of protocol converters.

ENTRANCE FEED

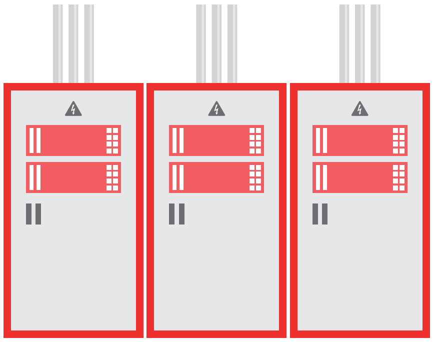

BEST PRACTICE 2 OF 5

UPS SYSTEMS

2

An uninterruptible power supply is one of the cornerstones of the data centre; it offers essential protection against disruptions that would otherwise result in downtime, lowered productivity, and even server hardware issues. Power costs are high and all power required by the critical IT load comes from the UPS, which means subtle gains in efficiency could mean savings of hundreds of thousands of dollars per year. According to a DCD Intelligence survey, the two highest drivers for investing in a new UPS were to reduce operating expenses and the need to increase capacity.

Legacy UPS systems, or systems 10 to 15 years old, are on average 5 to 15 per cent less efficient than their modern-day counterparts. In addition, modern UPS systems are much more efficient at smaller loads, which is typical with data centre designs (less than 40 per cent utilisation). Selecting the best UPS for your enterprise takes a thoughtful and methodical process that marries the ideal hardware configuration with analytics. Your UPS system should meet the unique needs of your computing equipment, as well as be aligned with your energy management and power distribution strategy.

The types of data centre UPS systems can be broken down into three categories:

- Offline
- Line interactive
- Online double conversion

Offline and line interactive UPS systems are rarely used by data centre operators, largely because of long switch-over times and load problems. Online double-conversion UPS systems constitute the vast majority of data centre UPS systems. That's because this technology creates new and clean wave versus an offline or line-interactive UPS that just filters the power, thus an online double-conversion UPS is better suited for mission-critical applications.

UPS design considerations

When it comes to the various technologies and design factors that make up a data centre UPS system, there are several innovative and highly useful variables that you should consider. The following is an overview of these design considerations with some guidance on how they may factor into your UPS purchasing decision.

> Legacy UPS systems are on average 5 to 15 per cent less efficient than their modern-day counterparts.

Modular UPS systems

One UPS design consideration that can help data centre operators decrease energy waste by scaling as the IT needs dictate is the introduction of modularity. The closer a UPS operates to its full load capacity, the more efficient it will be, so modularity can allow operators to adjust the size of the UPS system based on the needs of the IT equipment. Modular UPS designs can either scale through hardware modules or use software-based activation keys.

Modular UPS systems have several key benefits:

- Avoid term overspending in capital and defer capital expenditures (CAPEX) until needed
- Buy only the power components needed, reducing installation and maintenance costs

UPS SYSTEMS

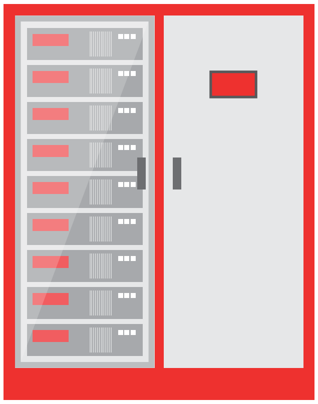

- Provide internal N+1 capability, reducing redundancy costs and floor real estate

The last point is an important one. If you choose a modular UPS system, make sure the architecture allows for the scaling of UPS capacity to closely match your actual load. Doing so will have the most substantial impact on efficiency.

Transformerless UPS systems

In the past, many legacy UPS solutions had a permanent transformer installed onboard that increased its physical size and weight and reduced its efficiency. There are reasons why a transformer might be required. For example, there might be a need to include stepping down voltage from 480 V to 208 V because the mains voltage is not the same as the voltage used by the IT equipment, which is common in North America.

Nowadays, many UPS solutions are omitting the transformer, thereby increasing UPS efficiency and allowing for greater flexibility for placement. Today's transformerless UPS systems are not only significantly smaller and lighter than transformer-based systems, but also more efficient, more reliable, and better equipped to limit fault current. In addition, they enable companies to capitalise on sophisticated features such as the energy saver system and variable module management system, which add reliability via reduction of mechanical complexity while lowering power costs.

It's because of these advantages that transformerless UPS designs are becoming more wildly adopted. In fact, today's transformerless UPS designs outnumber older technologies by a factor of two to one for new installations in North American data centres. According to IHS Technology, high single-digit growth rates are expected over the next five years in transformerless UPS systems, outpacing the growth of traditional transformer designs.

Figure 5: Flywheel UPS System

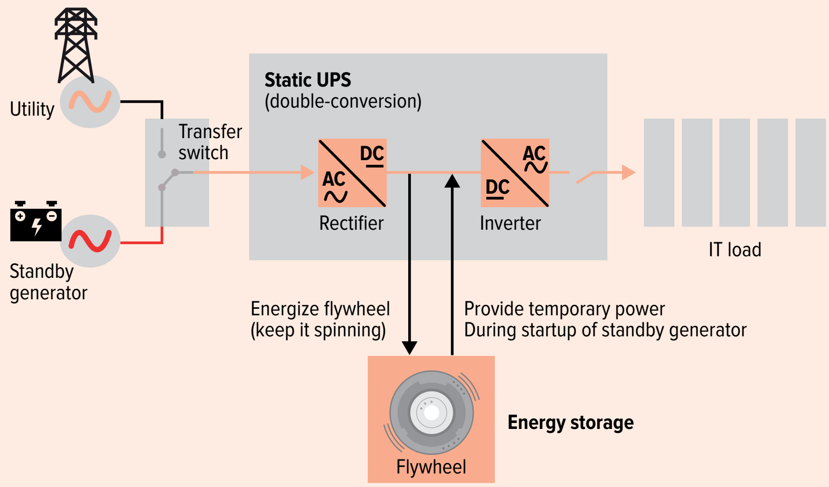

Source: Green Data Centre Design

Flywheel UPS systems

Flywheel systems store energy kinetically, using the inertia of a spinning mass to store and regenerate power. They are mainly used to provide load levelling for large battery systems, such as an uninterruptible power supply for data centres, because they save a considerable amount of space compared to battery systems.

Flywheel systems in production as of 2001 have storage capacities comparable to batteries and faster discharge rates. Newer flywheel systems completely levitate the spinning mass using maintenance-free magnetic bearings. There are several key advantages to flywheel UPS system use in a data centre:

- Less mass that creates a lighter, more compact footprint, saving premium data centre space and making weight less of a consideration in floor design
- Excellent energy efficiency with minimal heat generation
- Reduced noise, with the flywheel, coolant pump and fan typically only operating when needed

- Relatively simple installation
- Low maintenance, with routine replacement of bearings every five to 10 years, with some newer models eliminating mechanical bearing maintenance and failures

Diesel rotary UPS systems
Data centres rely on an uninterruptible and continuous power supply, and generally a diesel-generator back-up system is a requirement. Diesel rotary uninterruptible power supply devices (DRUPS) combine the functionality of a battery-powered or flywheel-powered UPS and a diesel generator.

Figure 6: Diesel Rotary UPS System

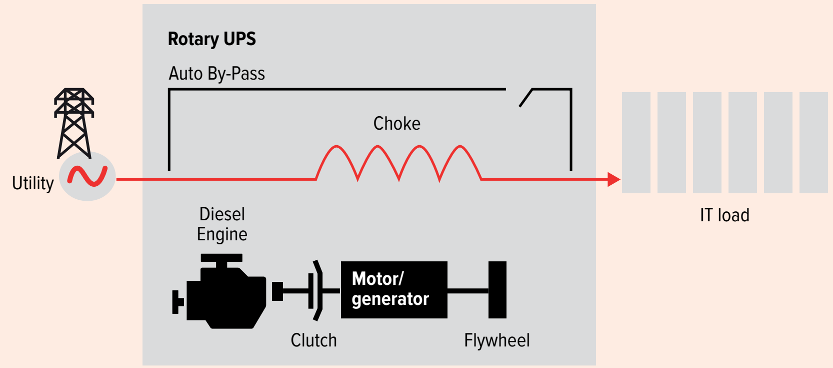

Source: Green Data Centre Design and Management

When mains electricity supply fails, stored energy in the flywheel is released to drive the electrical generator, which continues to supply power without interruption. At the same time (or with some delay, for example two to 11 seconds, to prevent the diesel engine from starting at every incident), the diesel engine takes over from the flywheel to drive the electrical generator. The electromagnetic flywheel can continue to support the diesel generator in order to keep a stable output frequency. Typically a DRUPS will have enough fuel to power the load for days or even weeks in the event of failure of the mains electricity supply. Some advantages of a DRUPS compared to a battery-powered UPS system include:

- Higher overall system energy efficiency
- Smaller footprint
- Use of fewer components
- Longer technical lifetime (no use of power electronics)
- No chemical waste (from batteries).

ENERGY STAR certified UPS

- The ENERGY STAR program specification for UPS systems establishes minimum average efficiencies based on different input dependency characteristics – voltage and frequency dependent (VFD), voltage independent (VI), and voltage and frequency independent (VFI)
- Rated output power goes from less than 1,500 kVA to greater than 10,000 kVA.
- Load profiles range from 25, 50, 75, and 100 per cent load

New, efficient UPS products generally range from 92 per cent to 95 per cent efficient. An ENERGY STAR qualified UPS can cut energy losses by 30 to 55 per cent. A 1,000 kVA UPS used in a large data centre could save 18,000 USD annually.

Eco-mode operation
The “eco-mode” with today’s data centre UPS system – also known as “high-efficiency mode,” “bypass mode” or “multi-mode” is a method of operating the UPS through the bypass line at a reduced power filtration in order to obtain improved electrical efficiency and save energy. Most surveys indicate that eco-mode (or similar technology) can provide a reduction of approximately two per cent in data centre energy consumption.

Eco-mode offers efficiency due to the bypass path, which is typically between 98 and 99 per cent, as compared to the base UPS efficiency of 94 to 97 per cent. The downside is that the IT load is exposed to raw utility mains power without normal conditioning from a double-conversion, online UPS. The UPS must continuously monitor the mains power and quickly switch to the UPS inverter when a problem is detected, before the problem can affect the critical load. More advanced UPS systems with eco-mode have integrated technology that mitigates these risks substantially.

Nevertheless, many data centre operators do not use eco-mode, mainly due to risks associated with electrical protection and reliability as described above. These risks vary depending on the design of a data centre's electrical architecture and the exact UPS eco-mode design approach and functionality. When the risks and rewards are weighed, some operators opt to forego the energy savings to avoid potential issues, while others will take steps to mitigate the risks and choose to use eco-mode. One crucial step that will help to avoid any negative impact of eco-mode is to ensure that the utility power is in an acceptable voltage tolerance of the UPS system's voltage settings.

Three benefits of running a UPS in eco-mode include:

- Less generated heat, which puts less strain on the cooling system
- PUE improvements
- Energy savings increases (exact savings dependent on data centre size).

Bottom line, eco-mode does involve some risks, but advances in eco-mode technology have significantly reduced these risks and have done so at only a small cost in efficiency. When evaluating a UPS system and the use of eco-mode, pros and cons should be weighted.

UPS and battery health monitoring

There are two areas of visibility that are vital to the success of any data centre: UPS monitoring and battery health monitoring.

To effectively manage and monitor UPS systems, consider these action items:

- Make sure the UPS system has Ethernet connectivity to simplify integration into existing monitoring systems.
- Integrate data into a building management system or a DCIM solution so when there is an issue both facilities and IT are notified.
- Look for UPS systems with an onboard LCD display that shows status information.

Battery health monitoring is the other important front. According to the Ponemon Study on data centre outages, 55 per cent of respondents who reported an outage attributed the issue to UPS battery failure. Additionally, the same report revealed that only 46 per cent feel that the UPS batteries are regularly tested and monitored.

There are four factors that can impact battery life:

- Ambient temperature
- Battery chemistry
- Cycling
- Maintenance

It's important to ensure guidelines for proper storage, temperature, usage, and maintenance, which contribute to battery life, are followed.

How long should batteries last? The Institute of Electrical and Electronics Engineers (IEEE) defines "end of useful life" for a UPS battery as the point when it can no longer supply 80 per cent of its rated capacity in ampere-hours.

Battery monitoring software can be a feature of a modern UPS or standalone. Look for solutions that provide some form of visual health indication so you know if the battery is functional. You also want to be able to access data remotely with real-time alerting.

SOME OF THE BENEFITS OF RUNNING IN ECO-MODE INCLUDE...

More efficient UPS generates less heat, putting less strain on the cooling system

PUE improvements

Energy savings increases (exact savings dependent on data centre size)

BEST PRACTICE 3 OF 5

ROOM DISTRIBUTION

3

Power distribution technology has advanced substantially in efficiency, density management, monitoring capabilities, and flexibility. The options available for data centres are greater than ever, allowing facilities to more accurately align the needs of IT. Changes in how power is deployed in data centres has been largely driven by the desire to be more efficient and better manage capacity.

Applying the concept of the intelligent power chain – an efficient hardware configuration planned for variable capacity demands and tuned for application optimisation – is paramount in this area because it touches upon so many different aspects of data centre functionality.

Modular PDUs
Modern data centres must meet rapidly increasing demands, so alternative power distribution approaches, such as modularity, are starting to become more prevalent. Modularity can help provide benefits such as greater reliability, flexibility, and easier manageability.

Modular options provide the following advantages:

- Integrated branch circuit monitoring
- Generally require less floor space than traditional PDUs
- Transformerless options increase efficiencies

Preterminated power whips
Overhead or underfloor cabling from the floor PDU can be run to feed the rack PDUs. The cables can be preterminated to simplify on-site deployment. A best practice in a raised-floor environment is to distribute power cabling overhead to allow for more efficient airflow delivery from perimeter cooling systems.

Busway distribution
The use of overhead busway systems in data centres is typically comprised of busway sections, or straight lengths containing busbars, as well as a slot for continuous access. Tap off boxes, or plug-in units containing circuit protection and wiring devices, are also typically integrated into a busway distribution system.

Busway systems can provide an overhead power distribution solution that integrates valuable features. These can include branch circuit monitoring via integrated or retrofittable metering units, as well as flexible design configurations that allow operators and data centre architects to simplify layouts, add components, and make any other needed changes quickly and easily. This flexibility also accommodates varying voltages and rack densities through the use of customisable tap off boxes.

In addition, a data centre horizontal bus distribution system:

- Provides visual control on rack loading, whips, and rack PDUs
- Speeds up and minimises coordination and costs for upgrades and changes
- Is simple and safe to install installation

The growth of busway systems is a direct result of these practical advantages. According to IHS Technology, overhead

ROOM DISTRIBUTION

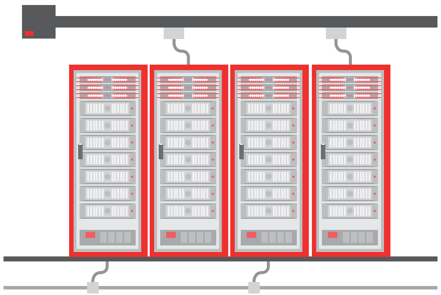

busbar systems are expected to grow in the high single digits within the next five years.

Higher voltage delivery

In North America, typical data centre voltages are 120 V and 208 V, which inherently have some inefficiencies when compared to the 230 V used in Europe and Latin America (240 V in Australia). IT equipment manufacturers typically design their power supplies to accept up to 240 V, due to the differences in voltages around the world. A general rule of thumb is the higher the current (amps), the higher the electrical losses and cost of energy.

In addition, higher overall efficiencies throughout the power chain result in:

- Less current and heat for lower cooling costs
- Reduced copper costs, due to the thinner wire required for higher voltage and less current
- Less transformers because you can consolidate floor PDU transformers into a single transformer sized to the UPS capacity
- Increases in floor space because you can move the transformer outside the room, freeing up space for IT equipment.

But higher voltage isn't without its downside. Like any attempt to maximise efficiency, complications may reduce the potential gains. One such complication is the added costs associated with running a full neutral connector within the system to all the distribution points. Another risk is that equipment needs to handle higher levels of available fault current, so more expensive branch breakers with higher interrupting current ratings are needed.

It is recommended due to the nuances associated with high-voltage power designs that data centre operators should:

- Apply higher voltage to new (greenfield) data centre designs, or self-contained "pods" within an existing data centre
- Consider rack PDUs that provide distribution breaker options capable of delivering high-interrupting capacities
- Use rack PDUs that provide current overload protection for connected loads

According to a recent report by IHS Technology, 400 Vac rack PDUs have recently grown in adoption in North America, which means higher voltage delivery from the floor PDU versus traditional 208 V power distribution. This trend has largely been driven by new, large data centre build-outs.

According to IHS Technology, use of 400 Vac rack PDUs is growing in North America largely due to new data centre build-outs.

Branch circuit monitoring

Power monitoring provides the data needed to:

- Prevent downtime
- Better manage capacity
- Improve efficiency

It's a tool that provides insight to use power more wisely and efficiently through greater transparency into energy usage. Today's operators understand the value of branch circuit power monitoring. However, to maximise its effectiveness, it's also important to implement best practices.

One best practice is using hardware that provides utility-grade accuracy, or within one per cent of the actual amount of power consumed. Although most data centre power meters claim to offer accuracy that is within five per cent of the actual power utilisation, a utility-grade level of accuracy enables co-located and other data centres to fairly rebill clients for the cost of energy.

It's also vital to incorporate both flexibility and adaptability within your data centre.

With the variety of power distribution products and suppliers, your power monitoring system has to be able to work with all of them. So it's vital to have a platform that ensures compatibility with your equipment as well as various amperage sizes and circuit configurations.

Another consideration is the use of standard, not proprietary, communication protocols for the hardware used to collect the power data. Whether SNMP, Modbus TCP, or BacnetIP, you need to make sure that you can integrate the meters with a DCIM or BMS system. Metering platforms should support all power distribution products, communicate easily with the software, and interact seamlessly with the other data centre components as well.

Finally, look for a power monitoring solution with robust and rich functionality. Most monitoring solutions use a complex and costly network of protocol conversions, middleware, and data interpretations to give operators and staff the most complete picture of power usage. Look for useful and practical features such as onboard Ethernet, onboard data logging, onboard alarming, and a Web interface that can reduce the failure points and cost associated with a monitoring deployment.

Metered power cables

Another monitoring option available today is the integration of metered power cables that can provide a host of valuable real-time information and analytics. These products feature the same power monitoring components found in many rack PDUs packaged in a unique power cord format. Intelligent power cables can provide a host of monitoring options, such as tracking temperature, humidity, and pressure.

Look for useful power monitoring features such as onboard Ethernet, onboard data logging, onboard alarming, and a Web interface.

These cables are typically self-powered by the line voltage with some products having the capability to wirelessly transmit detailed power information and self-configure with other nearby compatible cables. A key benefit to using this type of cable is with installations where the operator might not have the ability to install intelligent PDUs in the cabinet. For instance, some SAN solutions have everything integrated into the SAN cabinet. These cables are installed just upstream of the cabinet, allowing for data capture without voiding the SAN provider's warranty.

In addition, metered power cables are ideal for a retrofit (brownfield) environment where there are nonintelligent PDUs installed and a lack of budget dollar to replace them.

Metered power cables can be plugged into the PDU (replacing a standard power whip) to monitor, which makes for a highly cost-effective solution.

BEST PRACTICE 4 OF 5

CABINET DISTRIBUTION

4

Intelligent PDUs are growing at an exceptional rate in the data centre space. In fact, global PDU revenues are forecast to grow 5.6 per cent in 2015, according to a recently published IHS report. In addition, intelligent PDUs are growing twice as fast as basic (nonintelligent) PDUs highlighting the continued shift toward adopting intelligence in the IT cabinet.

There are several reasons why intelligent PDU adoption is growing globally. Intelligent PDU technology can help a data centre operator:

- Effectively monitor usage
- Report efficiency metrics
- Decrease power use in the data centre
- Enable capacity planning.

According to the Data Centre Alliance (DCA) Project, in a typical data centre the IT equipment consumes roughly 59 per cent of the total power to the facility. If you can use the data gained from an intelligent PDU to free up stranded capacity and measure how you can be more efficient in your cabinets, you can potentially make a significant impact on your operational expenses and even defer additional capital expenses.

Types of PDUs

Data centre operators have a choice of several types of PDUs that range significantly in price and function. The tier level and mission-critical nature of your data centre will often dictate the PDU you choose.

Basic – The most economical choice; provides simple power feeds to equipment with no additional features

Figure 7: **EU Data Centre Power Consumption (KW)**

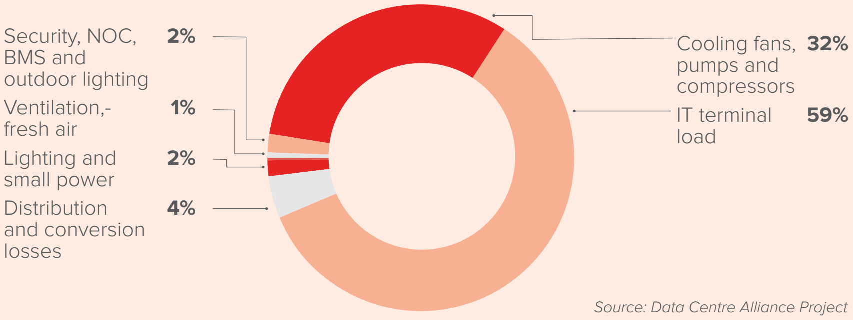

Source: Data Centre Alliance Project

Monitored – Provides a status on input current via an onboard display, a handy feature that allows the user to locally verify the load on the circuit or phase

Metered – Provides the ability to broadcast information via an onboard communications port allowing data to be collected remotely; allows PDUs to be monitored at the input or outlet level

Switched – Connects to the network and offers remote control of individual outlets. Typically monitoring is also available at the input or outlet level

Upgradeable – Gives data centre managers the flexibility to install the intelligence they require today with the option to upgrade technology as needs evolve

From a value-oriented perspective, basic and monitored can be considered entry-level choices, due to their low-cost, nonintelligent design. Intelligent PDUs (metered and switched) are forecasted

CABINET DISTRIBUTION

to grow twice the rate of nonintelligent PDUs, according to IHS Technology. Adoption rates currently are higher in Europe and the Americas, with steady growth in the Middle East as well.

PDU selection considerations
PDUs should be chosen based upon the product's ability to serve a facility's need for optimum power distribution. Here are some key factors to consider.

Application needs
These needs will drive the type of PDU you choose in many respects. If the supported application is mission-critical, then an intelligent PDU with more functionality is recommended. You may also be faced with a scenario whereby power is being monitored upstream at the breaker, and IT is gathering power usage data directly from the server, so the need is simply power distribution at the cabinet. In those situations, a basic or monitored PDU is all that's needed.

For mission-critical applications, you may also want to consider PDUs that have hot-swappable intelligence modules. That way, should a PDU intelligence module fail it can be replaced in the field without having to power down the entire strip, keeping the equipment online and mitigating the risk of downtime.

Business needs
Some IT departments will "charge back" the costs of their services that support the business. However, charging back the cost of power can be challenging. A metered by outlet PDU can accurately measure (within one per cent) power down to the individual outlet. That information can be pulled through a standalone power monitoring software or fed into a central DCIM solution. PDUs that claim that they have "utility" or "billing" grade accuracy should be certified to ANSI C12. 1-2008 or IEC 62052-11 or 62053-21 standards.

Something often overlooked when purchasing intelligent PDUs is the need to support the IP connectivity. Ethernet ports are very valuable on a network switch and using up a port per PDU can prove quite costly. In order to reduce those costs, consider PDUs that support one of the following technologies:

› Daisy chaining, which allows a limited number of PDUs to connect together under one IP address

› Wireless connectivity, which eliminates the need for copper cabling entirely and transmits information back via a proprietary gateway mounted in the IT cabinet

› Ethernet connection via a proprietary wired gateway mounted in the IT cabinet. The amount of PDUs that can be supported will be limited to the gateway itself. Generally, the gateway will also have additional ports designed to support temperature/ humidity sensors, cabinet door locks, and cameras

Equipment location
If there is equipment located in a remote location, with no IT staff locally, a switched PDU could help minimise potential downtime. Switched PDUs have individual outlet control functionality, so if a server isn't responding to the outlet it is connected to, it could be turned on and off remotely without IT staff intervention or requiring remote assistance. However, proper asset management practices need to be followed to ensure the accuracy of the server to outlet associations.

PDU COST

Typical cost for a 48-port Gigabit switch

$2500

or

$52

per port

Typical 100-cabinet installation contains

200

intelligent PDUs Each requiring one switch port

For a total cost of roughly

$10,000

for Intelligent PDU Connectivity

Cabinet density

The choice of PDU depends on the purpose of the individual cabinet (server, storage, and network) and the power density required. In cabinets that are relatively full of equipment, inrush current protection could be needed. In the event of power loss to a cabinet, once the power comes back online, the resulting high-current surge during start up (inrush current) can be several times greater than normal operating current. This could result in tripped fuses and circuit breakers. Switched PDUs can alleviate this concern by allowing the user to power outlets sequentially or groups of outlets that minimise the effects of inrush current.

Low-profile circuit breaker PDU designs are a good choice for higher-density cabinets. These designs minimise the space required to mount the PDU within the cabinet and allows operators more clearance when working within the cabinet. In addition, when PDUs are side-mounted within a data centre cabinet, not having low-profile breakers can sometimes lead to equipment accessibility issues when performing maintenance, particularly the outlets near the bottom of the PDU. That means having to remove the PDU to access certain IT equipment, potentially causing downtime.

DCIM integration

More and more data centres are deploying DCIM solutions in their environment. In order for DCIM software to be effective, data from intelligent hardware are needed. PDUs provide cabinet-level power data that are fed to DCIM software in order for managers to report on that information and make decisions.

PDUs can provide more than just power data. Various sensors, such as airflow temperature and pressure, humidity, water leak detection, and door contact all can be plugged into various dedicated ports on the PDUs themselves, feeding valuable information back to the DCIM software. Onboard USB ports further provide functionality such as the ability to use Wi-Fi dongles, and even video cameras.

Thermal management strategy

In 2008, the American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) expanded the recommended temperature range at the inlet of the server from 68° F (20° C)–77° F (25° C) to 64.4° F (18° C)–80.6° F (27° C).

As a result the IT exhaust temperatures also rise. So, depending on the thermal management strategy one adopts, along with the density of the cabinet, IT exhaust temperatures can reach upward of 120° F (49° C). The PDU selected needs to be rated above those temperatures to prevent any failures. A majority of outages at the cabinet are caused by human error. There are five simple and cost-effective choices to help reduce the potential for costly cabinet outages. Many of them use visual cues, such as colours, to help your staff easily identify power feeds, which partly ensures work is done on the right equipment.

The overwhelming majority of outages can be attributed to human error.

Table 2: Cabinet PDU Functionality

Category	Basic	Monitored	Metered	Outlet metered	Outlet switched	Outlet metered and switched
Branch/phase circuit breakers	●	●	●	●	●	●
Local display		●	●	●	●	●
Branch/phase monitoring		●	●	●	●	●
Temperature/humidity monitoring			●	●	●	●
Network IP address for remote monitoring			●	●	●	●
IP daisy chaining			●	●	●	●
DCIM integration			●	●	●	●
Outlet sequencing				●	●	●
IT device chargeback				●		●
Sombie device identification				●		●

Colour-coding on phase outlets, power cords, and intelligent PDUs can help to mitigate that risk.

Cabinet-level monitoring

Data centre operators and owners are under constant pressure to increase efficiency and reduce energy costs, but without proper monitoring tools in place to understand where your power is consumed or where there are hot spots, these goals are difficult to achieve. It's also vital to understand the impact of any power conservation changes within your data centre's eco system in order to avoid unintended performance degradation or downtime.

As stated previously, more than half of the power in the data centre is used at the IT load or within the equipment cabinet or rack (typically anywhere from 40 to 60 per cent of power usage within a data centre). The IT load also is the largest source of heat, making these two areas the obvious targets for increasing efficiency and lowering costs. Here we'll look at two different ways to monitor the IT load at the infeed (cabinet) and individual outlet (device).

Infeed (cabinet)

The power cable whips coming out of the remote power panel (RPP) are typically called branch circuits and are the power infeeds into the rack PDUs. This is the point in the power chain where you can measure the amount of power used within each cabinet and available power for new devices, helping you to better understand how to manage capacity. Software tools allow data to be plotted to view trends over time. Sometimes this information is already being monitored by facilities team at the RPP directly, but often this data resides with a building management system (BMS) that is not usually accessed by your IT group. Therefore, having this data at the rack level is useful so it can be accessed and compared by both facilities and IT.

Outlet (device level)

The use of power monitoring at the outlet is growing rapidly because it gives users a better understanding of information down to the individual device. Outlet-level monitoring of more than one outlet can be directly related to a particular device or server. Device power information has value because it helps you look at the power consumption of a group of similar devices to help determine those that are actually useful versus sitting idle and wasting power. Multiple devices or servers can often be tied to the power usage of a particular application or group within your organisation. This information is

1 Procure PDUs with locking receptacles when possible. Locking power cords can be used as a substitute if the PDUs that are already deployed do not have outlet locking functionality.

2 Color-coded alternating phase outlets help to take the guess work out of balancing the phases at the cabinet by distributing the phases on an alternating receptacle basis vs. grouping the phases together.

3 Color-coded power cords, with proper labelling, provide a simple visual identifier for different power feeds at the cabinet.

4 Color-coded PDUs provide visual cues to identify different power feeds and different voltages in a mixed-voltage environment.

5 Rack automatic transfer switches (ATS) allow devices with a single power supply to be connected to redundant A/B power sources. In the event that one power feed fails, the ATS will switch to the other feed.

also valuable if your organisation has considered billing back to different departments for their power usage.

Three benefits of monitoring at the cabinet

There are many advantages to monitoring at the cabinet, but these three top our list as the biggest benefits.

1. Identification of stranded capacity ("zombie servers")

Monitoring power usage and identifying additional capacity in the current facility over the significant costs of adding additional capacity via new infrastructure is simply a much more cost-effective approach. However, the status of data centre power is often based on allocated power without real-time monitoring data. Power is often being consumed by IT devices that are no longer in use, sitting there idle, still feeding off the data centre's power which costs you money.

Removing these "zombie" devices can be difficult because it is mainly a manual identification process that requires significant staff resources. However, PDUs with outlet level metering can streamline the identification process by providing the data needed to quickly locate the devices.

2. Outage prevention

The use of intelligent rack PDUs can help notify your team of issues before they occur. Warnings and critical threshold settings ensure that the rack PDUs do not experience overload conditions that could otherwise trip the breaker and the connected loads. Additionally, intelligent PDUs can provide environmental information, informing data centre staff when temperatures are high enough to cause hardware failure.

3. Improved capacity planning

There are many issues that can be resolved if you know how much power your cabinets need, as well as the amount available. Knowing the IT load is also a key parameter in power efficiency metrics like the Power Usage Effectiveness (PUE) parameter created by The Green Grid.

One of the ways power usage data can assist in capacity planning is access to reports that identify cabinets with power and space availability for new devices. In addition, managers can identify cabinets that have exceeded or will exceed their capacity in the future, based on the current growth rate.

Power monitoring software

Unfortunately, many managers don't have the IT equipment and site infrastructure power monitoring data and valuable insight to help increase the efficiency of their data centre and avoid possible outages. Without that data, it's difficult to have a focused strategy for maintaining and improving efficiency, as well as mitigating downtime. Power monitoring software can be a tool to help managers and operators take advantage of readily available opportunities to substantially reduce energy costs and prevent outages, which will save operational expenses too.

When evaluating power monitoring software, it is important to understand and measure the areas in the power chain that are the responsibility of your business. In the case of an enterprise-owned facility, it's critical to have intelligent hardware that can measure at multiple points discussed throughout this report: entrance feed, UPS systems, room distribution, cabinet distribution and IT equipment levels to feed into your software.

BEWARE THE ZOMBIES

According to Uptime Institute's estimates based on industry experience, around 20 per cent of servers in data centres today are obsolete, outdated, or unused. They also estimated that decommissioning one rack unit (1U) of servers can save an organisation on average:

- 500 USD per year in energy costs
- An additional 500 USD in operating system licenses
- 1,500 USD in hardware maintenance.

BEST PRACTICE 5 OF 5

IT EQUIPMENT

5

The various IT equipment running in data centres represent a significant power load and expense. Naturally, the more IT equipment and upgrades invested into a data centre, the more this hardware needs a reliable and efficient power infrastructure to support all these hard assets.

Applying Anixter's intelligent power chain concept of the right hardware configuration matched to analytics for greater intelligence can result in substantial benefits in this area because so much of the power usage comes from IT equipment. Focus on these key areas to increase power efficiency and maximise savings.

High-voltage power supplies

The choice of power supplies and the use of high-voltage distribution are often evaluated when trying to achieve energy efficiency and optimum power usage. The selection of the proper input voltage has a direct impact on power supply output capacity, conversion efficiency, thermal operation, and reliability. All of these variables impact the bottom line.

Some considerations for high-voltage power supplies:

- One less transformer needed to step down voltage at the 208 V level
- For each 1,000 W power supply, there's 1 to 2 percentage points of difference in efficiency
- 208 V U.S. 230 V globally

According to the U.S. ENERGY STAR program, substantial savings can be gained from a single watt saved.

Server virtualisation

Virtualisation can help data centre operators reduce the overall IT equipment physical footprint by virtue of system consolidation. With virtualisation, IT power consumption is designed to drop given the lower number of servers in operation. However, more power per square foot due to higher density or kW per rack is consumed.

It is therefore imperative to ensure the power and cooling infrastructure are adjusted to accommodate the higher densities and loads in order to maximise savings. Also, there are more applications running on a limited number of hardware appliances, which often means that it is even more critical to maintain a high level of availability, making intelligent PDUs with monitoring capabilities even more important.

Equipment efficiencies

Improving the energy efficiency of your data centre equipment can reduce server power use with no impact on your runtime. Efforts in this area also have the potential for offering increased computer capacity, power savings, and business continuity.

ENERGY STAR servers

In comparison to the rest of the equipment housed in a data centre, servers comprise a substantial share of power usage. Deploying more efficient servers throughout the data centre can be an effective method for reducing overall energy consumption. Additionally, as a byproduct of improved server efficiency, less heat is generated, which can reduce cooling costs.

The more IT equipment and upgrades invested into a data centre, the more this hardware needs a reliable and efficient power infrastructureto support hard assets.

IT EQUIPMENT

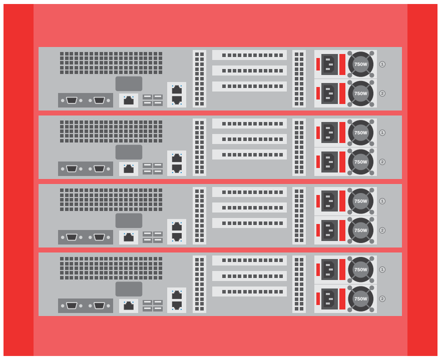

In 2009, ENERGY STAR released its first energy specification for computer servers. To earn the ENERGY STAR certification, servers must offer the following features:

- Efficient power supplies
- Improved power quality
- Capabilities to measure real-time power use, processor utilisation, and air inlet temperature
- Advanced power management features
- A power and performance data sheet that standardises key information on energy performance, features, and other capabilities

Power capping

Server workloads can vary depending on level of use. Data centre managers generally allocate enough power per rack to ensure performance in, in case all the servers are at 100 per cent utilisation. Unfortunately, this approach results in stranded power capacity that could be used somewhere else in the data centre. Power capping allows IT managers to set a threshold that servers cannot exceed. This ensures that the power consumption of all servers in the rack will not exceed the available capacity, which will usually trip a circuit breaker. The advantage to this approach is that it allows data centre managers to effectively allocate available power closer to actual usage, without having to worry that capacity would be exceeded. This frees up power to be used where it is most needed.

Intel® Data Centre Manager

One platform that may be particularly helpful with improving the efficiency of IT equipment, particularly in the area of power capping, is the Intel Data Centre Manager and Intel Intelligent Power Node Manager.

Intel Data Centre Manager (DCM) is a companion to intelligent PDUs. DCM provides insight into the server itself. It provides real-time, accurate power and thermal consumption data, enables management of data centre hotspots, and allows for power usage planning and forecasting. Also many DCIM software platforms are beginning to integrate directly with DCM so that data centre operators get the power and cooling info needed all the way down to the server itself.

A well-designed strategy for hardware refreshes for your data centre equipment is one of the cornerstones of improving efficiency.

Intel Intelligent Power Node Manager is an out-of-band power management policy engine. It enables regulation of individual server power consumption (power capping) through modulation of the processor's performance (P) and throttle (T) states.

Hardware refresh cycles

A well-designed strategy for hardware refreshes for your data centre equipment is one of the cornerstones of improving efficiency. Executing on it will:

- › Mitigate outage risks
- › Improve server performance in a smaller footprint
- › Simplify your IT infrastructure through standardisation
- › Improve capabilities, including energy-efficiency technologies and operating modes.

The never-ending challenge with hardware refresh cycles is balancing the performance and energy efficiency benefits with the impact on your IT budget. Make sure you weigh all the pros and cons, and take the time to carefully consider what is optimal in terms of both equipment choices and cycles time span for your unique data centre needs.

Grounding and bonding

Grounding should be addressed in the following areas: electrical distribution systems, IT cabinets and racks, and HVAC systems. If properly designed and built, the ground system is essentially a radial system from the electrical service entrance.

Building grounding systems should be directly bonded to all major power distribution equipment, including switchgear, generators, UPS systems, and transformers. Your facility should possess a building electrical main ground bus (MGB) where all the large-load feeder facility ground terminates.

Some additional grounding and bonding recommendations:

- › Make sure to ground all dead metal objects with the data centre.
- › Where there are multiple power service entrances, the ground ring conductor should be sized at 107 mm (4/0 AWG) minimum bar copper.
- › All below-grade grounding connection should be made by NRTL-approved methods, such as exothermic weld or high-compression connectors.
- › Ground bus bars should be placed to facilitate bonding and visual inspection.

Also, supplementary bonding and grounding methods should be deployed to improve facility and equipment performance. Examples of supplementary components include metallic raceways, racks and cable trays, under the raised floor or above the cabinet and rack metallic grid work, metal plates and sheets, multiple bonding conductors from equipment to a grounding, or bonding structure.

THE FUTURE OF DATA CENTRE POWER GENERATION

Data centres continue to get more efficient in the way that they consume energy. Yet, there are many significant challenges that lie ahead. With the trend in consolidating smaller data centres into larger facilities, how are operators going to get the power that they need to operate? Also, as data centres start to become larger, being inefficient starts to exacerbate and multiply power issues and expose vulnerabilities and weaknesses. And what about getting power from sustainable sources?

Sustainable energy trends
There is currently a transition to clean energy, partially due to social and governmental pressures. Additionally, because energy is the largest data centre cost and every environment is unique, different methods of energy generation are being considered.

Where will the energy that powers tomorrow's data centres come from? There's no crystal ball, and innovations in this area are fast and furious, so it's impossible to predict with certainty how facilities will be powered years from now. However, sustainable energy has growing momentum.

Here are some recent examples of data centres that use sustainable energy:

- Green Mountain Norway – hydroelectric
- Datadock France – geothermal
- Facebook Texas (2016) – wind

There has also been a rise recently in renewable energy investment. Examples of that are:

- 51,000 MW of wind was added globally in 2014, a 44 per cent increase over 2013.
- Solar costs have fallen 80 per cent globally since 2008.
- Solar and wind provided 55 per cent of new electricity generation capacity in the U.S. during 2014.
- China invested 90 billion USD in renewable energy in 2014, a 32 per cent increase over 2013, according to Greenpeace.

Fuel cell technologies
The adoption of hydrogen fuel cells is another fascinating area where there appears to be some significant investment. The trends suggest that adoption will occur first in applications where generator use is impractical, such as remote locations or confined environments.

For example, Microsoft currently has some IT equipment cabinets powered entirely by fuel cells bringing the power plant inside the facility thereby minimising power distribution losses.

Another telling statistic is that, in 2013, fuel cell industry sales generated revenues of approximately 1.3 USD billion, according to the U.S. Department of Energy. Fuel cell system revenues grew by 35 per cent over 2012, with significant growth seen both in North America, with a revenue increase of about 50 per cent over 2012, and Asia, with about 33 per cent growth over 2012. Europe showed a slight decline in fuel cell system revenues.

SUSTAINABLE ENERGY TRENDS

Hydroelectric

Geothermal

Wind

THE FUTURE OF DATA CENTRE POWER GENERATION

Battery technologies
According to IHS Technology, lithium ion batteries have recently been gaining more adoption in the market. However, they are in the beginning stages of development and have noticeable disadvantages when compared to lead acid batteries, such as a weaker safety record, low-current discharge, and higher costs. Currently, lithium-ion battery technology is enjoying a boost in research and development in the electrical vehicle market, which is helping to drive down manufacturing costs.

In recent news, Elon Musk and Tesla® have made headlines in their effort to potentially bring their Powerpack™ battery technology to data centres. Data centre operators will be more interested in the larger-capacity Tesla Powerpack, a 100 kWh battery storage system, which, according to Musk, is infinitely scalable. Several large-scale organisations have already started pilot programs with Tesla.

Figure 8: Fuel Cell System Revenue (US$M)

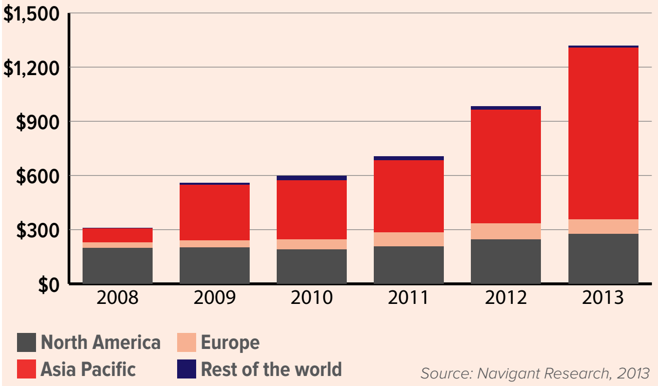

Source: Navigant Research, 2013

CONCLUSION

This report covered the complex nature of powering the data centre. The drive for improved efficiencies and managing variable IT capacity requirements are top priorities for all data centre operators and managers. The intelligent power chain achieved via the right hardware configuration coupled with intelligence and analytics can provide robust gains in five key areas. Additionally, intelligent data centre design decisions – whether brownfield or greenfield – can boost overall efficiencies and mitigate the risk of outages.

Breaking the intelligent power chain down into five best practices helps to simplify things, however each area requires attention to multiple variables, and sometimes a subset of variables.

But the key takeaway is not the complexity of all the moving parts that go into an intelligent power chain. The most important thing to hold onto is the fact that each piece of the puzzle can result in improvements in energy efficiency, performance, and equipment longevity. And these improvements can range from incremental to substantial in terms of impact and savings. Taken as a whole, and as you check off the list of what constitutes an intelligent power chain, your data centre can become an ultra-efficient and highly reliable facility that maximises energy usage to everyone’s benefit.

GLOBAL TECHNOLOGY BRIEFING

THERMAL EFFICIENCY BEST PRACTICES

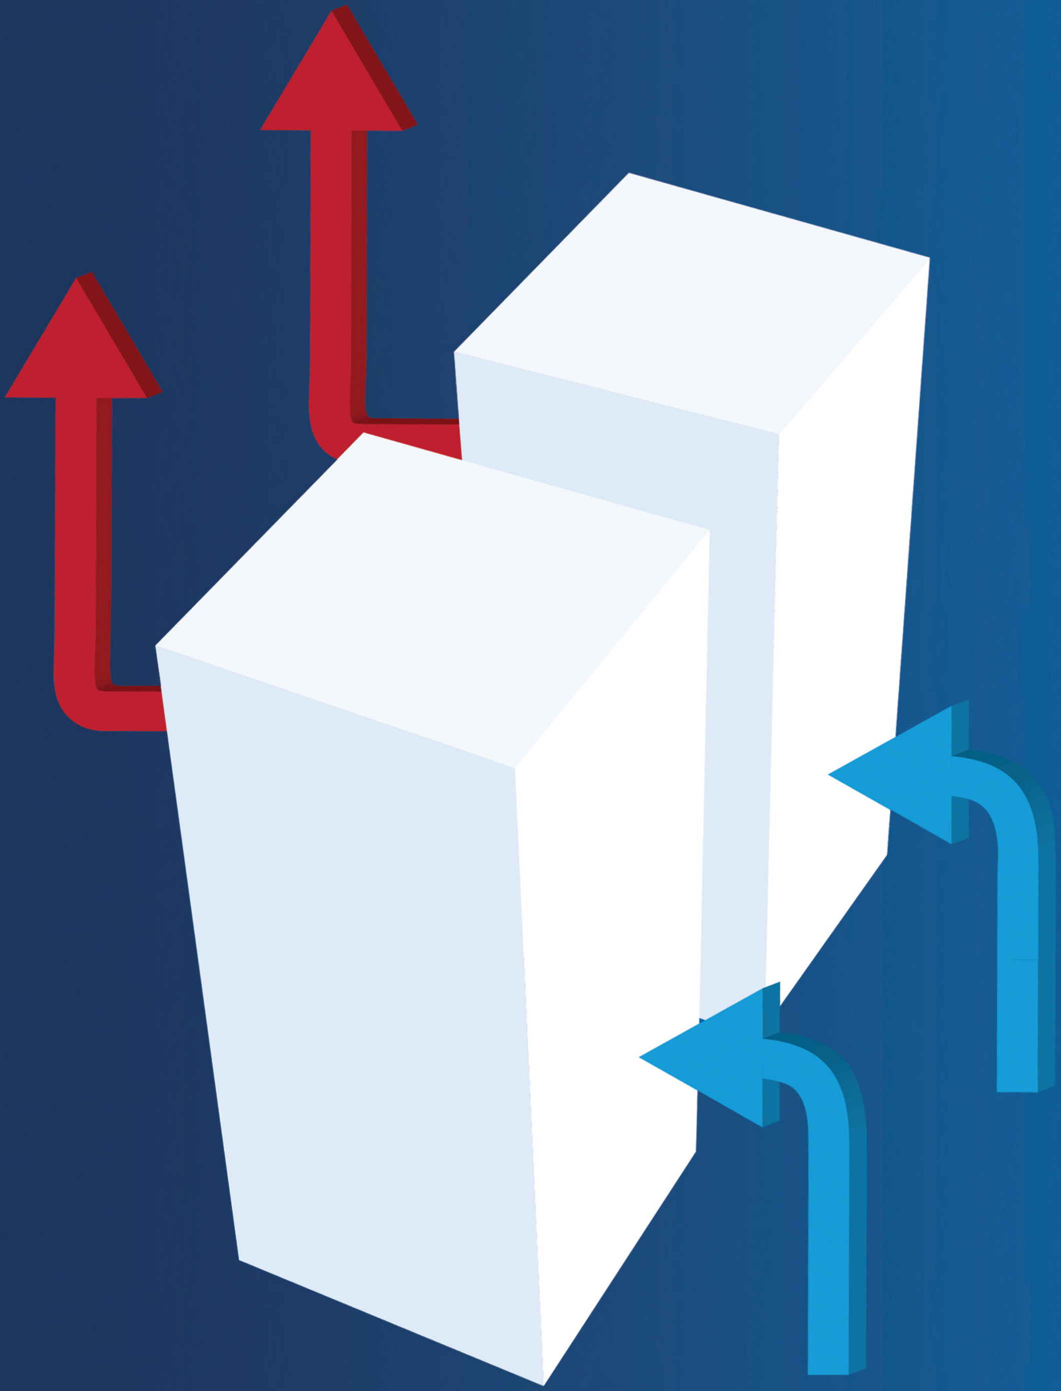

EXECUTIVE SUMMARY

Out of all the major facets that keep a data centre running and performing at optimum levels, thermal management is often considered a low priority compared to other functions, and it is often ignored unless problems develop. For those that work within the confines of a facility and understand the ramifications of the thermal environment upon highly sensitive, hot running equipment, having an efficient and effective thermal management strategy has become an important aspect of data centre management.

The method of cooling the equipment inside the data hall has always been a critical piece in any data centre's design. However, what is different today is that data centre facilities managers are under more scrutiny to cut costs and improve equipment reliability while supporting a growing IT demand. In order to meet those challenges, data centres are now being designed to grow with the business' needs. This includes the cooling system, which is being designed to scale as needed.

However, there are thousands of existing facilities out there that were designed 10, 15 and even 20 years ago, which still support their businesses and that need to adapt at a reasonable cost. This is where airflow management best practices can have an impact by making those environments more efficient and reliable to support today's IT equipment for a minimal investment.

In a 2014 report titled Trends in Data Centres, over 220 professionals were asked to describe one area they would change about their data centres. 19 percent cited more energy efficient and 15 percent stated a better cooling/HVAC system[1]. However, many facilities already have the cooling capacity on hand; it is poor airflow management that is not allowing the cooling system's full capacity to be used effectively. This report tackles the challenges, standards and best approaches to thermal efficiency in existing facilities largely due to them being more prevalent when compared to a new facility. You will be introduced to four best practices that are absolutely critical to achieving an optimal thermal environment. The report also explores the future of data centre cooling and introduces emerging technologies.

The approach to thermal management in the data centre is holistic, and each individual piece has a substantial effect on the entire ecosystem. The different variables that go into cooling a data centre are inextricably interrelated on a constant basis – even a brief lapse in the performance of one system can result in a breakdown in the rest of the areas. The many facets of thermal management will be explored from this all-inclusive perspective.

Over

200
PROFESSIONALS
WERE ASKED TO DESCRIBE

1

AREA THEY WOULD CHANGE
ABOUT THEIR DATA CENTRES

19%
CITED MORE
ENERGY EFFICIENT

15%
STATED A BETTER
COOLING/HVAC SYSTEM

[1] Insights Into What's Next: Trends in Data Centres 2014 – Mortenson

INTRODUCTION

WHAT'S DRIVING COOLING AS A PRIORITY

A lot has changed in data centre cooling in the last few decades. While once considered a simple design consideration akin to any building's comfort system, today's data centres now require a highly specialised approach to cooling, mainly driven by the overall increase in IT computing demands. However, there are several other driving forces that make thermal efficiency a high priority:

- In legacy data centres, the cooling system consumes a large amount of energy.
- Imbalanced room temperatures increase the likelihood of equipment failures.
- Increasing server rack densities can create unique thermal challenges.
- Changing IT requirements require cooling to be available on demand.
- High-availability environments need well-controlled temperature ranges for reliable performance.

Figure 1: E.U. Data Centre Power Consumption[2]

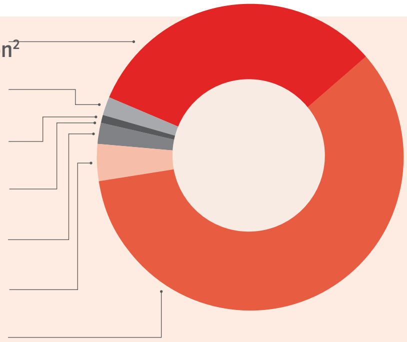

Cooling fans, pumps and compressors	**32%**
Lighting and small power	**2%**
Ventilation (fresh air)	**2%**
Communications	**0%**
Security, NOC, BMS and outdoor lighting	**2%**
Distribution and conversion lossess	**4%**
IT terminal load	**59%**

LEGACY COOLING PRACTICES

Before diving into today's thermal management challenges and solutions, it's important to create some historical context around data centre cooling. As data centres continue to evolve, the approach to cooling is often, unfortunately, based on outmoded cooling theories and methodologies, making some historical perspective essential.

A BRIEF HISTORY OF DATA CENTRE COOLING

From the first data centre built in the 1960s – mainly limited to space technologies for defence and other government entities – designers understood the two factors needed for successful operation were power and cooling. However, of these two factors, power most often took the hefty share of the limelight and attention, as facility owners focused on the quality, availability and reliability of power. Cooling was left to the IT equipment manufacturers to control through water-cooled CPUs.[3]

As data centres evolved alongside advancing power grids throughout the next two decades, the dependability and quality of power became less of a concern. Even though cooling started as a secondary issue, it evolved into a greater and greater challenge, especially as computing demands ramped up exponentially.

In the early 1990s, individual servers were introduced into the computer rooms, coming from offices and data closets. This move resulted in a shrinking form factor and an extraordinary increase in data centre power densities. The explosion of IT demands exacerbated the problem. No longer could a cooling system built for comfort also accommodate these new heat loads. To address this trend, more and more data centres were designed with massive chillers and air handlers. Because of the loud noise and other environmental issues this cooling equipment caused, isolated stand-alone computer rooms were developed to protect the sensitive IT equipment. This has forced facility, IT and maintenance staff to move to offices outside the computer rooms.

[3] Energy Efficient Thermal Management of Data Centres – Yogendra Joshi and Pramod Kumar 2012

LEGACY DATA CENTRES

In taking a closer look at how data centres were designed 15 to 20 years ago, there have been a number of factors that have contributed to a lack of efficient cooling. One common issue that often arose during the planning of a facility was the selection of an architectural and engineering firm that fully understood HVAC design practices but didn't have a complete understanding of the best practices around controlling airflow throughout the data hall itself.

As facilities evolved, the initial planning that went into building data centres has become more highly specialised. As with many other aspects of data centre design, engineers and planners without this specialised data centre airflow management training could forgo tuning the cooling system for maximum efficiency from the beginning.

In addition, with the many foreboding forecasts of dramatic data loads exploding, data centres were often designed to cram as much hardware onto a raised floor as possible. Maximising the footprint was the mantra, and cooling systems were designed to handle the maximum anticipated heat load. All perimeter cooling units were generally turned on once initially deployed, with cold inlet temperatures, which meant the room was getting far more cooling than the IT equipment required and which added significantly to the operational cost of the facility.

In terms of the type of cooling systems used most typically by data centres in the past, these are generally limited to two perimeter methods:

- Direct expansion (DX)
- Chilled water

DX COOLING SYSTEMS

DX cooling systems use a heat removal process that uses a liquid similar to what is used to cool a car or refrigerator. In the past, a majority of data centres have been cooled by packaged DX systems, which are often referred to as computer room air conditioners (CRAC). They were attractive to facility owners due to easy off-the-shelf options, the relatively small form factor and straightforward operation.

As data centre cooling technologies have advanced and became larger with higher total heat loads, more and more owners and operators during the past decade have chosen other technologies over DX that are capable of producing the same amount of cooling more efficiently.

Nevertheless, DX systems are still a viable option for many computer rooms. DX systems are now being configured to take advantage of economiser (free) cooling. However, this requires extra equipment to be installed with the cooling units, which increases the system's cost.

CHILLED WATER COOLING SYSTEMS

Chilled water designs use a centralised chiller to produce cold water. The chilled water is then piped to computer room air handlers (CRAH) with heat dissipation being handled by the cooling towers. Chilled water in the past was often the choice for large, multistory facilities, mainly because it was more efficient and – for the size of the facility – was less expensive.

CHAOS AIR DISTRIBUTION

The methodology of chaos air distribution is simple: arrange cooling units around the perimeter of the server room and produce a large volume of chilled air. The thought was that the turbulence of the air movement under the raised floor would create a homogenous supply air temperature.

This approach had the intention of cooling IT equipment while pushing out hot server exhaust air toward the facility's return air ducts. What was later discovered was the cooling units ended up creating a confounded airflow system under the floor, which resulted in no control over supply side temperature throughout the data hall.

Other air distribution inefficiencies associated with chaos air distribution include:

- **Bypass airflow**
  Cool supply air never enters the IT equipment. The supply air either is lost before it gets to the cold aisle, or it does enter the cold aisle but travels past the IT equipment.
- **Recirculation**
  Hot exhaust air re-enters the server intakes prior to being reconditioned by the cooling system. This is the leading cause of data centre hot spots.
- **Air stratification**
  Trying to direct cool air to the top of the rack face creates air masses in different temperature-based layers, which may force the lowering of set points on the cooling equipment.
- **Multiple underfloor temperatures**
  With each cooling unit being controlled locally, each unit could provide a different supply air temperature; these variations confound the ability to control the input air temperatures across the computer room.
- **Wasted cooling capacity**
  The additional cooling required to satisfy the heat load in the room is primarily due to the large amount of bypass airflow in the room.

Figure 2: Chilled Water Cooling System (Simplified)[4]

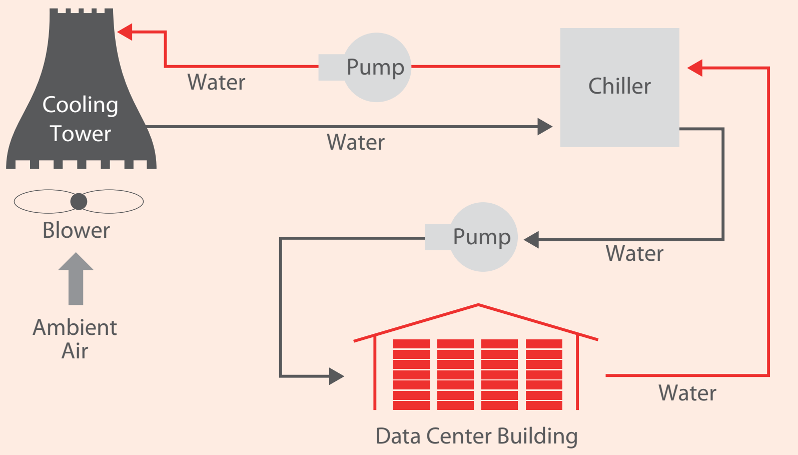

Figure 3: Imbalanced Supply Temperature[5]

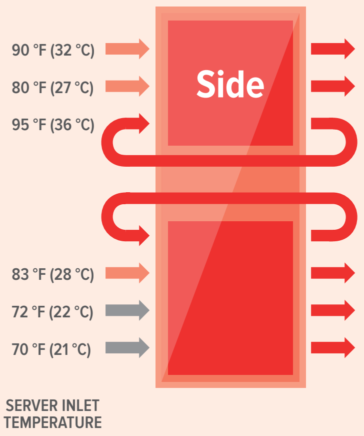

[4] Energy Consumption of Information Technology Data Centres - Madhusudan Iyengar 2010

[5] Properly Deployed Airflow Management Devices – ENERGY STAR

LEGACY AIRFLOW DESIGN

One of the hallmarks of legacy data centre airflow design is a proliferation of holes, gaps and obstructions in the supply air pathway, creating bypass airflow. Some examples include:

- Open cable cutouts
- Open holes underneath the cabinet
- Holes at the perimeter walls of the room under the raised floor
- Holes behind the air handling units
- Holes above the drop ceiling
- Additional perforated tiles than needed to properly cool the IT equipment

What is the impact of these tactics? All of these openings are the vehicles that produce bypass airflow, which represents conditioned air that never enters the IT equipment. In order to make up for the air that never reaches the intended heat load, the cooling units have to work harder to deliver enough air, which creates inefficiencies and can lead to unstable temperatures throughout the data hall. According to a 2013 Upsite Technologies survey, the average data hall had 48 percent bypass airflow.[6]

Much of these inefficient tactics in legacy data centres can be attributed to a lack of specialised knowledge on the science of airflow by data centre designers, mechanical engineers and facilities technicians.

LEGACY COOLING MENTALITIES STILL AROUND TODAY

You'll discover in the remainder of this report the importance of proper airflow management and the negative impact it can have on the cooling system.

> Fortunately, there is much that can be done in many existing data centres with a relatively minimal investment to improve the reliability and efficiency of the cooling system.

Yet, despite more informed approaches and the availability of better solutions, data centres take an approach that results in overprovisioning data centre cooling resources, instead of looking at the airflow throughout the rack, row and room. Why?

- **Lack of knowledge**
  Cooling data centres is a specialised field, and many facility owners simply aren't aware of the many subtleties and implications.
- **Aversion to risk**
  Many data centre managers avoid taking the risk of trying a new approach.
- **Complacency**
  Unless systems go down as a direct result of lack of cooling, many managers are hesitant to rethink their systems.
- **Fear**
  Many data centre managers simply do not want to risk modifying a cooling system for fear failures could lead to systems going down.
- **Apathy**
  This can be the case when the facility manager's function does not own the energy portion of the budget.

Inevitably, advances in cooling systems specifically designed to address the unique challenges of the data centre environment will force legacy data centres to address thermal efficiency. If not, increasing outages, simple economics, the need to cuts costs and the trend toward environmentalism will certainly force data centres to embrace better thermal management practices. Fortunately, there is much that can be done in many existing data centres with a relatively minimal investment to improve the reliability and efficiency of the cooling system.

[6] Lars Strong – Myths of Data Centre Containment – 2015

COMMON COOLING CHALLENGES

There are several challenges that make it difficult to keep data centres at an ideal temperature.

- **Increasing cabinet densities**
  Cabinet densities are on the rise for many data centres today. Although they aren't rising as much as once thought, there are several applications that require a large investment of power. These cabinets can require a different approach to cooling than the rest of the environment.

- **Operational budget cuts**
  Many data centre managers are being asked to reduce operational expenses and think that increased thermal efficiency requires significant capital investment.

- **Lack of knowledge of airflow management best practices**
  Just understanding the right techniques can be a challenge. The impact of deploying blanking panels, removing cabling from under the floor and using cable-sealing grommets can pay huge dividends.

- **Matching cooling to IT requirements**
  An efficient cooling system means that the right amount of cooling is being delivered to satisfy the IT equipment's demands. Because IT's requirements change dynamically, the cooling system should be adjusted frequently, but the information required to do that isn't always provided or accessible.

- **Overwhelming thermal design considerations**
  There are a lot of options and methodologies out there to cool a data centre. In addition, there are several options to separate supply and return air. In light of this, choosing the best approach can be difficult.

THE COST OF COOLING A DATA CENTRE

There is growing evidence to support the fact that IT equipment is no longer the primary cost driver of running a data centre. Operational expenses related to powering and cooling the facilities infrastructure are fast overtaking hardware costs. According to ASHRAE, part of the reason for this is the fact that server performance and efficiency has increased substantially in recent years while their costs have remained relatively constant. According to a report conducted by Christian Belady, during an eight year period server performance increased 75 times for the same hardware cost. In the same study, the performance per watt was shown to increase by 16 times in a typical server during that same period.

This is a paradigm shift in terms of how the C-suite perceives the economics of running a data centre, impacting operational decisions as they relate to power and cooling. It's all about reducing the overall total cost of ownership (TCO) and having an efficient thermal environment is becoming a huge factor in achieving that goal.

On average, the energy cost to cool a data centre is substantial. On the low end of the spectrum are estimates in the range of 30 percent or less, even though many data centres lacking efficient cooling systems and tactics can consume nearly 50 percent of energy costs.

Figure 4: Performance and Performance/Watt Increase in a Typical Server[7]

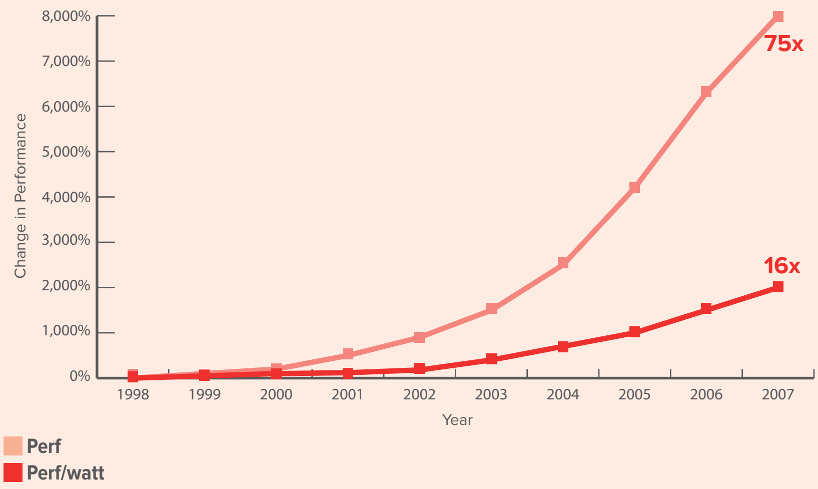

[7] In the data centre, power and cooling costs more than the it equipment it supports – Christian L. Belady 2007

HOW POOR COOLING AND AIR DISTRIBUTION STRATEGIES CAN IMPACT OPERATING EXPENSES

The cost of the cooling system can be greater than 30 percent of the total energy used in the data centre. The main culprit within the cooling unit is generally the compressor and fans. There are two ways to reduce the energy being consumed by the cooling units:

- Increase supply-side temperature (compressor energy)
- Decrease the volume of air being delivered to the IT equipment (fan energy)

Before any of those actions can be taken, best practices around airflow management must be implemented. If there is considerable bypass airflow and recirculation happening in the data hall, then increasing the supply-side temperature and decreasing the air volume being delivered can have devastating effects. The good news is that many existing data centres can implement these best practices and see significant gains in improvement of their systems.

In 2002, an Uptime Institute study found that an average of 60 percent of computer room cooling was going to waste through bypass air not passing through the IT equipment. In a 2013 follow-up study, Upsite Technologies found little improvement, estimating that an average of 48 percent of supply air is bypass airflow. What was also interesting was that in 2002 there was approximately 2.8 times the required cooling capacity available and in 2013 that figure rose to 3.9. Cooling capacity isn't the issue: it is inefficient cooling systems and poor airflow management practices.[8]

Table 1: The State of Airflow Management

2002 UPTIME INSTITUTE RESEARCH	
Bypass Airflow	60%
Hot Spots	10%
Cooling Capacity	**2.6x**
2013 UPSITE TECHNOLOGIES RESEARCH	
Bypass Airflow	48%
Hot Spots	20%
Cooling Capacity	**3.9x**

Cooling capacity isn't the issue. It is inefficient cooling systems and airflow management that is the problem.

[8] The State of Airflow Management – Upsite Technologies 2015

COOLING AND AIRFLOW MANAGEMENT GOALS

When talking about the cost of cooling (and thermal management in general), it's important to note the concrete differences between two concepts: active cooling and airflow management.

Active cooling systems and equipment are actually cooling the return air from the IT equipment. These systems generally include perimeter cooling such as CRAC or CRAH systems, close-coupled cooling (in-row) and rear door heat exchangers (in-rack). The goal of any cooling system is to supply enough conditioned air at as high a temperature as possible to allow for a reliable equipment operating environment.

Airflow management manages the supply air that comes from the cooling equipment to the IT load, and the return air that exits the IT load and goes back to the cooling equipment. Tactics include aisle containment systems, perforated floor tiles, blanking panels, airflow sensors and more. Airflow management's goal is to deliver the conditioned supply air required by the IT equipment, ensure it only passes through once and get it back to the cooling unit so it can be reconditioned.

A data centre manager knows when he has achieved an optimised thermal environment when:

- Bypass airflow is less than 10 percent
- The temperature measured at top and bottom of the IT equipment cabinet has a delta of less than 5 degrees Fahrenheit.

Figure 5: Once-Through Cooling[9]

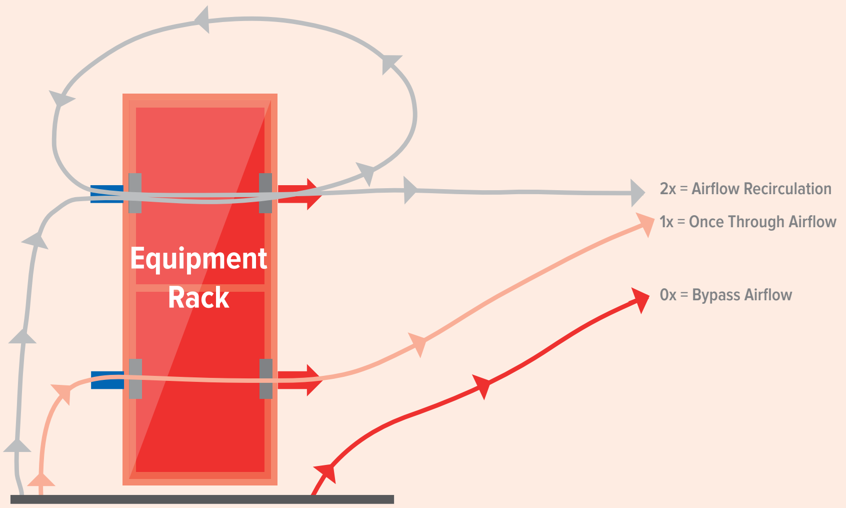

Room-Level Energy and Thermal Management in Data Centres: The DOE Air Management Tool – Magnus K. Herrlin 2010

Figure 6: Properly Balanced Supply Side Temperature

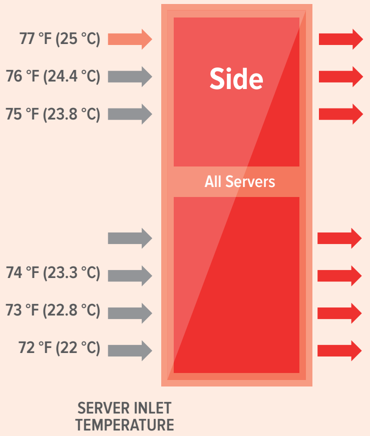

Better airflow throughout the data hall means higher temperatures and the less airflow volume is required to satisfy a given equipment load.

[9] Room-Level Energy and Thermal Management in Data Centres: The DOE Air Management Tool – Magnus K. Herrlin 2010

EFFECTS OF GEOGRAPHY ON FACILITY COOLING

While looking at aspects of data centre cooling, it's worth exploring the impact of the physical location of a facility on cooling performance and costs. Whether you own or manage an existing facility and are looking to take advantage of the potential cooling benefits that its climate may afford, or actually determining a location for a greenfield data centre, geography can play a major role in achieving a cost-effective thermal environment.

Chief among the geographical opportunities is the potential for a data centre to leverage free cooling systems, such as economisers These systems rely on outside air, provided the conditions are right.

There are four primary climate-related variables that can affect data centre cooling:

- Temperature: average outside air temperatures that are low enough to allow for the use of an economiser cooling system to be employed
- Humidity: ideal level is not so high as to require continuous dehumidification, and low enough to allow for adiabatic cooling
- Contamination: due to air quality issues
- Corrosion: due to poor air quality

AN IMPORTANT NOTE ON GEOGRAPHY

Of course, besides these four climate-related factors, it's worth noting that the geographic location of data centre as it relates to thermal management also hinges on the availability and quality of power today and in the future. Local energy rates are an important factor. Proper vetting of the power grid should be conducted before choosing a location or a cooling system.

Figure 7: **Biggest Factors That Drive New Data Centre Investments**[10]

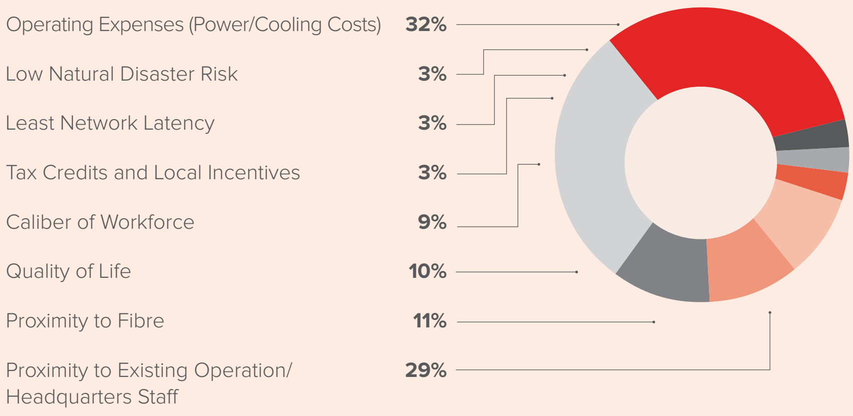

MEASURING COOLING SYSTEM EFFICIENCY

There are several different ways to measure the efficiency of a data centre's use of power, the most popular being Power Usage Effectiveness (PUE). However, while the use of PUE measurements are accurate for measuring power efficiency overall, there are several additional steps that need to be taken to accurately define the efficiency of the specific cooling system within a data centre. Below are some common methodologies to measure the effectiveness of the cooling system.

POWER USAGE EFFECTIVENESS

In 2007, the Green Grid Association published an energy-efficiency metric called Power Usage Effectiveness[11]. It is measured by dividing the amount of power entering a data centre by the power used to run the IT infrastructure. It's expressed as a ratio, with overall efficiency improving as the number is closer to one.

The challenge with PUE as it relates specifically to the cooling system is that it measures overall energy efficiency of the entire facility. This makes it difficult to quantify exactly how efficient or inefficient the cooling system is. Nevertheless, capturing and recording PUE can be a benchmark metric that will help gauge the effectiveness of any step taken to increase the cooling system's efficiency.

Figure 8: Power Usage Effectiveness[12]

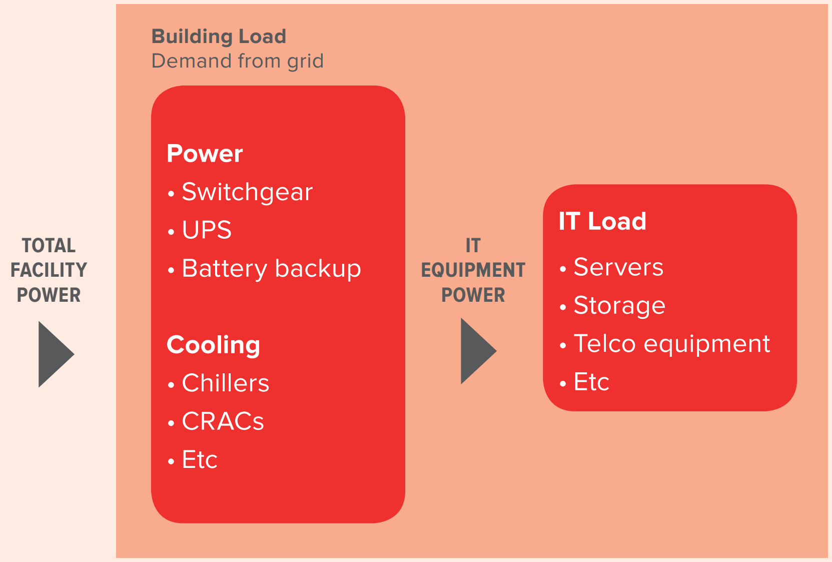

$$\text{PUE} = \frac{\text{Total Facility Power}}{\text{IT Equipment Power}}$$

COOLING CAPACITY FACTOR (CCF)

Cooling Capacity Factor or CCF[13] is a metric that was developed by Upsite Technologies to estimate the utilisation of the cooling system within the data hall. This metric allows a data centre manager to benchmark the rated cooling capacity versus how it is being utilised. Having this information allows data centre managers to understand if their cooling issues are due to a lack of available cooling capacity or poor airflow management practices.

OTHER IMPORTANT THERMAL EFFICIENCY METRICS

Here is a list of other metrics that are important to know when evaluating the overall efficiency of the cooling system[14]:

› Raised floor bypass open area
› Percent of data centre cabinets experiencing hot and cold spots
› Percent of perforated tiles properly placed in the cold aisle versus the hot aisle
› Ratio of supply air volume to IT equipment airflow volume

[11 & 12] PUE™: A COMPREHENSIVE EXAMINATION OF THE METRIC – The Green Grid 2012

[13] Cooling Capacity Factor (CCF) Reveals Stranded Capacity and Data Centre Cost Savings – Kenneth G. Brill and Lars Strong 2013

[14] Lars Strong – Myths of Data Centre Containment – 2015

DATA CENTRE COOLING STANDARDS

Standards for data centre thermal management have evolved substantially in the last few years, particularly in the areas of optimal temperature range as it relates to IT equipment. There are several organisations with useful published standards and guidelines, with the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) being highly prominent globally. As data centre owners and operators aim to achieve the best possible thermal environment, these organisations can be a valuable resource with well-researched, documented and methodical standards.

ASHRAE

ASHRAE is heavily involved with the development of modern data centre facilities. The ASHRAE Technical Committee 9.9 created an initial set of thermal guidelines for the data centre in 2004 with two follow-up editions in 2008 and 2011. Each new set of guidelines provided guidance on the different classes of servers and other data centre equipment that could function in higher acceptable levels of operating temperature and humidity for greater cost savings and efficiency. However, it's important to note that the standards have to be properly implemented in order to avoid heat-related outages and other potential issues.

THE GREEN GRID

The Green Grid is not a standard per se, but it is an international consortium dedicated to improving energy efficiency in data centres and business computing ecosystems. PUE and its reciprocal, DCiE, are widely accepted benchmarking standards proposed by the Green Grid and the Department of Energy respectively to help IT professionals understand how efficient their data centres are and to monitor the impact of their improvement efforts.

DIN EN 1377

DIN EN 13779 is a European standard, and encompasses ventilation for nonresidential buildings, as well as performance requirements for ventilation and room-conditioning systems. This standard focuses on achieving a comfortable and healthy indoor environment in all seasons with acceptable installation and running costs. It also specifies the required filter performance in a system to achieve good indoor air quality (IAQ).

EUROPEAN COMMITTEE FOR ELECTROTECHNICAL STANDARDISATION (CENELEC)

The European Committee for Electrotechnical Standardisation developed a series of European standards for data centres. This document specifies recommendations and requirements to support the various parties involved in the design, planning, procurement, integration, installation, operation and maintenance of facilities and infrastructures within data centres.

The EN 50600-2-3 standard focuses specifically on environmental control and specifies requirements and recommendations for such factors as:

- Temperature control
- Relative humidity control
- Particulate control
- Ventilation
- Energy-saving practices

ENERGY STAR

ENERGY STAR is a U.S. Environmental Protection Agency (EPA) voluntary program that helps businesses and individuals save money and protect the climate through superior energy efficiency. ENERGY STAR has collaborated with other data centre standards organisations to establish guidelines and best practices in the area of energy management. The organisation has published recommendations specific to cooling in the areas of temperature and humidity adjustments, hot- and cold aisle layouts, air-side economiser, proper airflow management and more.

OTHER STANDARDS AND RESOURCES

There are a number of other applicable standards, guidelines and best practice documents and research related to data centre cooling that may be helpful to designer, engineers, owners and operators:

- BREEAM
- 80 Plus
- Code of Conduct Data Centres
- BICSI
- Uptime Institute
- CEEDA
- US Department of Energy
- OSHA
- VDI 2054 (Germany)

DATA CENTRE COOLING APPROACHES

Data centre owners, designers and engineers have many choices when it comes to methodologies and the cooling topologies of their facilities. This section is not intended to provide heavily technical details or comparisons between these different choices, but rather a snapshot of each method and a few of their chief advantages. This informed perspective of the current data centre cooling landscape will provide the basis for the solutions and technologies that can be layered to further increase efficiency and cost-effectiveness.

METHODOLOGIES

There are many different methods for cooling data centres, some being a hybrid of technologies such as traditional systems with an economiser option or mode of operation. Here's a quick overview of choices broken down into a few major categories.

WATER

Chilled water systems are generally used in data centres with a critical IT load of 200 kW and larger with moderate- to high-availability needs. They may also be used as a high-availability dedicated solution. In many facilities, chilled water systems are often used to cool entire buildings, even though the data centre uses only a small portion of the total cooling capacity. These systems use large chillers to produce cold water which cools the supply air. The chilled water is pumped into the CRAHs in the computer room, while the heat is transferred to the cooling tower through condenser piping.

AIR

Air-cooled systems are two-piece designs that include an air-cooled CRAC and a condenser, which is also known as an air-cooled CRAC DX (direct expansion) system. Air-cooled CRAC units are widely used in data centres, and are considered the norm for small and medium rooms. In this type of system, refrigerant circulates between the indoor and outdoor components in pipes. Heat is transferred outdoors by using this circulating flow of refrigerant.

Air-cooled self-contained systems are AC units that are joined within an air duct. Usually these systems feature all the components of the refrigeration cycle in one enclosure, most typically within an IT environment. Heat exits this system as a stream of exhaust air, which is routed away from IT spaces to the outdoors or unconditioned spaces.

FREE COOLING

According to a report by IHS, the market for free cooling technologies will grow faster than any other cooling method over the next four years.

The term "free cooling" is a misnomer. It refers to the ability to cool a data centre without using a refrigeration cycle. There are still pumps and fans operating in these systems which consume energy. Free cooling takes advantage of local ambient conditions to cool the supply side air that feeds the IT equipment.

According to a report by IHS, the market for free cooling technologies will grow faster than any other cooling method over the next four years. Additionally, many of these technologies are being added on to existing DX or chilled water systems as an optional operating mode.

There are three main types of free cooling systems:

Air-Side Systems

Outside air is brought into the data centre directly through filters or indirectly through heat exchangers. In direct air-side systems, filtering of the outside air is important in order to avoid particulate or gaseous contamination as well as to provide some form of humidity control.

Water-Side Systems

This system leverages a cooling medium, such as water or a glycol water mixture that circulates directly through cooling towers and/or condensers rather than the chillers or compressors. Water-side systems segregate outside from inside air, providing cooling through a heat exchanger. Depending on the location of the data centre, there is opportunity to leverage cold water from rivers, lakes or oceans instead of using traditional closed-loop water chiller systems. Evaporative cooling systems use a number of methods to evaporate water into the airflow path of the airside economising systems in order to lower the air temperature entering the computer room. This system is used to extend the hours of free cooling in areas with hot and dry climates.

Higher Temperatures for Chiller Set Points

Raising chilled water set points is more of a retooling tactic than an actual methodology, but it's a growing trend for facilities that use chillers. Raising the cooling set point in the data centre allows operators to also raise the temperature of the water in the chillers, reducing the amount of energy required to cool the water and increasing the amount of free cooling hours.

It was widely reported that Facebook retooled its cooling system in one of its existing data centres and reduced its annual energy bill by $229,000[15]. Some of the ways Facebook was able to do this was by:

› Raising the temperature at the CRAH return from 72 to 81 degrees Fahrenheit.

› Raising the temperature of chiller water supply by 8 degrees, from 44 to 52 degrees Fahrenheit.

It was widely reported that Facebook retooled their cooling system in one of its existing data centres and reduced its annual energy bill by $229,000.

[15] Facebook Saves Big By Retooling its Cooling – Data Centre Knowledge – 2010

COOLING TOPOLOGIES

When designing a cooling system, there are three main topologies or layouts to consider: room, row and rack. Each topology has benefits and drawbacks, but the important thing to remember is that there isn't one method that is necessarily more efficient than the other. Any system can be made efficient with proper airflow management techniques. Selecting one topology over the other comes down to the physical layout of the room and the power requirements of the individual IT racks. It is also important to note that these topologies can be mixed in the same environment, so the cooling system is more tuned to the IT application requirements.

ROOM COOLING

Room or perimeter cooling is the traditional and most prevalent topology used in data centres today. It involves using a cooling unit such as a CRAC or CRAH system that delivers the supply air to the IT load with or without a plenum to the cold aisle.

There are two main floor designs that will affect the delivery of the supply-side air from a perimeter cooling unit: raised floor and slab.

RAISED FLOOR DESIGNS

Delivering supply air through perforated floor tiles can be easily arranged. This design tends to allow for much easier execution of an efficient air distribution strategy and gives greater control of where the supply air is going.

According to data centre thermal efficiency expert Dr. Robert F. Sullivan, some tips that might increase efficiency in a raised floor environment include:

› Keeping the underfloor as clean as possible
› Removing any power and data cables from under the floor and move them overhead
› Removing any items being stored under the floor
› Matching the number of perforated tiles properly to the airflow, which is matched to the heat load
› Sealing and plugging all cable cutouts using grommets and foam blocks

SLAB FLOOR DESIGNS

These designs offer flexibility in the location of the computer room, but it makes it harder to distribute air properly. According to Dr. Robert F. Sullivan, when distributing air into a computer room built on a slab the benefit of going to contained aisles occurs at a much lower power dissipation level (e.g., 5 kW per cabinet).

PLENUM AND DUCTWORK

Plenum and ductwork costs need to be considered when looking at a thermal management status. Plenum spaces are the open spaces above the ceiling or below the floor that are used for air circulation. Attention to these spaces is crucial for optimum air circulation in data centres. Inadequate ceiling heights, poor ductwork design or an undersized hot air return plenum can cause airflow management issues.

ROW-BASED COOLING

Row-based cooling goes by several different names including close-coupled cooling, row-oriented architecture and in-row cooling.

The true purpose of row cooling is to capture hot IT exhaust air, neutralising the potential negative impact before it can mix with surrounding air or recirculate to the front of the IT rack. Row cooling units that can capture 100 percent of hot air can improve energy efficiency and eliminate hot spots. Another benefit of row cooling is that the supply and return air doesn't have to travel as long a distance as it would have to in a perimeter- or room-cooled environment, which means smaller, more energy-efficient fans can be used to reduce the operational costs of the system.

Some important points to keep in mind when designing a row-based cooling topology:

› **It does not require a raised floor**
One of the benefits of row-based cooling is that it can be installed in environments that lack a raised floor. Because the cooling system is located in the same row as the IT equipment, a supply air plenum isn't required.

› **It needs overhead or under floor piping**
Row-based cooling still requires chilled water or refrigerant to operate, which means pipes will need to be installed overhead or under the floor for each cooling unit, which creates additional cost and complexity.

Figure 9: Example of Row-Based Cooling

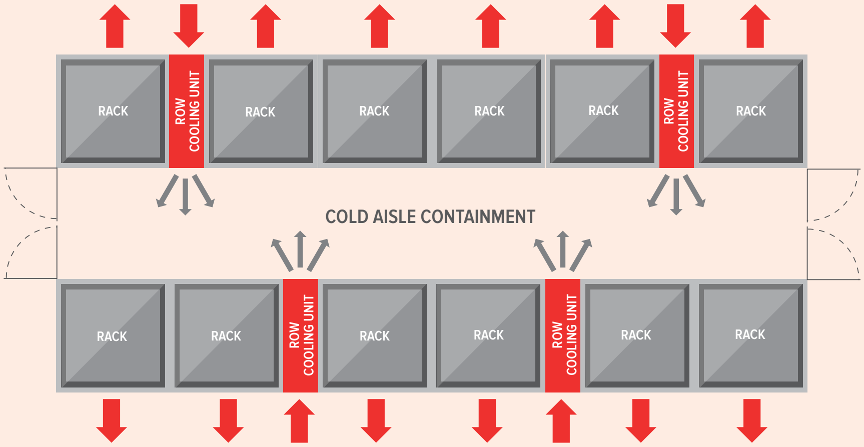

› **It's good for high-density applications and targeted cooling**
In environments that have high-density cabinet configurations, row-based cooling can be used to target those areas that can help alleviate strain on the perimeter cooling systems.

RACK COOLING

Rack cooling is when the cooling unit is brought inside the IT rack itself. There are two main drivers that relate to the implementation of rack cooling methodologies:

- High-density equipment cabinets
- Data centres built in nontraditional data centre environments

Sometimes high-density equipment cabinets produce too much heat for traditional methods used to cool the data centre. In those cases, rather than add additional perimeter cooling units, it might be more beneficial to add rack cooling.

Data centres that are constructed in rooms that might not be designed to support IT equipment can also benefit from in-rack cooling because like row cooling a raised floor isn't required. Additionally, the heat is largely contained within the rack itself and the exhaust air produced is conditioned air, which makes the rest of the room a comfortable temperature.

One method of rack cooling is the use of a rear door heat exchanger. Rear door heat exchangers use water or refrigerant to cool the IT exhaust air by passing it through a heat exchanger mounted directly to the rear of the cabinet. Fans then blow cool air out the rear of the cabinet and back into the data hall. By close coupling the rear door heat exchanger to the rack, the hot air is contained within the cabinet.

Figure 10: Example of Rack Cooling

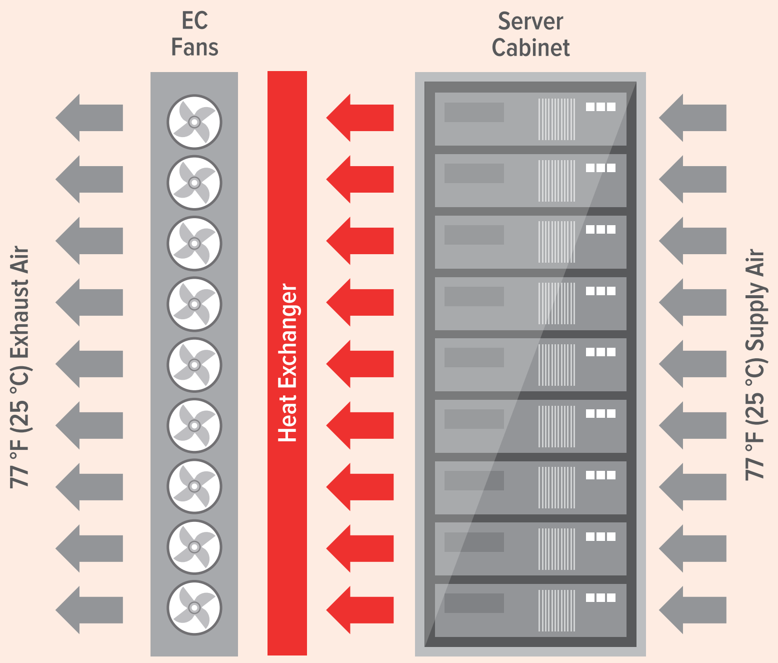

THE FOUR BEST PRACTICES OF CONDITIONAL ENVIRONMENTAL CONTROL

The goal of any data centre thermal management strategy is to make sure the room that the equipment resides in is at a temperature that is within range of the required operating temperatures specified by the equipment manufacturer. As stated previously, the old method of doing this is to flood the room with cold air, and if the room became too hot, add more perforated tiles, lower the temperature set points and finally add cooling capacity until the desired temperature is achieved. Due to the cost of operating a data centre and the rise in cabinet densities, data centre managers today are looking for a more efficient, reliable way of delivering air to the equipment in the cabinet.

Conditional environmental control is the process of delivering the exact amount of supply air at an ideal temperature and moisture content to maximise the cooling system's efficiency and improve equipment uptime.

There are several key drivers as to why a data centre would want to adopt this approach in a new or existing environment:

- Reduces the likelihood of hot spots and cold spots that can lead to equipment failure and added energy usage
- Regains stranded cooling and energy capacity, which allows for additional IT growth to support the business while minimising unnecessary capital expenditures
- Reduces operational expenses through the reduction in energy consumption by the cooling equipment
- Enhances the business' environmental image

The following sections will explain in detail these four best practices to achieve conditional environmental control:

1. Supply pressure
2. Supply temperature
3. Airflow segregation
4. Airflow control

Conditional environmental control is the process of delivering the exact amount of supply air at an ideal temperature and moisture content to maximise the cooling system's efficiency and improve equipment uptime.

BEST PRACTICE 1 OF 4

SUPPLY PRESSURE

A difficult thermal management challenge that many data centre operators have to face is how to deliver the proper amount of cooling that the IT load requires without oversupplying. Achieving a balance means the cooling system is using the minimum amount of energy required to deliver the supply air to where it needs to go, thereby reducing cost.

One reason why supply pressure is such a challenge is that the IT load is constantly in motion. Any change that takes place in the data centre — whether it is addition or subtraction of cabinets, IT equipment workload variation, containment systems, etc. — will have an impact on the airflow pressure within the data centre. Additionally, obstructions in the room, cabling under the floor in a raised floor environment and poor cable management inside of a cabinet can affect airflow pressure throughout the room. However, it is important that the amount of air being delivered by the cooling system is tuned to the environment so there isn't an over- or undersupply of air, which can lead to wasted cooling capacity or, worse yet, hot spots that can cause equipment failures through airflow recirculation.

Having a better understanding of the airflow pressure throughout the data hall will allow better identification of trouble spots, redirection of cooling to areas that might need it more, such as high-density equipment cabinets and verification of the effects of improvements made throughout the thermal management system.

This section of the report identifies some of the thermal management challenges that have an impact on supply pressure.

SUPPLY PRESSURE

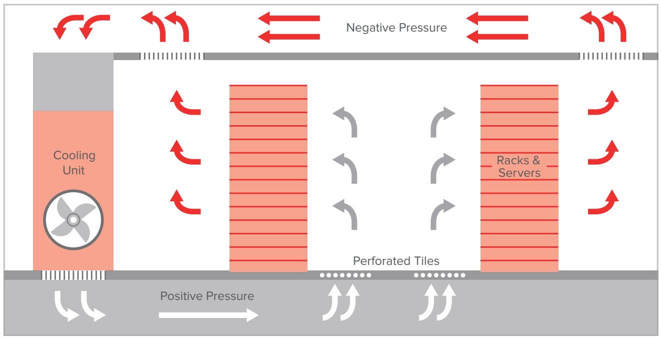

Figure 11: Bypass Airflow[16]

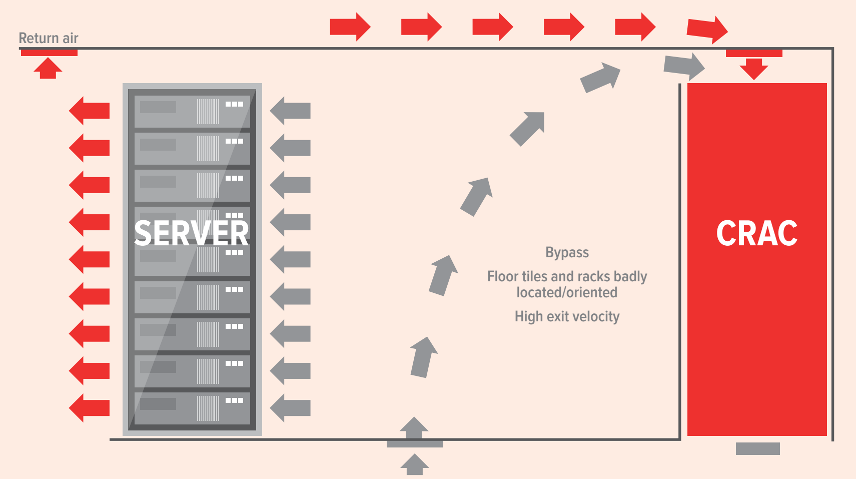

> A recent study of 45 computer rooms revealed that on average 48 percent of the conditioned air being generated by the cooling units is escaping from unsealed openings and inefficient or misplaced perforated floor tiles.[18]

Bypass airflow is conditioned air from a cooling unit that does not pass through the IT equipment before returning back to a cooling unit[17]. This is a problem because it reduces the static pressure under the raised floor, which means the cooling unit fans need to spin faster to deliver the same amount of air. This leads to excess cooling capacity because the actual heat load is being generated by the equipment to compensate.

Steps to reduce bypass airflow include:

- Measuring current bypass airflow by using weighted average temperatures from the cooling unit, server racks and ceiling return points — making sure to record measurements for each functional cooling unit [19]
- Walking through the data centre to visually identify the sources of bypass airflow, which can include cable openings in the raised floor, unsealed holes in the perimeter walls under the raised floor and holes behind the air handling units and at the bottom of building columns
- Sealing all identified sources of bypass airflow with floor grommets or fire-rated material for holes in the perimeter walls underneath the raised floor
- Ensuring the correct number of perforated tiles are only placed in the cold aisle, not in the hot aisle
- Remeasuring bypass airflow.

By reducing bypass airflow in the computer room, one can expect to achieve the following benefits:

- Reduced number of cooling units operating and the reduction in the operating expense of the cooling systems
- Improved equipment reliability through a reduction in hot spots
- Reduction in new cooling unit capital expenses
- Cooling system capacity gains for future expansion needs

[16] Improving Data Centre Air Management – Munther Salim and Robert Tozer February 2010

[17] Opportunities to Improve Cooling Capacity and Operating Costs – Upsite Technologies

[18] Upsite technologies research study

[19] Operational Intelligence Ltd, 2015

AIRFLOW RECIRCULATION

Airflow recirculation can be defined as hot exhaust air that passes through the IT equipment multiple times before it travels back to the cooling system. Recirculation is a problem because the supply air coming from the cooling equipment will heat up to a level that threatens the availability of the equipment, which can cause an outage or failure. These areas are generally referred to as hot spots. Some causes of airflow recirculation include:

- Absence of blanking panels in equipment cabinets
- Gaps in sides, tops and bottoms of cabinets
- Leakage through spaces left between cabinets
- Poorly designed airflow panels
- Noncontinuous rows
- Circulation over the tops of cabinets and rows
- Poor cable management inside the rack or cabinet.

Figure 12: Airflow Recirculation[19]

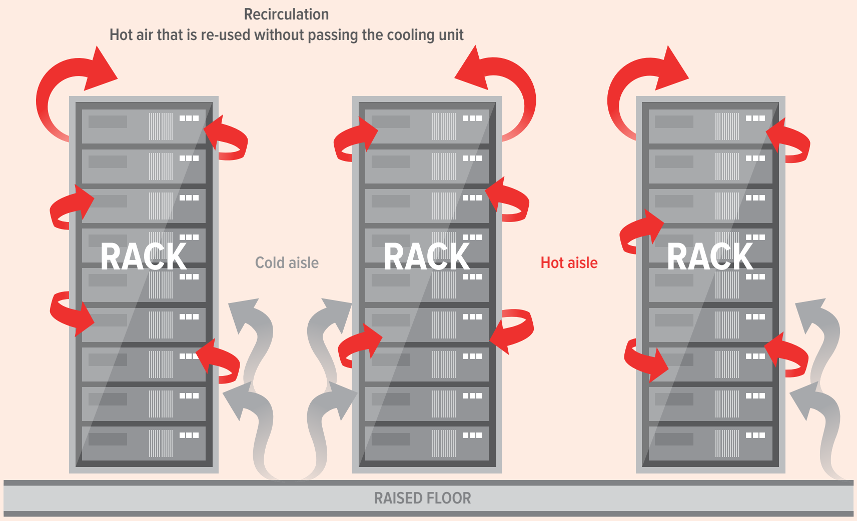

HOW TO DETECT AIRFLOW RECIRCULATION

The best way to detect if a data centre is experiencing airflow recirculation is to verify if there are any hot spots. To do that:

- Use an IR thermometer to identify a cold aisle with a wide range of temperatures or where temperatures are too hot
- Measure temperature at the top, middle and bottom of the rack by using sensors or a thermal heat gun
- Record the supply temperature from the cooling unit feeding that row (the best place to do this is under the raised floor in each cold aisle or at the base of the cabinet.)
- Determine if there is more than a 5 degree Fahrenheit difference between cooling system set point and the warmest temperature at the rack, which means there is energy waste.

Temperature sensors should be permanently deployed every third cabinet with one at the top, middle and bottom rack unit (RU). This allows a user to monitor the difference in temperature on a continuous basis. Ideally, the difference should be less than or equal to five degrees Fahrenheit. If the temperature variance is higher than that, then chances are there is airflow recirculation (mixing) or the perforated tiles are inefficient.

[20] Open Data Centre Measure of Efficiency – Paul Poetsma http://www.opendcme.org/opendcme-model/1-facility/14-recirculation

REDUCING AIRFLOW RECIRCULATION

Steps to reduce airflow recirculation include:

- Filling in all gaps inside the cabinet with blanking panels
- Plugging all gaps in between cabinets
- Adopting a hot and cold aisle layout
- Segregating supply and return air
- Using perforated tiles designed for heat dispersion
- Increasing Cubic Feet per Minute (CFM) to the cold aisle
- Replacing inefficient perforated tiles with tiles designed for higher heat transfer capability.

MONITORING AIRFLOW PRESSURE

> It is recommended that the differential air pressure trends be monitored over a period of time to establish a baseline.

After improving overall airflow pressure throughout the data hall, it is important to continue to tune the system. Monitoring the static pressure and airflow at different points in the data centre will help prevent significant pressure differentials, which could be a sign of bypass airflow or recirculation that ultimately lead to hot spots and equipment failure. It is recommended that the differential air pressure trends be monitored over a period of time to establish a baseline. The time allotted will depend on the data centre; if there is relatively few move, add and change work going on, a month should be sufficient; however, the ideal baseline would be a year-over-year comparison, which would then include seasonal variations that some data centres could experience due to extreme weather conditions.

Figure 13: Differential Sensor Placement in a Raised-Floor Environment

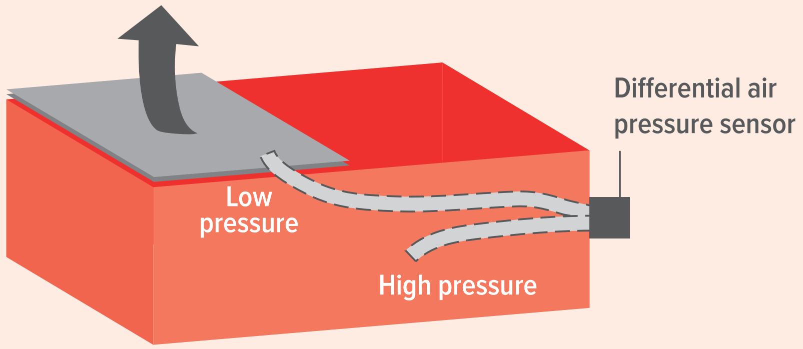

DIFFERENTIAL AIRFLOW PRESSURE SENSOR PLACEMENT

If the data centre is built on a raised floor or a slab, monitoring airflow pressure is still important. In order to collect accurate data, deploying sensors correctly is critical. The ideal placement for pressure sensors depends on the size and layout of the room, but here are some general best practices:

- Place sensors every 1,000 square feet
- Install a few feet above the floor
- Position sensor tubes in different pressure areas (e.g, above and below the raised floor and in the hot and cold aisle)
- When under the raised floor, install them no closer than 12 feet from a cooling unit

When collecting the information, use the lowest pressure reading as a guide because it is an indication of where the most air is required. In a nonraised floor environment, measure the static pressure difference between the hot and cold aisle.

BEST PRACTICE 2 OF 4

SUPPLY TEMPERATURE

The importance of achieving the right temperature is to help maximise the efficiency of the cooling system. It has been well documented that increasing the operating temperature will reduce the energy required to operate the cooling equipment thus providing substantial operational savings.

As stated earlier, the traditional method to cool server rooms in the past was to flood them with cool air and operate within the temperature range of 68 degrees Fahrenheit (20 degrees Celsius) to 71 degrees Fahrenheit (22 degrees Celsius). In 2004, ASHRAE increased the operating temperature range to 68 degrees Fahrenheit (20 degrees Celsius) to 77 degrees Fahrenheit (25 degrees Celsius) based on its findings and the advice from various IT equipment manufacturers. In 2008, ASHRAE added another change in the form of an addendum titled "Environmental Guidelines for Datacom Equipment" in which it expanded the recommended operating range to 64.4 degrees Fahrenheit (18 degrees Celsius) to 80.6 degrees Fahrenheit (27 degrees Celsius).

> Studies have shown that for every 1 degree Fahrenheit increase in server inlet temperature data centres can save 2 percent in energy costs.[21]

In 2011, ASHRAE published the third edition of the book, which contained two major additions. The first was guidance on server metrics that would assist data centre operators in creating a different operating envelope that matched their business values. Essentially, ASHRAE was providing the decision criteria and steps that data centre operators should take if they wanted to go beyond the recommended temperature envelope for additional energy savings. The second change was in the data centre classes. Previously, there were two classes that applied to IT equipment: classes 1 and 2. The new guidelines provided additional classes to accommodate different applications and priorities of IT equipment operation.

SUPPLY PRESSURE

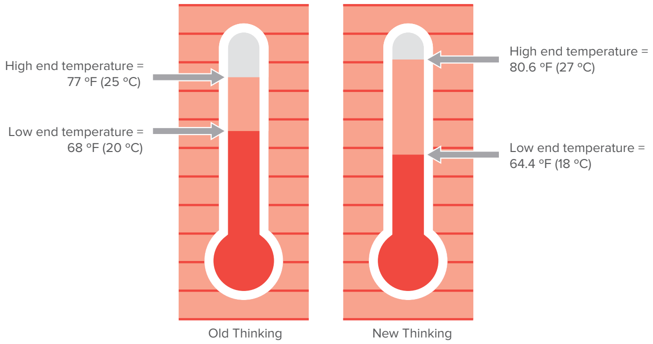

[21] Dr. Robert Sullivan – Anixter Interview 2015

Table 2: ASHRAE 2011 and 2008 Thermal Guideline Comparisons

2011 CLASSES	2008 CLASSES	APPLICATIONS	IT EQUIPMENT	ENVIRONMENTAL CONTROL
A1	1	Data centre	Enterprise servers, storage products	Tightly controlled
A2	2		Volume servers, storage products, personal computers, workstations	Some control
A3	NA		Volume servers, storage products, personal computers, workstations	Some control
A4	NA		Volume servers, storage products, personal computers, workstations	Some control
B	3	Office, home, transportable environment, etc	Personal computers, workstations, laptops and printers	Minimal control
C	4	Point-of-sale, industrial, factory, etc.	Point-of-sale equipment, ruggedised controllers, or computers and PDAs	No control

Based upon self-reported DCD intelligence census data, the global average of inlet air temperature is 72.5 degrees Fahrenheit (22.5 degrees Celsius). The proportion below ASHRAE's minimum recommended level of 64.4 degrees Fahrenheit (18 degrees Celsius) is higher (9.8 percent) than the proportion above the maximum level (2 percent), which means there still is capacity for data centres to reduce energy consumption and costs by raising air inlet temperatures.

THE DESIRE TO KEEP THINGS THE SAME

If raising the inlet temperature in the data centre has been proven to reduce energy thereby significantly reducing operational expenses, why aren't more data centres doing it? Some of the reasons that have been stated by various data centre operators are[23]:

- The current data centre is operational without any issues
- Legacy IT systems might not be able to cope with temperature changes
- The working environment will be hotter, and that isn't worth the savings
- Feeling it is difficult to change thermal management strategies in an existing data centre
- If there is a cooling system failure, IT equipment will shut down faster.

Table 3: Inlet Air Temperature 2014[22]

TEMPERATURE RANGE	PERCENT SPECIFIED
64.4 °F (18 °C) or cooler	9.8%
64.4 °F (18 °C) > 68 °F (20 °C)	15.8%
68 °F (20 °C) > 69.8 °F (21 °C)	20.6%
69.8 °F (21 °C) > 71.6 °F (22 °C)	24.1%
71.6 °F (22 °C) > 75.2 °F (24 °C)	12.4%
75.2 °F (24° C) > 77 °F (25 °C)	5.6%
77 °F (25 °C) > 78.8 °F (26 °C)	2.0%
80.6 °F (27 °C) or warmer	2.0%
Unable to specify	16.6%

[22] DCD Intelligence Census Survey – 2014

[23] Why Aren't Data Centres Hotter? http://www.datacentredynamics.com/critical-environment/why-arent-data-centres-hotter/94222.fullarticle - June 2015

CONSIDERATIONS WHEN INCREASING OPERATIONAL TEMPERATURE

Server Reliability
There is a perception in the industry that higher operating temperatures in the data hall can lead to a higher number of failure rates with the IT equipment. These concerns do hold some weight as it has been shown that there is a greater likelihood of failure when equipment is exposed to high temperatures caused by data centre hot spots. However, what isn't as widely understood is the long-term effect of temperature change on servers.

Figure 14: Relative Server Failure Rate with Temperature[25]

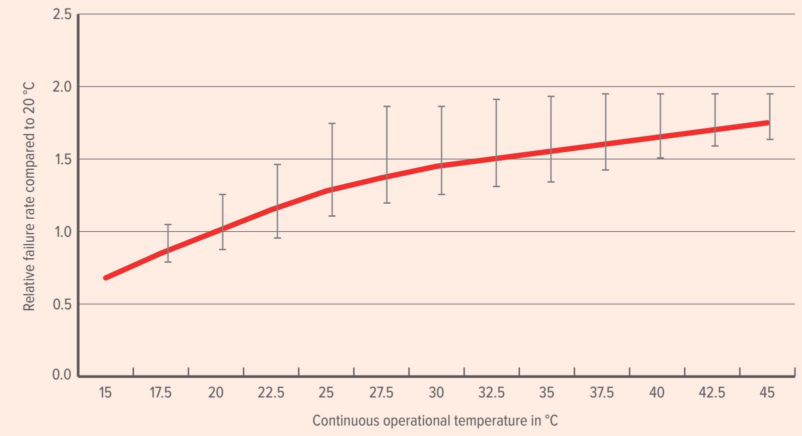

In 2011, ASHRAE released a paper titled 2011 Thermal Guidelines for Data Processing Environments – Expanded Data Centre Classes and Usage Guidance. One of the studies presented in the paper focused on relative server failure rates evaluated against temperature based on data from a variety of different hardware manufacturers. What was found was that at 68 degrees Fahrenheit (20 degrees Celsius), which was the baseline measurement one could expect roughly 10 servers to fail in a 1,000 server environment per year. If the inlet temperature was slowly increased to 80.6 degrees Fahrenheit (27 degrees Celsius), which is at the high end of ASHRAE's recommended temperature range, there would be an average increase of 2.5 server failures per year, for a total of 12.5 server failures for that same year. In other words, 10 failed servers out of 1,000 means 1 percent of the total servers failed; if 12.5 servers fail out of 1,000, that means 1.25 percent of the servers failed, or just a 0.25 percent difference. This data shows there isn't a significant increase in server failures over the long term when increasing inlet temperatures within ASHRAE's recommended guidelines.[24]

There isn't a significant increase in server failures over the long term when increasing inlet temperatures within ASHRAE's recommended guidelines.

[24] The Green Grid: Data Centre Efficiency and IT Equipment Reliability at Wider Operating Temperature and Humidity Ranges
[25] ASHRAE, 2011 Thermal Guidelines for Data Processing Environments – Appendix C, reformatted

Cooling System Ride-Through Time

Another commonly held misconception by data centre operators is that if they maintain a lower operating temperature in their server rooms they would increase the ride-through time of the cooling system. Ride-through time is the speed a data centre would heat up to reach a critical temperature causing IT equipment to shut down and/or fail during a cooling system failure. There is some validity to this depending on the thermal strategy and cabinet densities residing in the current data centre. Data centres that have deployed hot or cold aisle containment have a smaller contained air volume, so the temperature to the racks would rise faster.[26] Additionally, higher rack power densities will shorten the time for the environment to get to a critical temperature in the event of an outage. In data centres that operate in an open environment with no containment and lower cabinet densities, the ride-through time will be longer, but only by a few seconds. In environments where server densities are 6 kW or greater, the Uptime Institute recommends continuous cooling where the heat removal system is uninterrupted during the transition from utility power to generator power.[27]

Comparing server inlet temperatures with an operating temperature of 55 degrees Fahrenheit (13 degrees Celsius) versus an ideal operating temperature of 77 degrees Fahrenheit (25 degrees Celsius) would only buy seconds if there wasn't a UPS backup on the cooling system.

All things being equal, comparing server inlet temperatures with an operating temperature of 55 degrees Fahrenheit (13 degrees Celsius) versus an ideal operating temperature of 77 degrees Fahrenheit (25 degrees Celsius) would only buy seconds if there wasn't a UPS backup on the cooling system[28]. The amount of operating expense difference that would be incurred by running at the cooler temperature generally outweighs the 12 seconds that would be gained in that scenario.

Working Environment

There should be some consideration given to the working environment when deciding whether to raise the operating temperature in the data centre. The Occupational Safety and Health Administration (OSHA) and the International Organisation for Standardisation (ISO) both have guidelines for working in high-temperature environments.

OSHA calculates temperature exposure limits on the wet bulb globe temperature (WBGT), which also takes humidity into account. Because much of the hot aisle work is classified by OSHA as light work (adding removing cabling to servers), at 86 degrees Fahrenheit (30 degrees Celsius) continuous work is allowed. However, as temperatures start to climb past that, OSHA regulations could require employees working in the hot aisle to work less time and have more rest. For more detailed information on this, please refer to the OSHA technical manual Section III, Chapter 4.

In a traditional uncontained work environment where the IT equipment inlet air temperature is between 56-81 degrees Fahrenheit (13-21 degrees Celsius) where there is a rough assumed 40 percent bypass airflow (leakage) and 20 percent recirculation, the rest of the server room (outside the cold aisle) should be roughly 75 degrees Fahrenheit (24 degrees Celsius). At these temperatures the OSHA high-temperature guidelines wouldn't apply.

[26] Focused Cooling Using Cold aisle Containment – Emerson Network Power

[27] Uptime Institute – Implementing Data Centre Cooling Best Practices https://journal.uptimeinstitute.com/implementing-data-centre-cooling-best-practices/

[28] Facility Cooling Failure: How Much Time do You Have? David L. Moss - http://partnerdirect.dell.com/sites/channel/en-us/documents/facility_cooling_failure_white_paper.pdf

SIX STEPS TO ADJUST SUPPLY TEMPERATURE

There are a few steps that need to take place in order to safely increase the supply temperature in the data centre to the desired level.

1	DEPLOY TEMPERATURE SENSORS	There needs to be a baseline reading collected, and in order to do that, temperature sensors should be deployed across the data hall. Per ASHRAE guidelines, temperature and humidity sensors should be deployed every third rack, at the supply side with one sensor aligned with the top U, one in the middle U and one at the bottom.
2	MONITOR EXISTING TEMPERATURE	When collecting data from the sensors, it is important to establish a baseline of information and to understand the thermal dynamic of the current space. If there are hot spots currently in the data centre, it is important to get those issues resolved before turning up the set point. Temperature data should be monitored for at least a few days to account for swings in IT equipment utilisation throughout the day.
3	CORRECT AIRFLOW MANAGEMENT PROBLEMS	Bypass airflow and recirculation problems must be corrected and the proper number and distribution of perforated tiles need to be deployed before changing the temperature set points.
4	MOVE AIR HANDLER CONTROL POINT	With many legacy air handlers, the control point was located on the return air side. Moving the control to the supply side will give the data centre manager more control over the temperature that is being supplied to the IT equipment. There are kits available from various cooling system manufacturers to accomplish this.
5	ADJUST TEMPERATURE SET POINT	It is important that raising the temperature set point is done gradually, 1 degree Fahrenheit per week, until the desired inlet temperature has been achieved. It is important to watch for additional hot spots that could form throughout this process. The input air temperature should not exceed 80.6 degrees Fahrenheit (27 degrees Celsius). If there are areas where the temperature has been exceeded then the airflow volume from the air handlers needs to be increased or additional perforated floor tiles need to be placed in the cold aisle, always maintaining the proper number of perforated tiles for the total airflow volume in the computer room.
6	ADJUST CHILLED WATER TEMPERATURE	Raising the set point of the air handler generally means that chilled water temperature can be increased as well. Many data centres have chilled water set points between 42-45 degrees Fahrenheit (6-7 degrees Celsius) that have been raised in some facilities to 50 degrees Fahrenheit (10 degrees Celsius) or up to 18 degrees Celsius for substantial energy savings[29]. It is best to consult with the manufacturer of the chiller and air handler for guidance on the ideal temperature balance.

[29] Implementing Data Centre Cooling Best Practices –Uptime Institute https://journal.uptimeinstitute.com/implementing-data-centre-cooling-best-practices/

WHAT IS THE RIGHT TEMPERATURE FOR MY DATA CENTRE?

The short answer is: it depends. There are several considerations that go into understanding what the right temperature is for any particular data centre. Server manufacturers have agreed that inlet temperatures that fall within the range of 64-80.6 degrees Fahrenheit (18-27 degrees Celsius) have little effect on the reliability or performance, pending the temperature is not a sudden, but a gradual increase. ASHRAE has provided recommended and allowable guidelines, and classes within the data centre environment for additional guidance. It all comes down to what makes the particular data centre manager feel comfortable. A general rule of thumb is that temperature in the cold aisle should be even from the bottom of the rack to the top of the rack, and as warm as possible not to exceed ASHRAE recommended guidelines.

A general rule of thumb is that temperature in the cold aisle should be even from the bottom of the rack to the top of the rack, and as warm as possible not to exceed ASHRAE recommended guidelines.

A good start is to set a guideline, adjust gradually, monitor the change and report back the results. From there, if the equipment is still within the recommended temperature range, there have been no major complications and the employees are comfortable with the working conditions, the temperature can continue to be gradually raised. Remember, for every 1 degree Fahrenheit increase, there is roughly a 2 percent savings in energy cost.

BEST PRACTICE 3 OF 4

AIRFLOW SEGREGATION

3

The main purpose of any airflow segregation strategy is to isolate the cool supply air from the hot exhaust air, which prevents airflow recirculation. The less airflow recirculation there is in the room, the better temperature control the data centre operator will have and the less likelihood there will be any data centre hot spots. Airflow segregation, along with following the other best practices of conditional environmental control outlined in this paper allows for the following benefits:

- Prevents airflow recirculation, which reduces the risk of hot spots
- Allows for higher temperature return air to the HVAC equipment which increases efficiency
- Provides more consistent equipment rack inlet temperatures with the use of top rack units across the computer room

It is important to note that airflow segregation doesn't provide any efficiency gains on its own. What it does is allow for is greater energy savings and control over supply and return temperatures. In order for those efficiencies to be realised, data centre managers still need to increase rack inlet temperatures and reduce the volume of air being delivered to the IT load by reducing fan speeds or turning off cooling units.

It is important to note that airflow segregation doesn't provide any efficiency gains on its own. What it does is allow for greater energy savings and control over supply and return temperatures.

AIRFLOW SEGREGATION

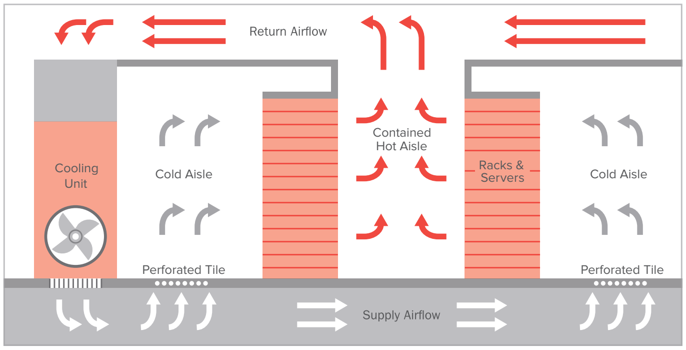

DIFFERENT METHODS OF AIRFLOW SEGREGATION

There are three main types of airflow segregation strategies available today:

- Hot aisle containment
- Cold aisle containment
- Vertical exhaust ducts

Each of these methods has multiple variations of how it can be constructed and deployed. For instance, hot aisle containment can be deployed with or without a roof, with one row or two rows of equipment cabinets, or with curtains or sliding doors. The following sections explain the differences between the three methods and considerations when selecting one versus the other.

Hot Aisle Containment

In a hot aisle containment deployment, the hot aisle is enclosed to capture the IT equipment airflow discharge with the rest of the data hall acting as the cold aisle. Generally in a hot aisle deployment, two rows of racks are set up with the "rear" of the racks facing each other so the IT equipment discharge flows into a hot aisle. From there, the hot air travels either through grates into a drop ceiling plenum, can be fully contained with a roof, or if the room has a high ceiling (greater than or equal to 24 feet), vertical walls can be constructed so the hot air is released back into the room well above the cabinets.

Cold Aisle Containment

In a cold aisle containment deployment, the cold aisle is enclosed to capture the supply air that comes from the active cooling equipment with the rest of the data hall acting as the hot aisle. In this environment, two rows of racks are set up with the "front" of the racks facing each other so the supply air from the cold aisle can flow into the IT equipment in each row. Generally, cold aisle containment has a cap or roof structure installed at the top of the cabinets.

Vertical Exhaust Ducts

Vertical exhaust ducts (VED), or chimney systems as they are sometimes referred to, is another hot aisle containment approach. The difference, however, is that each cabinet has an exhaust duct (or chimney) mounted directly to the top rear to allow for the heat exhausted from rack-mounted equipment to flow directly into an overhead ceiling plenum or if the ceiling is high enough into the data hall and pulled back into the cooling unit(s). With this system the entire room becomes a cold aisle.

Table 4: Containment Physical Assessment Checklist

FACILITY PHYSICAL ASSESSMENT	YES	NO
Is the data centre arranged in a hot/cold aisle alignment?	□	□
Are the cabinets a uniform height/width?	□	□
Is there a drop ceiling installed?	□	□
If no, is there clearance to install a drop ceiling?	□	□
Is the data centre installed on a raised floor?	□	□
Is the data centre installed on a slab?	□	□
Is there overhead cabling installed?	□	□
Is there an overhead bus way system installed?	□	□
Are there any columns inside a row of racks?	□	□
Is perimeter cooling system installed?	□	□
Is a close-coupled coupled cooling system installed?	□	□

HOW TO SELECT THE RIGHT FORM OF CONTAINMENT

Research has shown that from an efficiency standpoint hot, cold or chimney containment does not hold one advantage over the other.[30] So the question becomes, how do you know what is best for your data centre? The first step is to assess the current facility. Table four on page 35 contains a checklist that will provide the key information needed to better determine the right form of airflow segregation in an existing data centre.

> Research has shown through research that from an efficiency standpoint that hot, cold or chimney containment does not hold one advantage over the other.

Once the facility has been assessed, then the data centre operator can begin to evaluate one form of containment versus the other.

WHEN TO CONSIDER COLD AISLE CONTAINMENT

- ✓ It’s generally the simplest solution to deploy in a retrofit.
- ✓ There is a raised floor already installed.
- ✓ The data centre is installed in a hot and cold aisle configuration.
- ✓ There are few stand-alone cabinets, and most cabinets are in uniformed aisles and have similar levels of power dissipation.
- ✓ Higher overall data centre working conditions are acceptable.

WHEN TO CONSIDER HOT AISLE CONTAINMENT

- ✓ There is a drop ceiling with a return air plenum installed
- ✓ There are many stand-alone IT cabinets.
- ✓ The data centre is installed in a hot and cold aisle configuration.
- ✓ Cooler overall working conditions are preferred.
- ✓ With hot aisle containment, a single airflow control can be implemented for a normal size computer room. For large rooms, multiple control zones might have to be implemented.

WHEN TO CONSIDER VERTICAL EXHAUST DUCTS

- ✓ There is a drop ceiling with a return air plenum installed.
- ✓ Building columns interfere and make aisle containment impractical.
- ✓ There are limited overhead obstructions.
- ✓ High-density cabinets are scattered around the data hall.
- ✓ There is no raised floor installed.

[30] Datacentre 2020: hot aisle and cold aisle containment reveal no significant differences

MEETING FIRE CODE COMPLIANCE

One of the challenges with aisle containment is understanding the impact of the fire and safety requirements on retrofits and new construction data centre deployments. Many data centres are implementing containment solutions due to the energy-efficiency benefits, but they are unaware of the potential risks as it relates to the fire and safety standards.

The National Fire Protection Association (NFPA), while being focused mainly in North America, can provide guidance globally and has two standards that address aisle containment fire protection in data centres:

- NFPA Standard 75 (2013), Standard for Fire Protection of Information Technology Equipment
- NFPA Standard 76 (2012), Standard for the Fire Protection of Telecommunications Facilities

No longer are fusible links (heat activated/ temperature sensitive) compliant; the roof of the containment system needs to be removed via smoke or fire system activation.

It is highly encouraged that data centre managers purchase both copies of the standard and review them in their entirety. Here are some highlights from the standards[31]:

- Elements of the containment system must be constructed with materials that have a flame spread index less than 50 and smoke development less than 450 in accordance with one or more of the following: ASTM E 84 and UL723.
- Smoke detectors need to be rated for the intended temperatures of the hot aisles if they are located there.
- Aisle containment systems are no longer considered plenums, so plenum-rated construction materials are unnecessary.
- In a retrofit environment, if automatic sprinklers are present and the aisle containment creates an obstruction to the sprinklers, they need to be modified to comply with NFPA 13.
- No longer are fusible links (heat activated/temperature sensitive) compliant; the roof of the containment system needs to be removed via smoke or fire system activation.
- Panels can no longer drop into the aisle, which can potentially obstruct egress.

Having a better understanding of these requirements before designing an aisle containment solution can help reduce costs that would be incurred through additional labour should the original design not pass a local fire code.

SIDE BREATHING EQUIPMENT

Equipment that does not breathe front to rear or rear to front can have a detrimental impact on any containment system or the equipment itself. As data centre managers start to raise the set points and contain their aisles, ensuring this type of equipment gets adequate cooling becomes even more critical. For these types of situations here are some recommendations:

- Deploy a cabinet with adjustable side intake ducts that support mixed airflow equipment in the same cabinet.
- In existing cabinets that cannot support side breathing equipment, deploy rack mountable intake devices matched to the type of airflow required by the device. These are available in active (fan assisted) or passive models.

[31] Complying With NFPA's Aisle Containment Requirements – ASHRAE Journal, September 2015

BEST PRACTICE 4 OF 4

AIRFLOW CONTROL

Airflow control can be defined as the process of directing airflow from the cooling unit to the IT equipment and back. Air flows much like water and will generally take the path of least resistance. That is the reason it is so critical to help guide the air to where it needs to go. Making sure the supply pressure is balanced, the temperature is at an ideal level and the supply and return air are segregated, which will improve overall efficiency and make for a more reliable environment for the IT equipment. Control over the airflow will ultimately provide data centre managers the stability they need for the entire thermal management system. This area of evaluation compared to the other three best practices is probably the simplest and most overlooked.

New studies being conducted by several universities funded by the National Science Foundation (NSF) are showing significant effect on airflow and temperature control by the geometric design of airflow panels. For the past 50 years, the method of cooling data centres has been to determine the full BTU (kW) load generated from the electrical usage and provide enough CFM of cold air to offset the heat generation.

It's not the volume of air that comes out through the floor tiles that matters, what is most important is deploying floor tiles that can create a turbulent airflow which has a greater impact on heat transfer.

The current way to classify perforated tiles is based on the flow area offered. Typically tiles with a 25 percent open area are used for low-density applications that require flow rates less than 1,000 CFM/tile (1,667 cubic meters/hour), and tiles with up to 60 percent open area are used for higher rack densities that require flow rates greater than 1,200 CFM/tile (2000 cubic meters /hour).

What the game-changing discovery has been is that the vent design of the perforated floor tile plays a critical role in the distribution of cold air to the rack. Depending on the vent design, the distribution of cold air can vary significantly across tiles that have an identical open area.[32] In other words, it's not the volume of air that comes out through the floor tiles that matters. What is most important is deploying floor tiles that can create a turbulent airflow which has a greater impact on heat transfer.

[32] Complying With NFPA's Aisle Containment Requirements – ASHRAE Journal, September 2015

Figure 15: Different Types of Air Movement

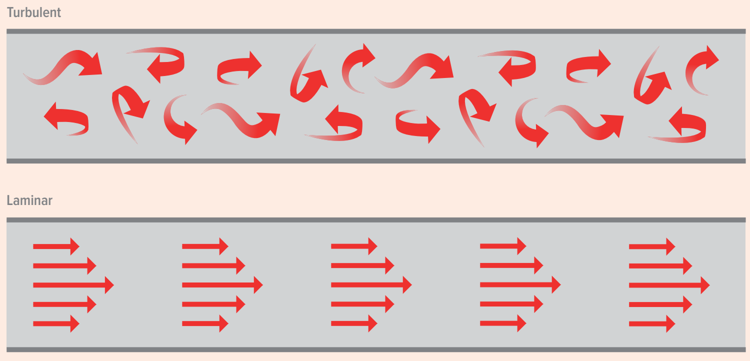

AIRFLOW CONTROL

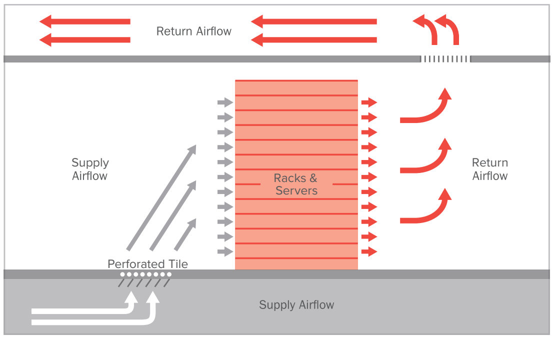

TYPES OF AIRFLOW DISPERSION IN THE COLD AISLE

- **Puddle effect**
  Airflow distribution has poor velocity or is short cycling at the front edge of the panel. This is normally seen with 25 percent airflow panels but may also be seen in some high-density designs due to geometrical design of panel construction.
- **Jet stream effect**
  Airflow distribution is moving at too fast of a velocity. This is seen quite often in high-density (50 percent or more open area) panels that are streaming the air and lacking turbulent flow characteristics.
- **Heat transference effect**
  Airflow is directed toward a rack but with a designed turbulence and speed that allows cooler air to be pulled in effectively by server fans. This type of airflow dispersion is created by perforated tiles designs that have a lower anterior direction fin protruding below the bottom surface of panel.

SELECTING THE RIGHT PREFORATED FLOOR TILE

- Constructed in a way that ensures positive flow out of every part of the tile (no short cycling)
- Diverts air to the servers so it can penetrate the boundary air of heat on the front side of the cabinet or rack
- Should be able to stratify the air to the top of the cold aisle
- Contains fins underneath, which help to create a heat transfer flow
- Correct load rating to prevent damage from cabinet installations

ELECTRONICALLY COMMUNICATED (EC) AND VARIABLE SPEED DRIVES (VSDs)

After the mechanical cooling system, fans are the next largest energy consumer on computer room air conditioning units (CRACs). In traditional cooling systems, the fans are often operating at a fixed capacity — 100 percent. Moving to a variable speed fan technology can save fan energy consumption by as much as 76 percent by enabling cooling systems to throttle fan speed up or down to meet the changing IT heat load.

Energy consumption by the fans is related to the cube of the speed of the fan. For example:

- A 10 percent reduction in fan speed would result in a 27 percent energy savings
- A 20 percent reduction in fan speed would result in a 49 percent energy savings[33]

[33] Smart Ideas for your Business Variable Speed Drives – Commonwealth Edison Company 2012

Figure 16: Comparing Heat Transfer, Puddle and Jet Stream Airflows in the Data Centre Cold aisle

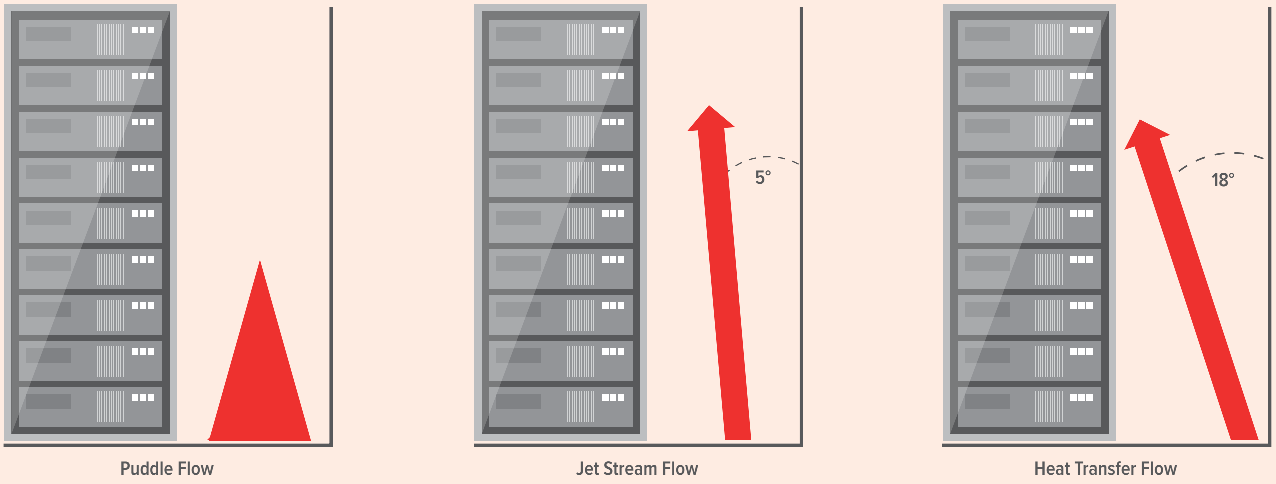

COMPARING THE DIFFERENCE BETWEEN VARIABLE FREQUENCY DRIVES AND EC FANS

In order to evaluate what would be best for a particular environment, it is important to know the difference between a VFD and an EC fan. Both have particular benefits, but it comes down to the amount of capital that can be invested up front and the type of cooling system currently in place. Most manufacturers have retrofit kits that can be purchased; however, it is recommended that they be installed by the manufacturer or a certified mechanical contractor. Below are the major differences between the two solutions:

VFDs

- Drives added to current centrifugal fans in CRAC system
- Achieves speed control by varying the frequency of the electrical current
- Enables fan speed to be adjusted based on IT heat load demands

EC Fans

- Replaces the fans and motor assemblies in an existing CRAC
- Provides a different, more efficient design when compared with traditional centrifugal fans
- Achieves speed control by varying DC voltage delivered to the fan
- Enables fan speed to be adjusted based on IT heat load demands

Table 5: VFD and EC Fans Compaison[34]

	VARIABLE FREQUENCY DRIVES	ELECTRONICALLY COMMUNICATED FANS
Energy Savings	Save energy when operated below full speed	Save energy at any speed due to more efficient design
Cooling Unit Type	Better suited for larger systems with ducted up flow cooling units	Better suited for down-flow units
Maintenance	Similar to traditional fans	Reduces maintenance due to no fan belts and less dust creation
Installation	Can be retrofitted on existing cooling units	Can be retrofitted on existing cooling units
ROI	Faster payback due to lower upfront cost	Lowest total cost of ownership, but higher upfront cost

[34] Choosing Between VSDs and EC Fans – Emerson Network Power

THE FUTURE OF DATA CENTRE COOLING

The cooling systems of today might not be what the cooling systems of tomorrow look like. As data centres continue to grow and evolve, there will be new technologies that come to the forefront. As resources such as water become more expensive, cooling system manufacturers are continuing to come up with new, more efficient ways to cool the data centre. Here is a look at some future technologies that might be used more in the future.

LIQUID COOLING

There are different approaches to liquid cooling. The two main approaches are to either bring a liquid into the server itself via an enclosed system that features pipes and plates or by immersing the entire server chassis into a fluid.

There have been predictions of a total energy savings (cooling plus IT) for a 2.5 MW data centre in New York, Boston, Philadelphia or Toronto in the 30 percent range for annual power savings based on $0.10/kWh, of more than 500,000 USD.

One challenge with liquid cooling is the high upfront cost associated with it. If the cost doesn't come down, then many data centre managers are likely to avoid it. Another challenge is simply the fear of bringing a liquid closer to the IT equipment itself. This has been something generally avoided so there are some negative perceptions with doing such.

HEAT WHEEL COOLING

Heat wheels have been around for years, but recently heat wheels are starting to find their way into some data centres. A heat wheel is generally made out of aluminium honeycomb. Hot air from the servers is pulled into the heat wheel by fans. The wheel absorbs the heat from the air as it passes through. The cooled air is then sent back into the data hall. In a separate plenum, outside air is pulled into the cooling unit by fans, passes through the wheel absorbing the heat and discharges it back outside. The wheel rotates between plenums to effectively transfer the heat.

Heat wheel cooling systems are efficient, are an indirect airside cooling system and have been installed in a variety of different geographies around the world. There have been predictions of a total energy savings (cooling plus IT) for a 2.5 MW data centre in New York, Boston, Philadelphia or Toronto in the 30 percent range for annual power savings based on $0.10/kWh, of more than 500,000 USD.[35]

Some of the drawbacks to heat wheel cooling are the amount of space it requires cost; as conventional cooling is still required if outside temperatures are too high.

FREE COOLING ADOPTION

Gartner defines free cooling as any technique used to reduce the energy consumed by the cooling systems or the time that the cooling units run by using the outside temperature of air or water to cool the data centre or other facilities. Free cooling solutions are being adopted all over the world today. IHS forecasts strong growth globally for the free cooling market through 2018.

There are two main types of free cooling technologies: air-side and water-side economisers. As data centres start to run warmer, it allows for more hours that a free cooling system can be used, which reduces the overall operational expense of the cooling system and making free cooling systems more attractive to data centre managers.

[35] Robert E. McFarlane TechTarget – 2010

CONCLUSION

This report covered a wide spectrum of variables that impact the data centre thermal environment. Increasing efficiency, reducing total operating costs and avoiding costly outages are top of mind for all data centre operators and managers, and cooling is an area that directly impacts all these objectives. Using the foundation set forth in this paper provides a better understanding of today's challenges, methods, standards and more. Additionally, having knowledge of the four best practices for optimising a data centre's thermal environment provides a clear roadmap for a more efficient and reliable facility.

This report covered a lot of ground with this complex topic, trying to bring some important – though not always obvious – flash points to the forefront. On the surface, keeping hot-running IT equipment cool seems pretty straightforward that data centre owners and facilities managers can solve by simply purchasing additional cooling equipment and pumping in large volumes of conditioned air. However, what has been discovered in many data centres, is that not following the best practices of airflow management has a detrimental impact on the entire cooling system. Additionally, most existing data centres have a surplus of cooling capacity that can be reclaimed by reducing bypass airflow and preventing recirculation.

The key takeaway is that thermal management needs to be addressed as an integrated, holistic and ongoing effort, not in a piecemeal, disparate and short-term way. The good news is that once these four key areas are addressed properly in the data centre, it's possible to achieve an energy-efficient, cost-saving, high-performing thermal environment.

GLOBAL TECHNOLOGY BRIEFING

DCIM ENABLEMENT BEST PRACTICES

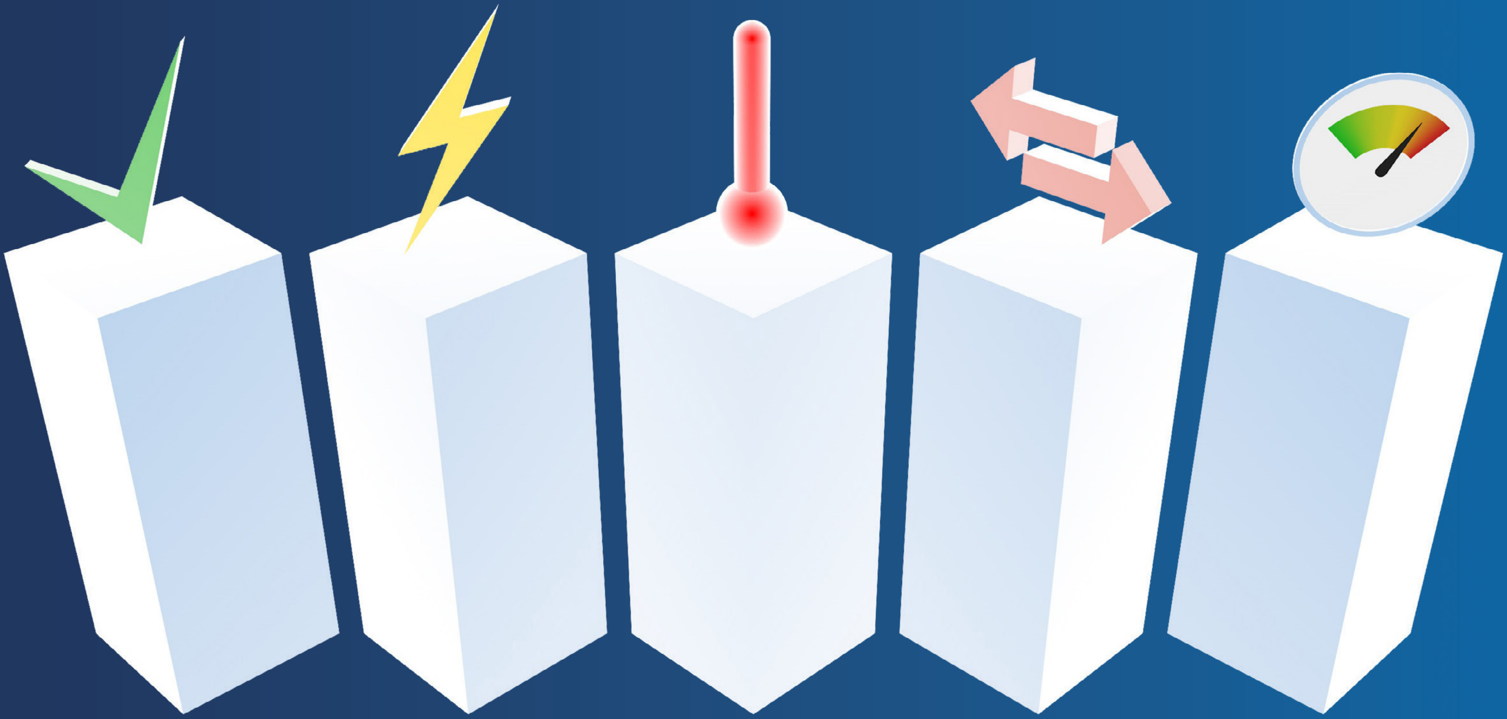

EXECUTIVE SUMMARY

In the IT universe, “get more from less” is a common mantra. Usually, this mandate comes from the executive suite and applies to the IT budget, staff, and other resources. However, when it comes to a data centre’s infrastructure — power, cooling and space — how can more efficiency and lower costs be achieved while maintaining availability? Is there a way to effectively manage a data centre’s infrastructure in a way that makes sense to the IT team and fosters collaboration with facilities to achieve the goals of the business?

Data centre infrastructure management (DCIM) can hold the answers to those questions and more with one caveat — the data centre team must understand which pieces of DCIM will solve their business challenges and which can be readily integrated with existing technologies. It is critical that organisations looking at such solutions need to be ready for DCIM by defining a scope that will solve the most pressing challenges while making sure that it is focused enough to be achievable.

DCIM is still a relatively new and evolving area and is subject to industrywide confusion. With barriers to adoption and common industry challenges, it’s important to create a plan for the selection and preparation of a DCIM solution. This is why Anixter created the five senses of DCIM — a thorough treatment of key areas that data centre managers should know as they seek to implement a solution.

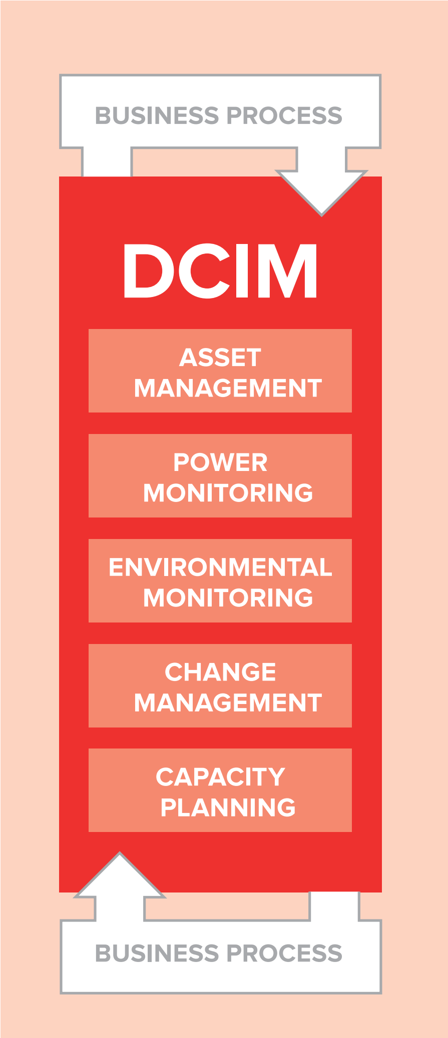

INTRODUCTION

DEFINING DCIM

Although some stand-alone infrastructure management products were being used in larger data centres more than a decade ago, DCIM was first widely used in 2007. At the time, existing data centre infrastructure management products focused mainly on power monitoring and environmental conditions. Products that addressed asset management and mapped IT processes for managing move, add, and change work were less common.

During this time, the facilities management team almost exclusively monitored and managed data centre power and cooling systems. By 2010, the continued growth of data centres made understanding the performance and capacity of these functions critical for IT teams. This was especially true with the largest enterprises, which were experiencing subpar efficiency in power, cooling and physical space usage, resulting in thousands or even hundreds of thousands of dollars in energy and resource waste and gaps.

Throughout its short history, DCIM has often been defined by individual manufacturers and suppliers, often with a bias that tilted its functionality in a certain direction. However, most industry experts agree on a broad definition of DCIM as a solution that helps data centre managers track and analyse information about their facilities' operational status, assets and resource use, such as space, power and cooling.

Figure 1: The Flow of Information Management

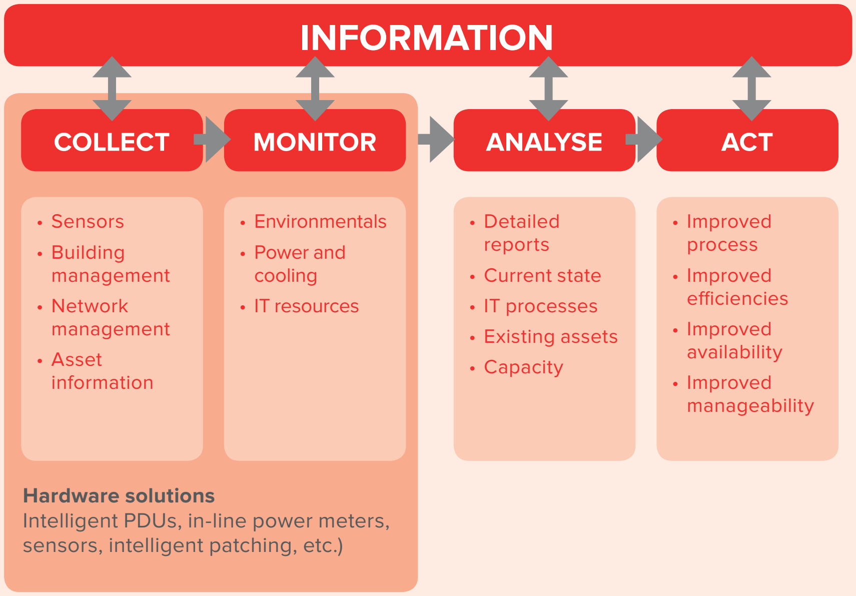

Gartner defines DCIM as "tools that monitor, measure, manage, and/or control data centre use and energy consumption of all IT-related equipment (such as servers, storage, and network switches), and facilities infrastructure components (such as power distribution units [PDUs] and computer room air conditioners [CRACs])."

A more aspirational and forward-thinking view of DCIM is one that positions it as a way to gain a complete picture of the current state of the data centre, providing actionable and intelligent insight that drives energy efficiency and strategically plans for future capacities, including space, power and cooling resources.

A comprehensive and well-integrated DCIM solution will collect and monitor a data centre's assets, resource use and operational status, capturing thousands of data points. This data can then be analysed and acted on by staff to meet business and service-oriented goals as well as maximise performance.

DCIM systems can help data centre managers:

- Execute technical and business goals and changes
- Reduce waste and unnecessary over-provisioning
- Plan for new data centre capacity through forecasting
- Decrease downtime risks
- Increase energy efficiency.

DCIM DRIVERS

MODERN DATA CENTRE COMPLEXITY

With thousands of IT assets under their supervision, data centre managers are constantly under pressure to deliver results. To deliver those results, they need tools that can transform the inherent complexity of running a data centre into information and insight they can act on. The increased complexity of the data centre architecture, including higher densities and virtualisation, has exceeded the capabilities of managing through the use of traditional means, such as simple spreadsheets.

Today's data centre managers need to see, understand and optimise the complex interrelationships that drive the modern data centre, arguably one of the most complex types of facilities in the world. Existing management systems and techniques often have limited capabilities, and data centre managers need to understand what is happening in the data centre. IT asset management systems fall short of showing the interrelationship of data centre assets. In addition, ad hoc monitoring and reporting is growing increasingly unwieldy as data centres grow.

Those charged with running and managing complex data centres need a holistic and transparent view of their entire IT infrastructure that outputs meaningful and actionable data instantly in real time and with accuracy.

In short, those charged with running and managing complex data centres need a holistic and transparent view of their entire IT infrastructure that outputs meaningful and actionable data instantly in real time and with accuracy.

INTERNET OF THINGS

In addition, the trend toward the Internet of Things (IoT) has added to data centre management complexity with the sheer amount of data that needs to be analysed and used. In fact, Gartner estimates that by 2020, the IoT will grow to 26 billion units, creating enormous challenges and even greater complexity for data centres.[1] These IoT deployments will generate a tremendous amount of data that will require processing and analysing in real time.

Today's data centre managers need to be more forward-looking with capacity management tools at their disposal, such as DCIM, to be able to actively meet the business priorities associated with IoT.

[1] Gartner – Gartner Says the Internet of Things Will Transform the Data Centre, 2014

CAPACITY PLANNING TOP-OF-MIND

On the surface, it may seem pretty straightforward that the top driver for DCIM is energy efficiency. After all, there are financial pressures, particularly when it comes to the need to decrease energy costs. The drive toward higher efficiency is also being pushed through legislation and industry standards, including the EPA ENERGY STAR program for data centres and the European Union Code of Conduct.

According to an Uptime Institute survey,[2] capacity planning is the main reason many organisations are adopting a DCIM solution. However, it is important to point out that capacity planning is really the goal for DCIM, not a starting point. DCIM generally starts with better monitoring and management of power, cooling and IT assets.

To fully understand capacity, data centre managers need to know how much power and cooling they are consuming out of the total available, as well as asset utilisation. As a byproduct of better capacity planning, a data centre manager will better understand how to maximise physical space (asset utilisation), power and cooling systems (energy efficiency and reduce operational expenses [OPEX]) and reduce risk (minimise outages).

Figure 2: DCIM Adoption Drivers[3]

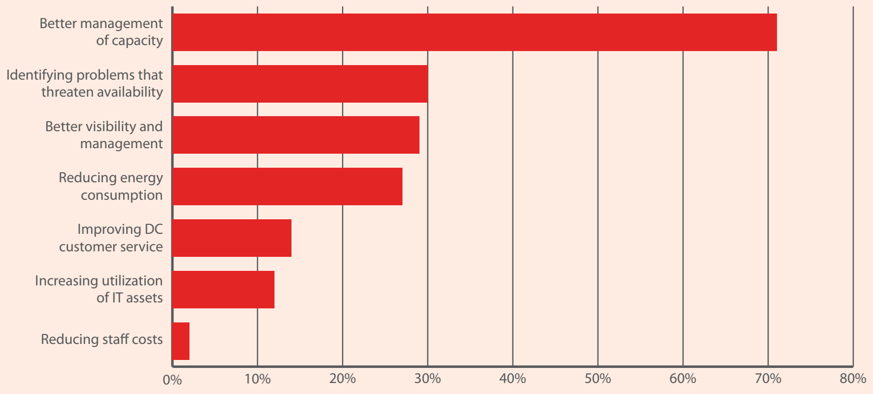

[2] Uptime Institute Global Datacentre Annual Survey, 2013

[3] DCD Intelligence, 2015

LACK OF ADOPTION

Data centre complexity, Internet of Things and the need for capacity planning are at the top of the list to make DCIM deployment a priority. Yet, despite all the evidence that makes the business case for DCIM seem pretty strong, there are still challenges.

DCIM MISUNDERSTANDING

Throughout the industry, there is a prevalent lack of understanding on what exactly constitutes a DCIM solution as well as its functionality. Adding to the confusion are differing views on DCIM models. Multiple DCIM models have been put forth by analyst firms such as Gartner, Forrester and the 451 Group. Even though they are similar in many respects, there are subtle differences between the various views of DCIM.

Part of the misunderstanding surrounding DCIM is the fact that it has often been (and still is) a very broad term. It is often promised by less credible suppliers and manufacturers to do more than its current capabilities allow.

COST A TOP BARRIER TO ADOPTION

An Uptime Institute survey found that more than 60 percent of respondents cited cost as the chief barrier to DCIM adoption. Although there is a wide variance in the cost of DCIM solutions, many budgets don't account for this expense.

Another barrier as it relates to cost is implementation. These differ from the cost of the product, yet are as much of an obstacle for DCIM deployment as any initial product investment. Facility managers need to consider the "costs" beyond deployment such as:

› Time
› Dedicated internal resources
› Education
› Installation
› Integration.

IT and facility team collaboration (or lack thereof) can have an impact on moving forward with DCIM. That's because getting the budget finalised for such a large purchase could mean competing interests and priorities. Although DCIM's functionality hopefully bridges the gap between these traditionally siloed groups — IT and facilities — sourcing funding from these two groups can be problematic.

Figure 3: DCIM Barriers to Adoption[4]

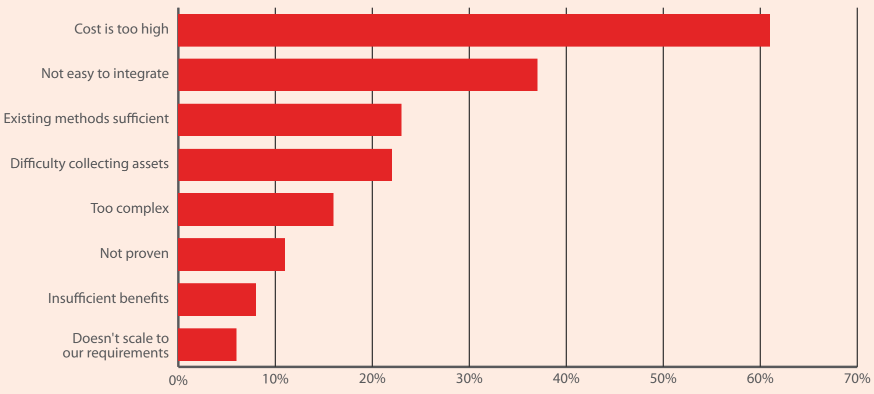

[4] DCD Intelligence, 2015

DIFFICULTY PROVING RETURN ON INVESTMENT

DCIM return on investment (ROI) is achieved through decreased downtime, eased management, increased flexibility and improved transparency. However, measuring these changes is not always easy, which makes proving the rationale for purchasing DCIM to the executive suite and purchasing decision makers difficult.

In general, proving ROI figures is tricky, especially with something like DCIM. When businesses try to bring everything under one centralised solution, and try to take on all issues at once, it makes proving ROI even more difficult. Sometimes the best approach is to find small wins or to solve one particular challenge first. Once the challenge is identified along with an estimate of how much it costs the business, facility managers can also calculate how DCIM would help the business fix the problem. Those can be hard numbers (e.g., inefficient cooling systems tied to electricity costs) or they can be softer numbers (e.g., productivity gains from better business processes gained from the insight that DCIM offers).

> Sometimes the best approach is to find small wins or to solve one particular challenge that has been the biggest issue first.

Calculating ROI definitely depends on the issues that need to be solved. For example, ROI can be calculated by looking at better asset management practices:

› How difficult is it to maintain current asset inventory?
› How accurate is the asset information currently?
› Have there been outages associated with not knowing asset information?
› How long does it take on average to find a particular asset in the data centre?
› Is there an asset inventory process? If so, how long does it take to complete?

ROI AND GETTING APPROVAL FROM FINANCE

There are some key lessons that can help gain approval for DCIM:[5]

› Understanding common, costly data centre problems that can be fixed with DCIM
› Identifying the high-value problems inside a data centre
› Measuring what these problems actually cost
› Learning to align the DCIM project to corporate objectives
› Building, identifying and selling an ROI model to the decision makers

[5] Data Centre Knowledge – How to Develop ROI for DCIM Projects – Bill Kleyman

IT PROCESSES SLOW TO CATCH UP

In order to gain the full benefit of DCIM, previously stand-alone IT processes such as asset provision need to be fully integrated.

Integration requires effort, particularly in bridging protocols. For example, IT systems move information in a digital format and rely on the TCP/IP protocol, while building systems communicate using analogue network protocols, such as BACnet and LonWorks. In order for DCIM to bring differing systems together, bridges need to be built between the two protocols, starting at the network layer.[6] The good news is that there are top-tier DCIM manufacturers and suppliers with tools and solutions that address these integration issues.

This is why working with a trusted partner is paramount for those evaluating DCIM.

FLASH POINT: OVERPROMISING AND UNDERDELIVERING

As mentioned briefly, an unfortunate issue that often surfaces with a young and still-evolving technology platform is the pervasiveness of manufacturers and suppliers that make inflated claims of the advantages of DCIM. Suppliers that fail to deliver on their promises are damaging this fast growing market.

With more than 70 suppliers that claim to have a DCIM solution,[7] no industrywide standards and a lack of a formalised open-architecture approach, this will likely continue to be an issue. As the DCIM market gains traction and grows in maturity, the truly credible and reliable DCIM providers will stand out.

This is why working with a trusted partner is paramount for those evaluating DCIM. Because DCIM is such a significant investment of time and resources, it's important to work with an established and reputable partner that can map a company's challenges into the right solution agnostically with no product bias, inflated claims, or empty promises.

[6] Data Centre Knowledge - How to Develop ROI for DCIM Projects – Bill Kleyman

[7] TechTarget – DCIM Vendor's Promises vs. Data Centre Realities – Paul Korzeniowski, 2014

THE CHALLENGES OF SELECTING A DCIM SOLUTION

When selecting a DCIM solution for a data centre, there are a number of challenges that may arise. Through years of experience in working with customers as well as collaborating with industry experts, these are some of the most common issues that surface in the process of choosing the right DCIM solution.

LACK OF PRODUCT EVALUATION TOOLS

With an estimated 70 suppliers in the market who say they have some form of a DCIM solution in their portfolio,[8] it's hard to decipher the real and credible solutions from the rest. Additionally, there aren't any standards developed around DCIM, which makes evaluating the different solutions difficult.

UNDERSTANDING IMPLEMENTATION COSTS

Fully understanding all the implementation costs so they can be integrated within the IT budget (the most likely source of the purchase dollars) and properly presented to the purchasing decision makers can be a hefty task. Of course, it's important to factor in the effort and cost of implementation when looking at the overall cost of a DCIM solution.

INTEGRATION WITH OTHER SYSTEMS

Selecting the right DCIM solution relies heavily on how well the technology integrates with existing and future systems. Determining integration capabilities can be an obstacle because there are many different ways to implement DCIM. Some argue in favour of a loose framework of best-of-breed components, others for tighter solutions from a single supplier. In the same vein, there are many approaches to integrating DCIM tools into the wider digital infrastructure, and the degrees of integration can vary widely.

[8] 451 Research - The DCIM Market is Still Open to New Entrants – Rhonda Ascierto, 2015

INTEGRATING SOFTWARE AS PART OF DAY-TO-DAY BUSINESS PROCESS

As the software component of DCIM continues to develop, data centre operators committed to using it to manage their facility effectively will have to address the many processes this new technology touches on a day-to-day basis. DCIM is only as good as the action taken to improve discovered inefficiencies in power, cooling, asset utilisation and workflow. To take action, processes need to be carried out. Ideally, these processes are discussed on the front end prior to searching for the right DCIM solution so the solution can be built around the business process versus building the process around the solution, which can lead to a lower adoption rate.

DCIM is only as good as the action taken to improve discovered inefficiencies in power, cooling, asset utilisation and workflow.

INTERDEPARTMENTAL CONFLICT

The ideal DCIM solution represents a true convergence between IT, facilities, and their corresponding management systems. This approach allows the teams to work together to meet organisational objectives. Unfortunately, this convergence doesn't always readily translate into the human equation, as the potential (and common) interdepartmental conflicts may occur between IT and facilities teams.

Figure 4: Data Centre Systems That Can Integrate With DCIM[9]

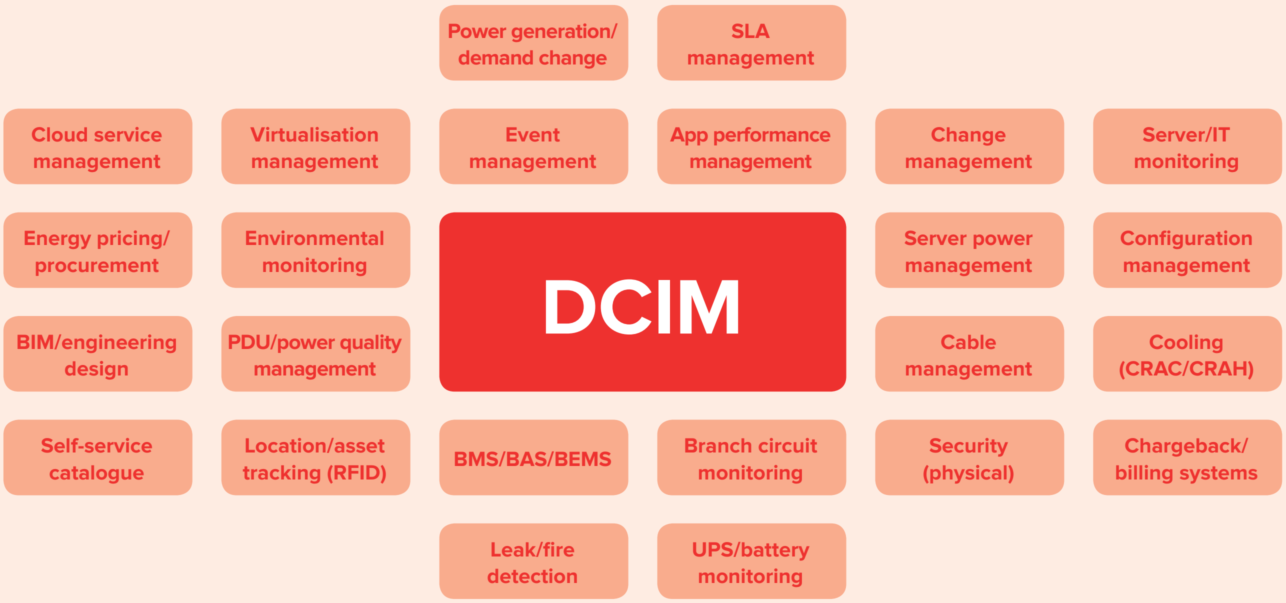

[9] 451 Research – Next Generation Datacentre Management, 2014

IMPORTANCE OF INFORMATION MANAGEMENT

The core goal of DCIM is useful information, in other words, actionable data. Data centre managers are under pressure to run a more efficient operation and doing more with less, and, as the old adage states, you can't manage what you don't know.

THE RISK OF A PARALYSED OPERATION

Collecting and analysing data centre infrastructure information is critical to the decision making process. Not doing so can enhance the risk of data centre paralysis, which means it becomes difficult to:

› Make predictive decisions
› Plan for the future
› Have an understanding of the data centre's current state.

For example, without collecting and monitoring the right information, there can be no basis for current and future planning, which can potentially lead to a state of confusion or disorder in a critical environment. Without timely and comprehensive reports, it's difficult to highlight inefficiencies that are costing the business. Without live feeds, there is an overall lack of awareness, keeping the team out of touch with its environment.

Figure 5: Risk of Data Centre Paralysis

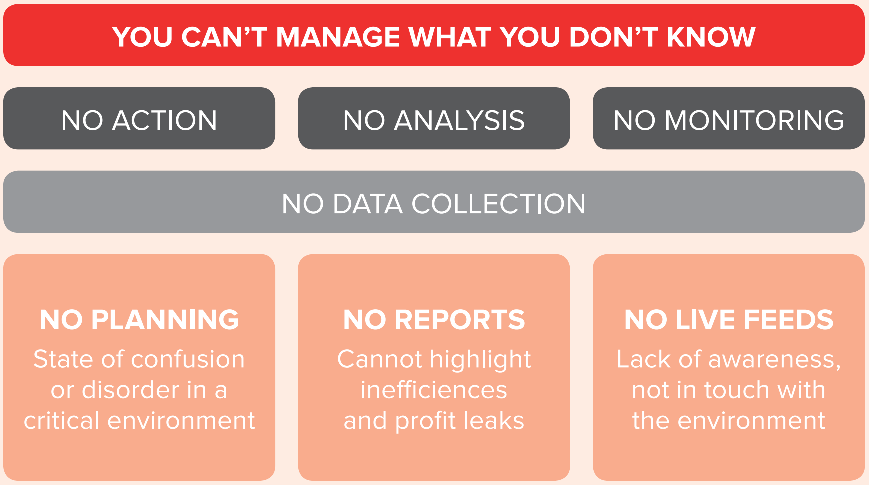
YOU CAN'T MANAGE WHAT YOU DON'T KNOW
NO ACTION
NO ANALYSIS
NO MONITORING
NO DATA COLLECTION
NO PLANNING
State of confusion or disorder in a critical environment
NO REPORTS
Cannot highlight inefficiences and profit leaks
NO LIVE FEEDS
Lack of awareness, not in touch with the environment

IS THE DATA CENTRE READY FOR DCIM?

Before implementing DCIM, it's best to prepare operations in a few key areas. Doing so will help establish a more cost-effective and smooth deployment as well as reap rewards on an ongoing basis.

OPEN THE LINES OF COMMUNICATION BETWEEN IT AND FACILITIES

DCIM holds the promise to open the lines of communication between the IT and facilities groups. However, IT will need to take certain steps in order for DCIM to translate into better data centre efficiency.

IT and facilities teams don't always make talking and working from common data sources a priority. When DCIM is determined to be the way forward, it should be integrated into existing systems so duplication of functionality doesn't occur, which could potentially create a whole new set of data silos.

> As both the IT and facilities teams embrace DCIM together, these separate visions should become simply one lens by which everyone will view the data centre.

Of course, facilities, and IT systems merging does potentially create new managerial issues. Typically, the building management group and the data centre department have worked separately, and often at an arm's length. In order to prepare an organisation for DCIM, there has to be some internal effort for the two teams to work better and more closely together.

It's worth noting that the DCIM solution itself may help mend the fences between IT and facilities groups. Facilities management teams often see the data centre as just another building to be managed. The IT team tends to see the data centre as the entire universe of its professional life. As a result of these differing perspectives and values around the data centre, one group's priorities may not match the other's. The language that each group speaks can also be subtly — sometimes not so subtly — different. As both the IT and facilities teams embrace DCIM together, these separate visions should become simply one lens by which everyone will view the data centre.

A UNITER, NOT A DIVIDER

A good DCIM solution should give all teams the opportunity to collaborate and work toward a common goal of optimum data centre efficiency.[10]

For example, DCIM solutions have the capability to monitor various alerts and alarms, bringing building management and even the security team together. Even though these teams may not have DCIM training in particular, it's a solution that can potentially benefit everyone.

DEPLOY INTELLIGENT HARDWARE

In order for most DCIM tools to work properly and at their maximum potential, the current data centre must feature at least a baseline foundation of intelligent hardware that will collect data at the level of granularity required to solve the business' specific challenges.

For example, if the DCIM software can't communicate with a particular cooling unit or UPS, it will be unable to determine accurate capacities or their current status, which makes planning impossible in real time. Without the intelligent hardware that communicates with the DCIM solutions, answering important questions critical to data centre operations will be challenging. These questions could include:[11]

› Where should I place the next server?
› When will my cooling capacity be exhausted?
› Where am I losing power efficiency?
› What is the impact of this particular change I'm thinking of making?

Reporting and dashboard functionality can also be gravely hampered if inputs are missing or are wrong. For example, a data centre's PUE metrics, which are often reported through a DCIM dashboard, rely almost entirely on the collection and understanding of the connections of many lower-level sensor readings. If the system is unable to communicate with all of the necessary sensors, the PUE metrics reported would be incorrect. Not having a clear picture of power, cooling and environmental conditions at the rack leads to an inaccurate picture of infrastructure capacity and status upon which the DCIM software will make erroneous assumptions, calculations and recommendations.

It's crucial that a data centre makes the investment in physical infrastructure devices and intelligent hardware that can feed the DCIM solution what it needs.

The most effective DCIM solutions work off the continuous input of live data from the physical infrastructure devices and other management systems. These may include intelligent hardware pieces that could communicate to the DCIM server on an on-going basis in order to effectively monitor and plan:

> It's crucial that a data centre makes the investment in physical infrastructure devices and intelligent hardware that can feed the DCIM solution what it needs.

› UPSs
› PDUs
› Power meters
› Environmental sensors and probes
› Security cameras
› Cooling units
› Flow meters
› BMSs
› CMDBs.

At a minimum, UPSs, cooling units, rack PDUs, temperature and humidity sensors should be enabled for network communication.

[10] TechTarget – Data centres define DCIM, not the other way around – Robert McFarlane, 2015

[11] Avoiding Common Pitfalls of Evaluating and Implementing DCIM Solutions – Patrick Donovan, 2012

DETERMINE AND APPORTION STAFF LEVELS FOR DCIM DEPLOYMENT AND UTILISATION

It's paramount for all stakeholders, including management, to agree and commit the necessary resources to implement and operate the solution. All of this upfront discussion and buy-in allows for ongoing cooperation and participation well beyond the implementation phase.

Also, owners for the tools and their associated processes should be explicitly named before the system is implemented. This may be tricky because facilities personnel may be unfamiliar with IT systems while IT personnel may have little knowledge of power and cooling. For this reason among others, it is recommended that evaluation and operation teams include people from both sides to help close any knowledge gaps.

Working close with manufacturers to understand staffing and workforce requirements will help to make the system work effectively. This information will help the evaluation team decide what level of manufacturer (or consultant) -provided training and support might be needed.

DEFINE ACTIONABLE DCIM BUSINESS PROCESSES

Having business processes in place to take action on the gathered intelligence is important. Without processes and a plan to act on the information, all that is left is data — good data, albeit, but no clear path and resources in place to use that data to solve problems and increase efficiency.

DCIM solutions, in themselves, are a good way to fill gaps in operational processes. A credible DCIM solution can help streamline, facilitate and provide a clear dashboard of IT and facilities systems, transforming a complex and disparate environment into something much more cohesive and manageable. However, such a crystallised vision still relies on business processes to implement, operate and maintain the DCIM solution.

A DCIM-ready business process should be addressed in a manageable way, perhaps starting out with core functions and features that are most important, as opposed to attempting to address all processes at once, which will have an impact on the chosen solution. Choosing solutions that are modular will allow the data centre team to address what it needs currently and scale later as required.

Figure 6: Turning Information Into Action

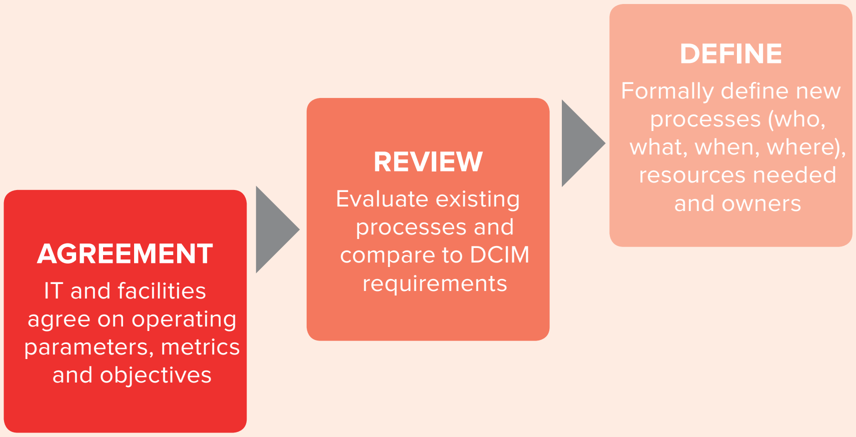

AGREEMENT
IT and facilities agree on operating parameters, metrics and objectives

REVIEW
Evaluate existing processes and compare to DCIM requirements

DEFINE
Formally define new processes (who, what, when, where), resources needed and owners

CONSIDERATIONS FOR DCIM SELECTION

DEVELOP INDIVIDUAL GOALS

DCIM is about data. So, what sort of data is needed to meet team goals for managing a data centre effectively and efficiently? If the goals are defined, diving deeper and answering much more specific questions should naturally lead to a list of individual objectives needed from a product choice.

That being said, there are several considerations that a data centre manager should think about before even looking at specific DCIM solutions. Those considerations are as follows:

- What problems are trying to be solved?
- Of these problems, which are particularly the most pressing issues?
- Why are the current methodologies not working?
- What is the desired end state?

As these questions are answered, the scope of DCIM requirements will be defined. This requirements list should focus on the information needed to manage a data centre based on set goals. Use these requirements and goals as a starting point to evaluate DCIM solutions.

CONSIDER OTHER STAKEHOLDERS' GOALS

It is important for facilities, IT and management teams to work together early on and come to an agreement on the adoption and use of DCIM tools in conjunction with their existing tools.

Conversely, it's a mistake for management to decide to use a DCIM system without the buy-in from those who will be required to implement and operate it. All sides should be involved in the early evaluation phase to make certain everyone's needs and expectations are met. Not only will this secure the right selection for the entire facility, but it's also a positive step in nurturing collaboration with other stakeholders and teams.

> It is important for facilities, IT and management teams to work together early on and come to an agreement on the adoption and use of DCIM tools in conjunction with their existing tools.

ESTABLISH SOME BASELINE CRITERIA

No matter what solution is selected, DCIM tools should have certain "must-have" attributes in order to be prepared for the future and be effective today:[12]

- Scalable, modular, flexible system
- Open communication architecture
- Standardised, pre-engineered design
- Active manufacturer support structure

Using these four characteristics as a high-level baseline for evaluating DCIM tools may certify that the business' processes, data and methods will be in line with expectations moving forward.

[12] Avoiding Common Pitfalls of Evaluating and Implementing DCIM Solutions – Patrick Donovan, 2012

DETERMINE INTEGRATION NEEDS

When selecting a DCIM solution, a key consideration is how this integration will be achieved and how will it be supplied. It is important to highlight that the more systems that need to be integrated, the more expensive and complex a project becomes and the longer the project will take to implement.

DON'T EVALUATE IN A VACUUM – OTHER STAKEHOLDERS SHOULD PROVIDE INPUT

It is important to involve the facility team in all aspects of DCIM deployment and management going forward. This applies directly to the evaluation and selection process because these initial stages will help provide strong buy-in from facilities counterparts.

Consider involving all stakeholders in the selection process:

- Infrastructure and operations
- Facilities (data centre and corporate)
- IT architecture
- Business and technology analysts
- CSR
- Finance (IT and corporate)

Involving all potential stakeholders can lead to a strong consensus on the value of DCIM which helps in funding the investment.

START WITH THE BASICS, THEN MOVE FORWARD

A lot of DCIM implementations become stalled because businesses try to take on too much at once. They attempt to pull everything together under one system in a short time, but find the difficulties of such an endeavor overwhelming. This can lead to frustration and a lack of clear wins along the way, causing the deployment to stall or stop entirely.

According to a DCD intelligence survey, over 35 percent of those who planned to deploy DCIM in 2013 stopped in 2014.[13]

Regardless of the scope of the DCIM solution – comprehensive or on a smaller scale to address a specific issue – it's important to start with the most essential elements that are the highest priorities for an organisation. Ideally, any DCIM solution should be modular, so it's relatively easy to accomplish the most basic integrations in a sequential manner, building upon successes one step at a time.

As the solution is evaluated, consider the most important dashboards that are a priority for the different stakeholders that were involved in the selection process.

Taking a more simplified, realistic and pragmatic approach will help avoid overwhelming an organisation – both in terms of costs and workforce hours – as well as prevent information overkill and project fatigue.

Figure 7: How 2013 Intentions Toward DCIM Panned Out in 2014: Percent Census Samples 2013 and 2014

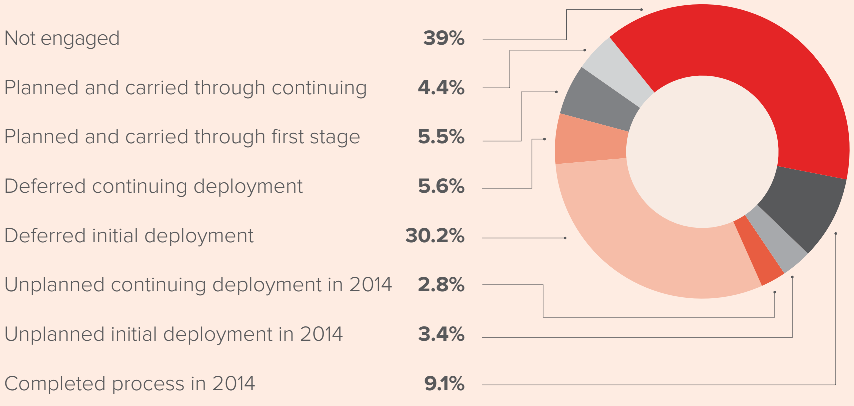

[13] DCD Intelligence – 2013 – 2014 Census Survey

THE FIVE SENSES OF DCIM

In the same way humans have a variety of senses, DCIM can be looked at similarly: information is collected and interpreted to solve specific business challenges. Anixter has defined five senses of DCIM, each one solving a particular business challenge:

› Asset management
› Power monitoring
› Environmental monitoring
› Change management
› Capacity planning

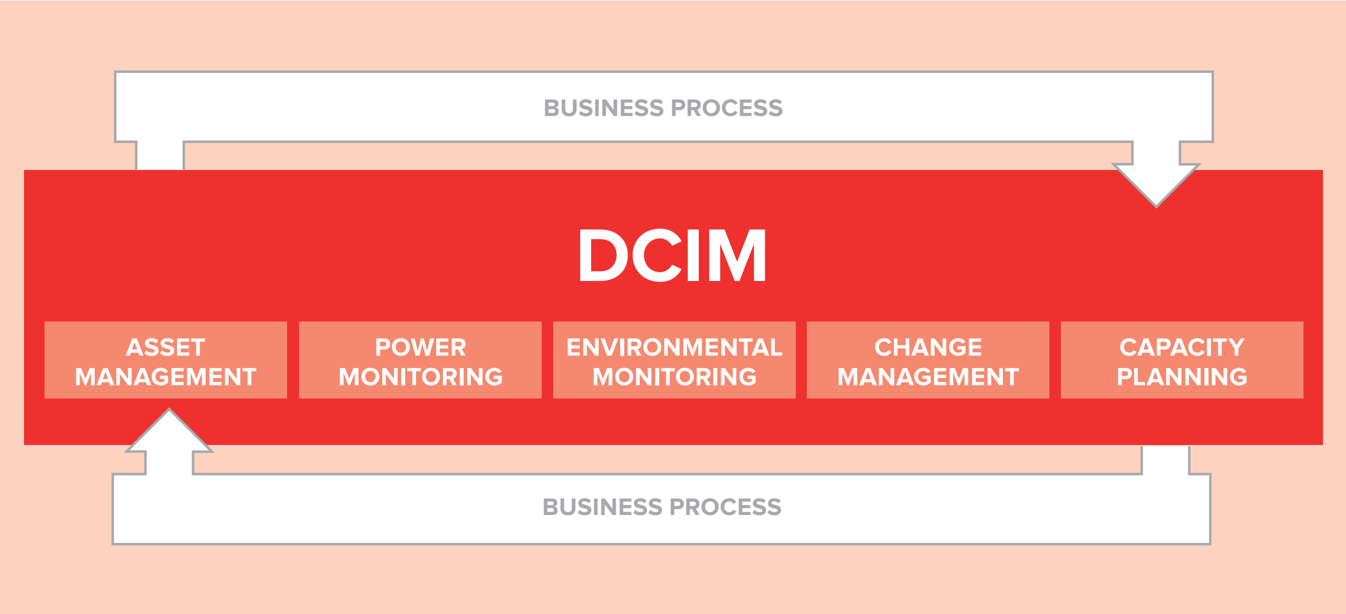

THE FIVE SENSES OF DCIM 1 OF 5

ASSET MANAGEMENT

1

Asset management is a broad topic, but for the purpose of this paper, it will focus on the assets that are generally governed by a DCIM tool, which are the hardware assets throughout the data centre. IT asset management (ITAM) entails collecting inventory, financial and contractual data to manage the IT assets throughout their life cycle. ITAM depends on robust processes with tools to automate manual processes.[14]

Having and maintaining an accurate asset database is critical to the data centre. Arguably, asset management is the most important function of all of the five senses of DCIM. Not having a clear understanding of where IT assets are located throughout the data centre can negatively impact the daily operations of the facility.

CONSEQUENCES OF POOR ASSET MANAGEMENT

INCREASED DURATION OF IT HARDWARE DOWNTIME	Hardware downtime is inevitable; it is going to happen. However, the ultimate goal is to alleviate the problem as quickly as possible. Not knowing for sure where a potential server or piece of IT hardware is physically located can increase the time it takes to diagnose and fix the outage.
INCREASED TIME TO DEPLOY NEW EQUIPMENT	When deploying new IT equipment, an operator needs to understand the ideal place to put it. If a physical walkthrough is necessary to determine available rack space, network connectivity and available power, it increases the amount of time it takes to deploy the equipment.
INCREASED FINANCIAL RISK	Asset management goes beyond just physical locations of devices. Not knowing what an asset is, who it belongs to and what the impact of losing that particular asset has on the business can pose a significant financial risk.
INEFFICIENT USE OF AVAILABLE CAPACITY	In the data centre, capacity is everything. Once that capacity runs out, a new data centre needs to be built. Not knowing what is in the environment, how it is interconnected, how much power and cooling it consumes, and how much physical space it takes up makes planning for the future almost impossible.

ASSET MANAGEMENT

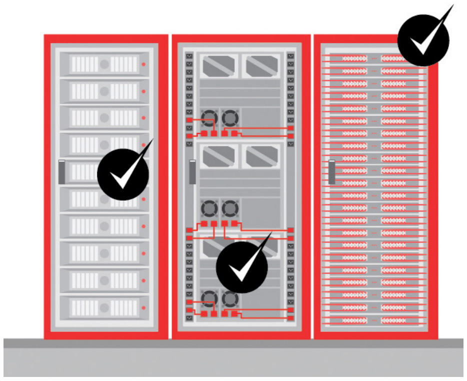

[14] Gartner IT Glossary

IT ASSET MANAGEMENT GOALS

It is important to define your objectives when looking for the right asset management solution for the data centre. Many DCIM solutions have asset management capabilities, but some are more robust than others. Generally, the main goal for most data centres in regard to asset management is to have a single source of truth for all hardware within the server room. The other aim is to build a virtual model of that data. That virtual model helps a user make decisions on new and existing assets faster.

What is going to be different, depending on the goal of the business and individual, is how much information is required per asset. Table 1 breaks down the areas asset management can define into three main categories: physical location, asset configuration and asset ownership.

> Generally, the main goal for most data centres in regard to asset management is to have a single source of truth for all hardware within the server room.

Table 1: What Asset Management Can Help Define

PHYSICAL LOCATION	ASSET CONFIGURATION	ASSET OWNERSHIP
Room name and location	Power connections	Purchase date
Rack name and location	Network connections	Purchase price
Rack unit number	Unit size and weight	Supplier purchased from
	Virtual hosts	Department owner
		End of life date
		End of lease date

ASSET MANAGEMENT BEST PRACTICES

Once it has been agreed that a centralised database is required to monitor all IT and facilities assets, the first thing that needs to be determined is what the process is going to be to track current and future equipment. Without having a solid process in place and mandating that the process be followed, maintaining the data integrity of a future asset management solution becomes difficult.

Some things to consider when building the asset management process:

- Having a process review to uncover gaps
- Restricting entry into the data centre
- Dedicating centralised resources to keep asset information up to date by maintaining records
- Mapping out desired future state – what it will look like after the tool is implemented

Once the process for current and future asset management has been discussed and agreed upon by all involved parties, the next step is to implement that process. It's best to start small, which could mean one location or several new IT equipment orders. This is also a good time to determine if there are existing tools in place today that are being used by different departments to maintain and track assets and how those tools should interface with the new tool being looked at.

Next, a complete asset inventory is required to make sure that all asset information is as accurate as possible and in a format that it can be easily imported into the new tool. Many times, asset software will provide templates of the necessary information needed. Additionally, many manufacturers will offer asset management services to perform this task if it cannot be completed by the data centre staff.

It might make financial sense to have the asset collection performed by a third party. The typical cost to collect "readily visible" data (manufacturer, model, location, serial number and device name) is 15 USD per device. Figure 8 illustrates example costs of a data centre with 8,000 assets.[15]

Figure 8: Data Centre Asset Collection Costs in an 8,000 Asset Data Centre

MANUFACTURER COST	DOING IT YOURSELF
Initial collection of basic device information could cost $120,000.	Collecting basic device information would require 40 weeks of effort from a workforce.
Including detailed system data, it could drive the data collection cost to $600,000 or more.	Collecting detailed system data would require 200 weeks of effort.

[15] Data Centre Knowledge – Data Centre Infrastructure Management, David Cole

Once the asset inventory is complete, deployment of the new tool should happen quickly. The reason for this is to maintain asset integrity; that way it is as close to 100 percent accurate once cutover to the new tool happens. It will be important at this time to train anyone who needs to have access to the tool. To make sure that the new tool is being used properly, spot checks on inventory could be performed after the first month or quarter of use.

Finally, to be successful it is critical to maintain asset integrity. This can be done by performing, at a minimum, annual asset audits. To make the inventory process faster, assets should have some type of bar code (asset tag) or even RFID (radio frequency identification).

BENEFITS OF RFID

- Automates asset moves, adds and changes
- Decreases risk of physical asset theft
- Speeds up physical inventory audits

Figure 9: Asset Management Best Practices

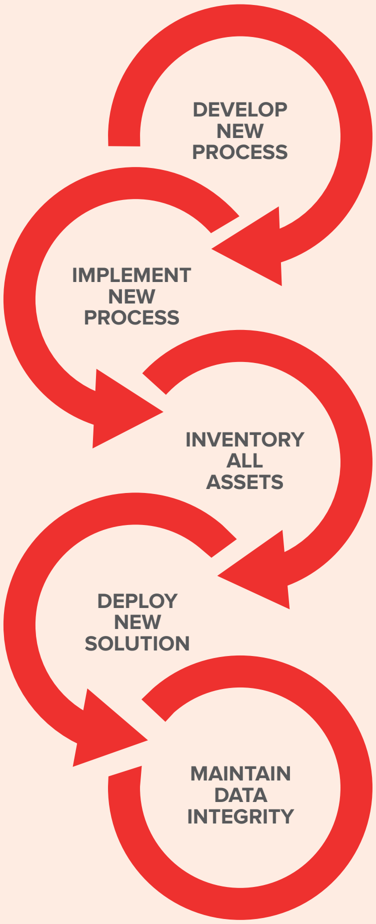

THE FIVE SENSES OF DCIM 2 OF 5

POWER MONITORING

2

Monitoring facility power systems has been around for quite some time. However, what is changing is the desire to monitor closer to the IT equipment to get a better understanding of how power is being consumed at the cabinet. This drive to take a more holistic approach to monitoring the entire power chain can be seen by an increase in adoption of intelligent cabinet PDUs. Today's data centre is consuming more power per square foot than it ever has in the past,[16] which is being driven due to an increase in IT hardware demands. This increase in IT demand while using the same amount of physical space has led to an increase in the cost of an outage. In a 2013 survey conducted by the Ponemon Institute, it was found that the cost of outages ranged from 45 to 95 USD per square foot[17] which was a 41 percent increase from 2010.

This sharp increase in the cost of a data centre outage as well as the increasing demands to reduce operational expenses are major influencers in a facilities manager's decision to invest in power monitoring solutions in the data centre. In fact, according to DCD Intelligence, the number one DCIM capability investment that data centres were making was in real-time monitoring.

Figure 10: Cost of Data Centre Outages

SURVEY OF 67 INDEPENDENT DATA CENTRES		
Facility manager	CIO	DC manager
CTO	IT compliance	Security officer

The cost per square foot of data centre outages now ranges from 45 to 95 USD	Minimum cost of 74,223 USD to a maximum of 1,734,433 USD

POWER MONITORING

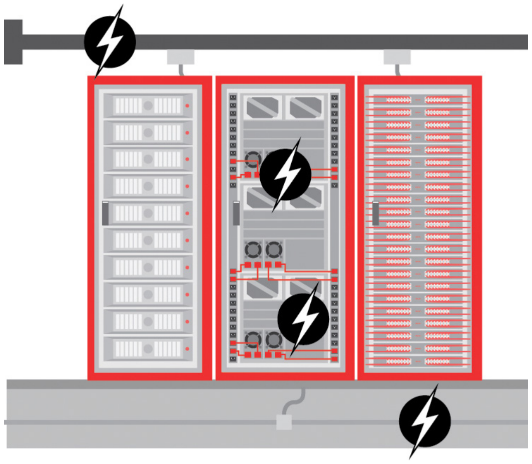

[16] DCD Intelligence

[17] Ponemon Institute – Cost of Data Centre Outage, 2013

POWER MONITORING GOALS

Prior to implementing a power monitoring solution, it is important to understand the business' goals. Some of the goals that can be accomplished through following power monitoring best practices are:

- Increasing energy efficiency, thereby reducing operational expenses
- Deferring capital costs on new equipment by maximising available power capacity
- Reducing the amount and duration of future outages.

CONSEQUENCES OF NOT MONITORING

DIFFICULT TO UNDERSTAND CAPACITY NEEDS

Measurement of the available capacity throughout the power distribution system is vital to make certain that there is room to support future IT needs. Not planning for the future can lead to outages and delays in deploying new business applications.

POOR VISIBILITY INTO POWER SYSTEMS

Having visibility into the entire power chain, and all systems which support it, is important to make sure the system is being run as efficiently as possible. According to The Green Grid Association, to have full insight into an infrastructure's energy efficiency, multiple components from the utility entrance through the IT equipment should be monitored. This gives facilities and IT the full picture to make better deployment decisions supporting future projects.

INCREASED RISK OF UNPLANNED OUTAGES

Better measurement and data can aid an IT or facilities manager in understanding what caused an outage and potentially prevent an outage from happening in the future. Not having this information can paralyse decision making due to a lack of visibility into what is happening real time.

Figure 11: Investment in Specific DCIM Capabilities – Percent of 2013 Census Respondents[18]

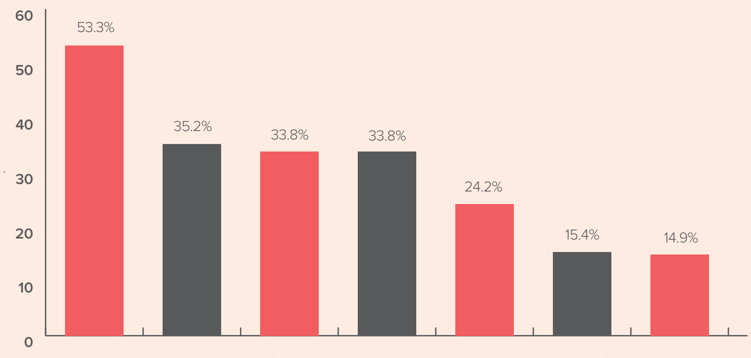

[18] DCD Intelligence – 2013 – 2014 Census Survey

It is also imperative to define who owns different aspects of the power chain. For example, if the data centre is in a co-location facility then the delivery of power from the entrance of the building to the data hall generally falls to the owner of the co-location, and the responsibility of the enterprise inhabitant is usually within the IT cabinet itself. The power chain can be monitored at multiple points, which are:

- Entrance feed
- UPS systems
- Room distribution
- Cabinet distribution
- IT equipment.

> Establishing organisational goals versus individual goals is important to gain acceptance across multiple departments of a potential new power monitoring solution.

In order to meet the business' goals, it is important to know how these systems are currently being monitored, and if they are being monitored, it is important to aggregate the data into one centralised tool.

Lastly, establishing organisational goals versus individual goals is important to gain acceptance across multiple departments of a potential new power monitoring solution. For example, if the individual goal is to reduce the need to send IT staff to remote locations to troubleshoot single phase UPS failures, what would that mean for other departments? Perhaps those remote locations support important financial data for accounting, which would make the IT solution beneficial for them as well.

POWER MONITORING BEST PRACTICES

Once goals have been established, it becomes important to assess the data centre's current situation:

- Current hardware — can it collect the required information needed to meet the organisation's targets?
- Current monitoring tools — what methods are already being used? To what extent should these tools integrate into the proposed power monitoring solution?
- Top power related challenges — target the problems that are causing the most issues first.
- Power chain connectivity map — is there one in place today? Is it accurate? Having this information is critical to understanding the impact of changes to the entire power chain.

Figure 12: Power Management Software Systems That Can Integrate With DCIM[19]

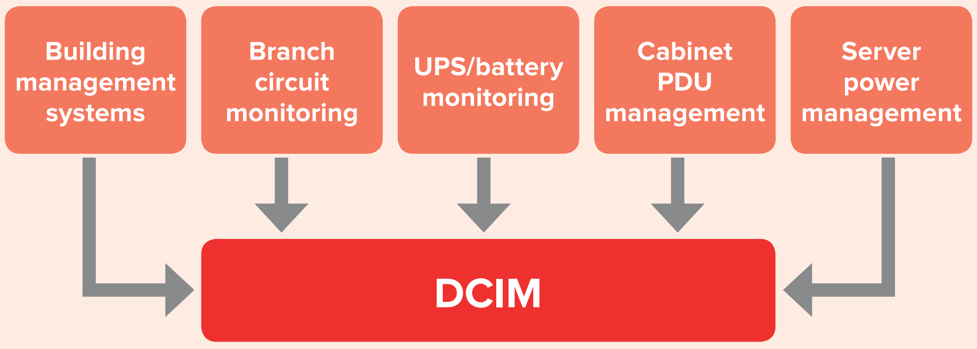

[19] 451 Research – Next Generation Datacentre Management, 2014

Using the agreed-upon business objectives, the scope for the power monitoring application should be constructed to clearly highlight the purpose of power monitoring to all stakeholders. Having a narrow scope will allow the business to get faster returns on the initial investment, which can make a stronger case for deeper integration later. Once the scope has been established, acceptance criteria should be defined to measure success against. Next, a deployment schedule is required to identify who owns the completion of each project stage.

Having a narrow scope will allow the business to get faster returns, which can make a stronger case for deeper integration later.

Daily or weekly status reports help to keep the project moving and all the stakeholders informed. Some of the questions that might come up while planning the deployment of the system are:

- Does intelligent hardware (e.g., metered cabinet PDUs) need to be installed prior to deployment?
- Will deployment affect any production systems?
- What other departments need to be involved throughout deployment?

Finally, the last step is to examine the ROI by determining if the acceptance criteria have been met and if the DCIM solution was worth the investment.

Figure 13: Power Monitoring Best Practices

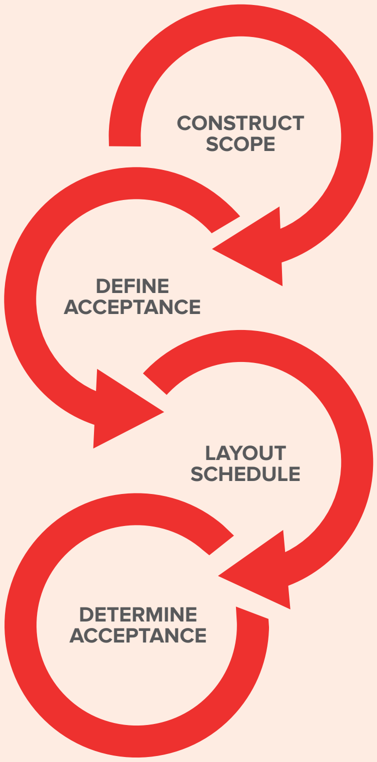

THE FIVE SENSES OF DCIM 3 OF 5

ENVIRONMENTAL MONITORING

3

In a 2014 report titled “Trends in Data Centres”, more than 220 professionals were asked to describe one area they would change about their data centres. 19 percent of them stated that they would like to make their data centre more efficient.[20] Additionally, in a census study performed by DCD Intelligence, the two largest drivers associated with the adoption of DCIM were investing to improve energy efficiency and investing to improve operational efficiency.

Figure 14: Investment Drivers Associated With DCIM: Percent Census Samples 2013 and 2014[21]

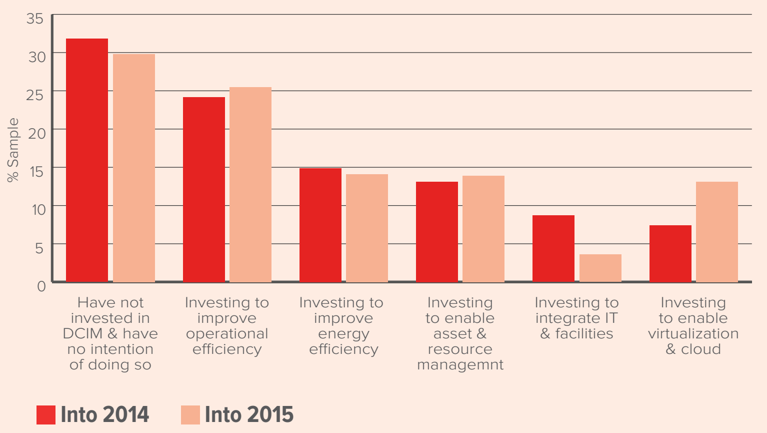

The approach to environmental monitoring in the data centre is holistic, and each individual area is inextricably interrelated to another separate area. Much like power monitoring, the areas that are owned by the business should be looked at as a whole to get a complete picture. Making adjustments to one part of the system can have a major impact on other areas throughout the data centre.

There are several factors that make thermal monitoring a high priority in any data centre:

- In legacy data centres, the cooling system consumes a large amount of energy.
- Increasing rack densities can create unique thermal management challenges.
- Constantly changing IT requirements require cooling to be available on demand.
- High-availability environments need well-controlled temperature ranges for reliable performance.

ENVIRONMENTAL MONITORING

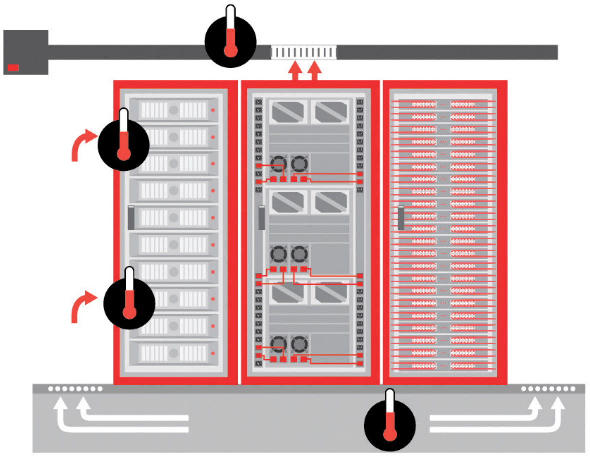

[20] Mortensen – Insights Into What’s Next: Trends in Data Centres, 2014

[21] DCD Intelligence – 2013 – 2014 Census Survey

CONSEQUENCES OF NOT MONITORING

INCREASED OPERATIONAL EXPENSES

The cost of cooling the data centre can be greater than 30 percent of the total energy consumed by the facility.[22] Poor thermal management practices mean an increase in energy consumption by the cooling equipment.

INCREASED CAPITAL EXPENSES

In a 2013 study performed by Upsite Technologies, an average of 48 percent of supply air is bypass airflow.[23] Bypass airflow is cool, conditioned air that never reaches the IT equipment, and if undetected, it can lead data centre managers into believing their room requires additional cooling units to meet the needs of the IT load.

INCREASED RISK OF EQUIPMENT FAILURE

IT equipment requires a consistent temperature to secure reliability. Not monitoring the environment throughout the data centre can lead to imbalanced temperatures. Temperature fluctuation at an IT equipment rack is airflow recirculation. Airflow recirculation can be defined as hot exhaust air that passes through the IT equipment multiple times before it travels back to the cooling system. It can be detected through temperature deltas from the bottom of the rack to the top of greater than five degrees Fahrenheit.

Figure 15: Imbalanced Supply Temperature[24]

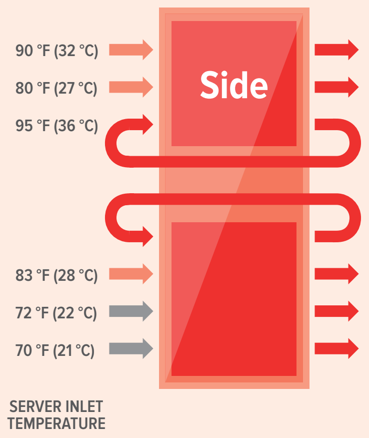

Table 2: The State of Airflow Management

2011 CLASSES	2008 CLASSES	APPLICATIONS	IT EQUIPMENT	ENVIRONMENTAL CONTROL
A1	1	Data centre	Enterprise servers, storage products	Tightly controlled
A2	2		Volume servers, storage products, personal computers, workstations	Some control
A3	NA		Volume servers, storage products, personal computers, workstations	Some control
A4	NA		Volume servers, storage products, personal computers, workstations	Some control
B	3	Office, home, transportable environment, etc.	Personal computers, workstations, laptops and printers	Minimal control
C	4	Point-of-sale, industrial, factory, etc.	Point-of-sale equipment, ruggedised controllers, or computers and PDAs	No control

[22] Data Centre Alliance Project

[23] Upsite Technologies – The State of Airflow Management, 2015

[24] ENERGY STAR – Properly Deployed Airflow Management Devices

ENVIRONMENTAL MONITORING GOALS

The purpose of a data centre's thermal management strategy is to make sure the room that the equipment resides in is at a temperature that is within range of the required operating temperatures specified by the equipment manufacturer.

The traditional method of doing this was to flood the room with cold air, and if the room became too hot, add more perforated tiles, lower the temperature set points and finally add additional cooling capacity until the desired temperature is achieved.

The purpose of a data centre's thermal management strategy is to make sure the room that the equipment resides in is at a temperature that is within range of the required operating temperatures specified by the equipment manufacturer.

Conditional environmental control is the process of delivering the exact amount of supply air at an ideal temperature and moisture content to maximise the cooling system's efficiency and improve equipment uptime. The only way to truly gain insight into the thermal environment is to monitor different points in the cooling system.

There are five areas that should be monitored to achieve an ideal thermal environment in the data centre:

- Supply airflow pressure
- Supply airflow volume
- Supply airflow temperature
- Bypass airflow
- Airflow recirculation

Figure 16: Bypass Airflow[25]

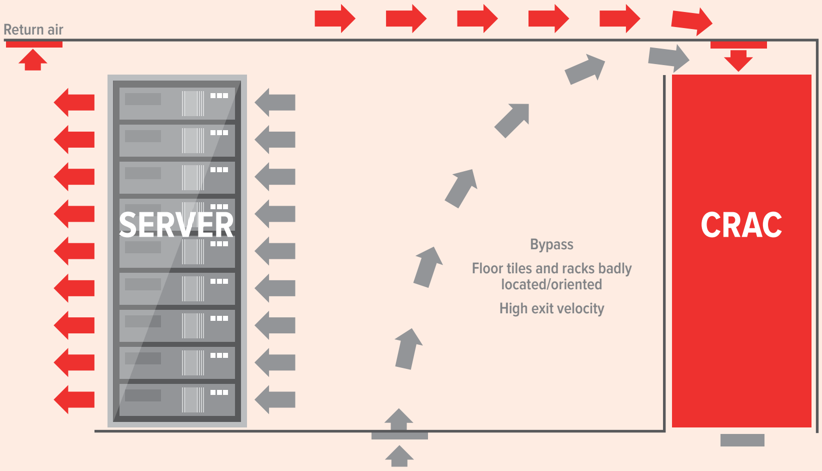

Figure 17: Airflow Recirculation[26]

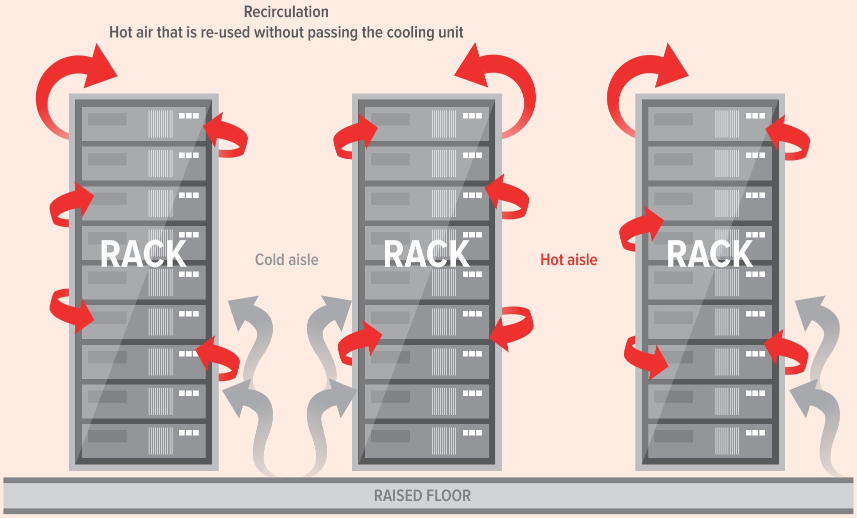

[25] Improving Data Centre Air Management – Munther Salim and Robert Tozer, 2010

[26] Open Data Centre Measure of Efficiency – Paul Poetsma

A data centre manager knows when they have achieved an optimised thermal environment when:

- Bypass airflow is less than 10 percent
- The temperature measured at top and bottom of the IT equipment cabinet has a delta of less than 5 degrees Fahrenheit
- The temperature at the hottest point in the data centre falls within 60.4 and 80.6 degrees Fahrenheit (18 and 27 degrees Celsius)

Figure 18: Properly Balanced Supply Side Temperature

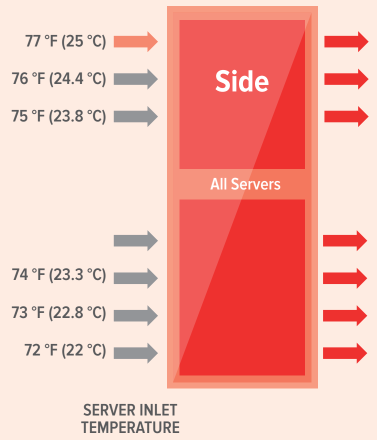

Figure 19: Differential Sensor Placement in a Raised-Floor Environment

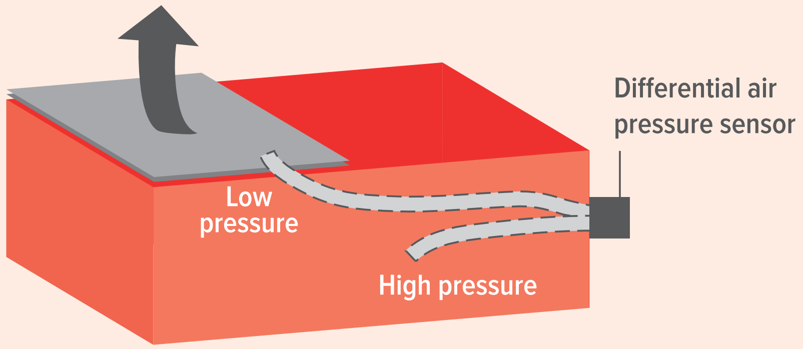

ENVIRONMENTAL MONITORING BEST PRACTICES

Environmental monitoring can range from understanding why certain cabinets are experiencing higher temperatures to reducing operational costs across the facility by 10 percent. Again, a larger goal could require integration with existing systems. Additionally, because power and environmental monitoring are generally part of the same software package, it will be important to map in those business goals that help with the software selection process.

Generally, in the context of a DCIM solution, environmental monitoring is deployed in the data hall itself, and then later integrated with the cooling units if the business desires real-time control. Prior to making any changes to the cooling system, it is important to make sure the proper data collection hardware is deployed throughout the data hall, which includes airflow pressure and temperature sensors.

Airflow pressure sensors in a raised-floor environment should be placed every 1,000 square feet a few feet above the subfloor. The sensor tubes should be positioned in different pressure areas (e.g., above and below the raised floor and in the hot and cold aisle). Additionally, they should be installed no closer than 12 feet from a cooling unit.

Airflow temperature sensors should be deployed per ASHRAE guidelines, every third rack, at the supply side with one sensor aligned with the top U, one in the middle U and one at the bottom.

It is important to get a baseline reading throughout the facility, which will provide a roadmap to help reach the data centre's thermal management objectives.

Once a hardware and software solution has been installed, it is important to get a baseline reading throughout the facility, which will provide a roadmap to help reach the data centre's thermal management objectives.

Figure 20: Environmental Monitoring Software That Integrates With DCIM[27]

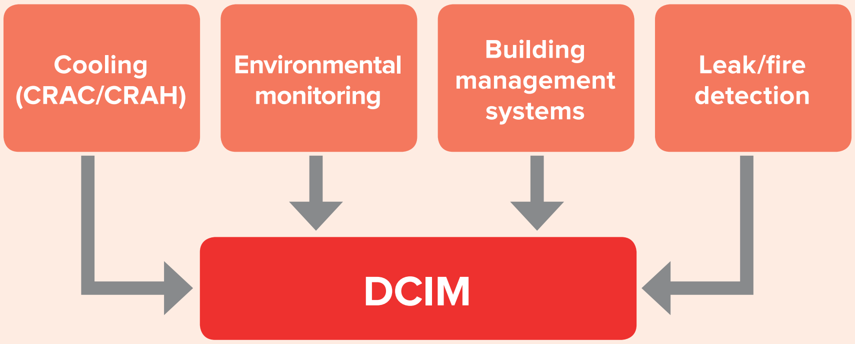

Airflow pressure sensors coupled with information from the cooling units will help determine if there is bypass airflow, which means the volume of air leaving the cooling units isn't reaching the intended IT load. The temperature sensors will show if there is airflow recirculation by showing the temperature delta from the bottom to the top of the cabinet.

After a baseline reading has been established, the next step is to perform a physical walkthrough to correct the airflow management problems identified by the environmental monitoring solution that was deployed.

To reduce bypass airflow:

- Walk through the data centre to visually identify the sources, which can include cable openings in the raised floor, unsealed holes in the perimeter walls under the raised floor and holes behind the air handling units and at the bottom of building columns.
- Seal all identified sources of bypass airflow with floor grommets or fire-rated material for holes in the perimeter walls underneath the raised floor.
- Ensure all perforated tiles are in the cold aisle; there should be none in the hot aisle.

To reduce airflow recirculation:

- Seal all gaps in the cabinets with blanking panels.
- Seal all gaps in the sides, tops and bottoms of cabinets with accessories from the cabinet manufacturer.
- Walk through the data centre and look at the rear of the cabinets to see that the cables are properly dressed and do not interfere with IT exhaust vents.
- Adopt a hot- and cold-aisle layout if possible.

Figure 21: Environmental Monitoring Best Practices

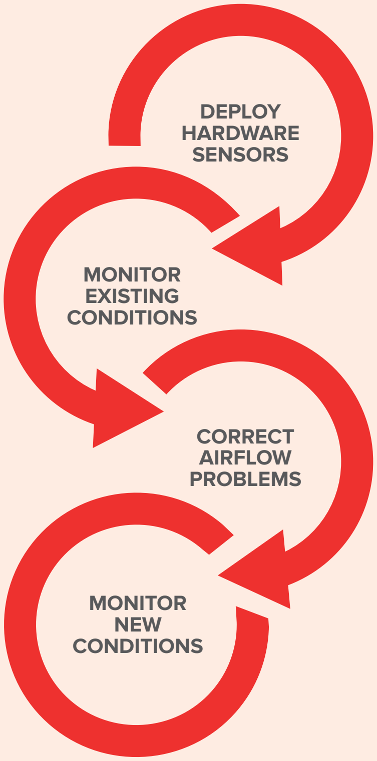

Once the airflow pathway concerns have been addressed, then it is important to get another reading to see how the changes have impacted the data centre. After this knowledge has been gained, then a data centre manager can begin to:

- Adjust temperature set points
- Adjust chilled water temperatures
- If EC or VSD fans are installed at the cooling units, adjust the volume of supply air
- Remove or add perforated floor tiles
- Turn off cooling units.

[27] 451 Research – Next Generation Datacentre Management, 2014

THE FIVE SENSES OF DCIM 4 OF 5

CHANGE MANAGEMENT

4

Change management, sometimes referred to as workflow management, has two important functions in the data centre: tracking any change that is made to an asset throughout the facility and providing insight in how a change to one asset affects other associated assets. The change management aspect of DCIM takes the analytical information that is gathered from asset information, power and environmental monitoring and helps to streamline the execution process.

CONSEQUENCES OF A LACK OF CHANGE MANAGEMENT

INCREASED RISK OF OUTAGES	According to the IT Process Institute's "Visible Ops Handbook", 80 percent of unplanned outages are due to ill-planned changes made by administrators (operations staff) or developers. Many of these outages are caused by a misconfiguration in a hardware installation or not understanding the impact an IT asset has on another asset.
60 PERCENT OF AVAILABILITY AND PERFORMANCE ERRORS ARE THE RESULT OF MISCONFIGURATIONS.[28]	
DECREASED SPEED OF DEPLOYMENT	Not having an automated system for change requests means that tasks for different departments need to be created and managed manually. Additionally, a physical walkthrough is required to determine the ideal location to place new equipment based on power, cooling, space and network connectivity.
DECREASED STAFF PRODUCTIVITY	IT and facilities staff find it more difficult to focus on strategic initiatives as a result of needing to manually process and and update change orders before they can move on to other groups for completion.

The change management aspect of DCIM takes the analytical information that is gathered from asset information, power and environmental monitoring and helps to streamline the execution process.

CHANGE MANAGEMENT

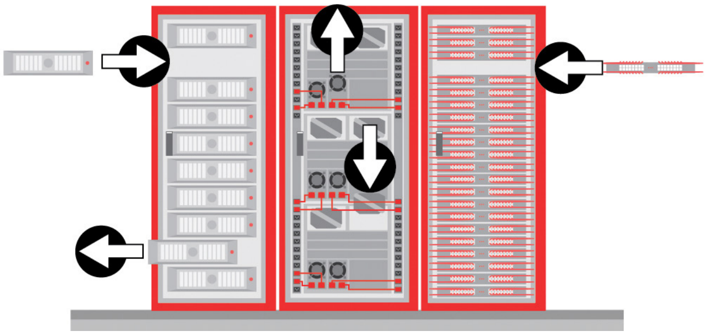

[28] Enterprise Management Association

CHANGE MANAGEMENT GOALS

According to the Information Technology Infrastructure Library (ITIL), the goal of the change management practice is to establish that standardised methods and procedures are used for efficient and prompt handling of all changes in order to minimise the impact of change-related incidents upon service quality and consequently improve the day-to-day operations of the organisation.

As it relates to the data centre, using a DCIM solution along with integrating into existing ticketing/change management systems should:

- Increase the productivity of IT and facilities staff
- Improve the accuracy of hardware installations
- Improve the communication and collaboration between different IT and facilities teams
- Provide better visibility into the impact of changes before they are made, thereby reducing outages
- Improve or establish a clear process for all move, add and change work.

Generally, most data centres have some form of change management system in place. Some examples would be ServiceNow and BMC Remedy. Those systems are designed for IT and service management. Being able to pull workflow information such as approvers, who is performing work for a given ticket, order and receiving statuses, and then prepopulating the asset information into the DCIM solution to automate reserving rack space, power, and network connections can decrease time to deploy and limit configuration errors.

BENEFITS OF CHANGE MANAGEMENT INTEGRATION

INCREASE THE PRODUCTIVITY OF IT AND FACILITIES STAFF

IMPROVE THE ACCURACY OF HARDWARE INSTALLATIONS

IMPROVE THE COMMUNICATION AND COLLABORATION BETWEEN DIFFERENT IT AND FACILITIES TEAMS

PROVIDE BETTER VISIBILITY INTO THE IMPACT OF CHANGES BEFORE THEY ARE MADE, THEREBY REDUCING OUTAGES

IMPROVE OR ESTABLISH A CLEAR PROCESS FOR ALL MOVE, ADD AND CHANGE WORK.

CHANGE MANAGEMENT BEST PRACTICES

The current workflow solution deployed in the data centre will help to scope what is required from a change management functionality standpoint in a DCIM solution. DCIM can either integrate with existing ticketing systems that have workflow built in or provide predefined workflow models. When looking for change management in a DCIM solution, it is important to identify solutions that meet these requirements:

- Flexible – Every organisation's needs are different and a predefined workflow should be tailored to support it.
- Open – In order to get the information required to manage data centre workflow processes, the DCIM solution may be required to interface with ticketing systems, IT and facilities management systems, financial systems and asset management systems.
- Robust – Because change management touches so many different systems, it is important that the solution provide a robust set of integration and reporting tools to provide data centre operators and management the views that they need into the status of particular projects as well as provide the data needed to improve current processes.

Prior to implementing a change management solution, it is important that the first three senses of DCIM – asset management, power monitoring and environmental monitoring – are in place.

Prior to implementing a change management solution, it is important that the first three senses of DCIM – asset management, power monitoring and environmental monitoring – are in place. Change management is the key link between organisational planning and execution. The data gathered from those three areas will provide the information needed to streamline the execution process. For example, it is difficult to do cause and effect analysis by removing a UPS from the environment without knowing what equipment is connected downstream (asset management) and how the increase in power load will affect other UPS systems (power monitoring).

Figure 22: New Server Deployment Typical Workflow[29]

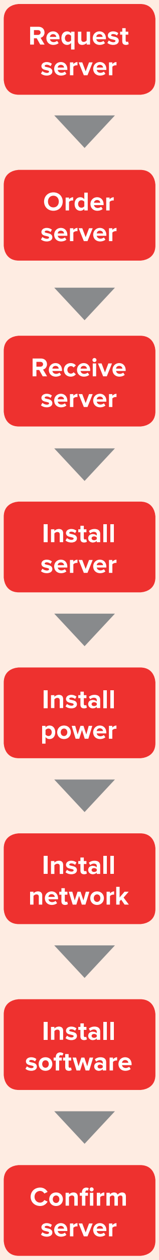

[29] Data Centre Handbook – Hwaiyu Geng

THE FIVE SENSES OF DCIM 5 OF 5

CAPACITY PLANNING

5

Monitoring and insight into the data centre's current state is extremely important to the day-to-day operations of the facility. However, having insight into the future, understanding when a facility is going to run out of power, cooling, network connectivity, or physical space is invaluable to the business. Not only does it help with planning for the future, but having detailed knowledge of the current data centre's capacity can help defer unnecessary capital costs through better understanding of how it can be reclaimed by eliminating hardware that is no longer needed and allocating that capacity to new business applications.

> 56 percent of manual planners need to devote more than 40 percent of their time, every month, to capacity planning and forecasting.

The majority of data centres today are estimating their capacity for new IT equipment; however, it is largely being done manually. According to a recent Intel survey of 200 data centre managers across the U.S. and U.K., 43 percent rely on manual methods for planning and forecasting.[30] The problem with manual entry is the shear amount of time it takes to record and maintain the information. In the same study, Intel noted that 56 percent of manual planners need to devote more than 40 percent of their time, every month, to capacity planning and forecasting.

In order to accurately understand the capacity within a data centre, the first three senses of DCIM need to be implemented to collect the necessary information required. It is impossible to report on physical space available within a given IT cabinet without knowing what is inside that cabinet, that goes the same for available power and cooling to that cabinet as well. That information needs to be gathered from asset management, power and environmental monitoring systems.

CAPACITY PLANNING

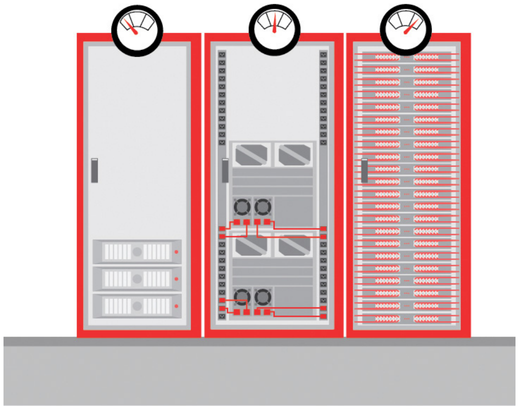

[30] Intel – Inefficiencies cost data centres time and money

CAPACITY PLANNING GOALS

Capacity planning allows data centre managers to plan for the future more effectively through the use of granular data on power, cooling, network connectivity and physical space. DCIM can allow for the aggregation of this information from different disparate systems to allow for more informed decisions, which will help:

- Quickly respond to changing IT demands
- Postpone new data centre construction until it is required
- Understand where inefficiencies lie and where capacity can be reclaimed
- Speed up new hardware deployment
- Minimise risk of downtime by preventing future outages.

Figure 23: Senses of DCIM Required for Capacity Planning

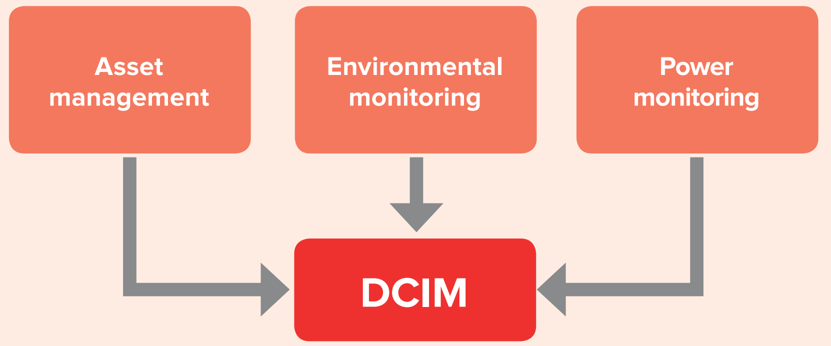

CONSEQUENCES OF POOR CAPACITY PLANNING

UNPLANNED CAPITAL EXPENSES	Not knowing if capacity is running out and then having to react to the needs of the business can lead to unplanned hardware purchases that aren't budgeted for.
DELAYED BUSINESS GROWTH	Additional computing capacity should be available when the business needs it to spur growth. Not being able to support a new application because there isn't available capacity in the data centre can negatively affect potential revenue streams.
INCREASED OPERATIONAL COSTS	Failing to understand how much power or cooling is being used at a given point means that those resources might be consumed by hardware that is no longer required to support the business. These devices are consuming available capacity that could be used for additional expansion.

CAPACITY PLANNING BEST PRACTICES

In order to have an accurate representation of capacity throughout the data centre, it is important to first figure out what is most important to track. For instance, if the data centre is in a co-location facility, cooling capacity might not be important to understand because it isn't being managed directly, the co-location is managing that. However, in an owner-operated data centre, it is important to have visibility into the power, cooling and networking systems as well as have access to the physical asset data to be able to accurately forecast for the future.

Once there is an understanding of what is important to track, the next step is to aggregate data from various IT and facilities monitoring solutions into a centralised solution. This will help uncover any gaps in the data and also help to establish the current baseline of the facility. Also during this stage, it is critical to make sure the data that is being collected is accurate. Understanding current and planning for future capacity is only as good as the integrity of the data. If there is doubt, manual audits should be performed on the power, cooling and assets to be certain they are accurate.

After the data has been aggregated to a centralised point, then it becomes important to analyse that information to uncover some focus areas of improvement. Target improvement of the immediate business needs in order to get some wins before implementing wide scale changes.

The last step is continually improving and optimising the data centre environment. Because IT demands are always changing, it is important to continue to manage and optimise the capacity within the data centre. This will allow for expansion without over-provisioning as well as provide business continuity.

Figure 24: Capacity Planning Best Practices

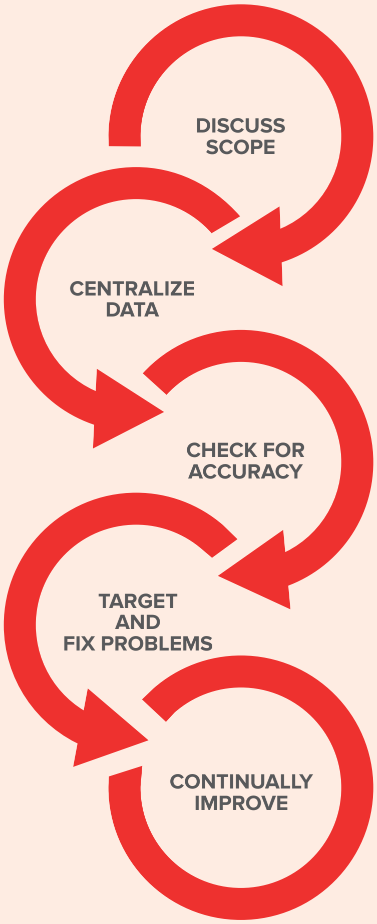

THE FUTURE OF DCIM

72 percent of data centre managers polled by Gartner responded that they would consider smaller DCIM manufacturers, especially if innovative solutions were offered.

Despite the challenges and obstacles laid out, DCIM has a bright future with research analysts predicting growth as high as 60 percent penetration within a year.

Although the DCIM market has some large players in the space, there are a number of smaller manufacturers bringing innovative ideas and solutions to an eager customer base. In fact, 72 percent of data centre managers polled by Gartner responded that they would consider smaller DCIM manufacturers, especially if innovative solutions were offered.

As DCIM continues to mature, new functionality and capabilities will inevitably come to the forefront. Some of the more promising innovations – some of which are already available in one form or another – are in the following areas:

› **Automated asset location**
Some systems now offer automated systems and RFID tagging, eliminating manual location tracking, with new innovation likely to provide more advanced options in the future.

› **Asset auto-discovery and change management**
Advances that allow for detailed information about an asset to be captured automatically, such as server configuration data, are continuing to become more prevalent.

› **Mobile applications and touch-based technology**
Many DCIM manufacturers have already begun to offer tools adapted to these platforms, allowing for on-the-go monitoring and more.

› **Integration with other data centre management tools**
More and more DCIM manufacturers are beginning to open their system up to outside integration, allowing for the possibility of greater and greater options and capabilities.

› **Control loops**
The evolution of DCIM is moving towards automated, closed-looped control systems, helping teams move from being reactive to being proactive when resolving issues.

› **“What if” scenarios**
Planning tools that run potential scenarios help to analyse the impact of new equipment, technology refreshes, equipment failures and more, will be a growing innovation within tomorrow’s DCIM solutions.

CONCLUSION

As the complexities of running and managing a data centre infrastructure continue to grow, so do the vast opportunities to enhance and improve the overall data centre environment with additional efficiency.

As a relatively new technology sphere, DCIM is subject to a lack of standards and contradictory definitions which can result in varied expectations, so is critical to define and set your own companywide criteria and success metrics for a DCIM solution. Once a DCIM path is defined and implemented, the enhancement to current capabilities and future potential can far exceed the alternative of continuing with disparate and inconsistent data centre management that is not aligned with company objectives.

The importance of information management, including tips on how to ready an organisation for DCIM as well as some insight to the selection process is included in Anixter's five senses of DCIM, which are crucial for the core functionality and value that it provides.

DCIM can help data centre managers run their facilities more effectively and efficiently, providing the process of selecting and evaluating a solution that adheres to many of the principles and guidelines outlined in this report. The right DCIM solution should be adaptable to future data centre expansion, as well as a vital tool in helping achieve growth in the fastest possible way.

ANIXTER®

SPONSORED BY ANIXTER'S TECHNOLOGY ALLIANCE PARTNERS[SM]

Anixter's Technology Alliance Partners provide solutions designed to connect the world's most important systems. Our partners help organisations operate more efficiently and securely while maximising value.

AUSTIN HUGHES® AXIS® COMMUNICATIONS COMMSCOPE® CORNING CPI CHATSWORTH PRODUCTS

PANDUIT® Raritan® A brand of legrand® TRANSITION NETWORKS® VIAVI

GLOBAL REACH. LOCAL ADVANTAGE.

With Anixter, you get a true local partner around the world. No other distributor of our kind can claim an in-country presence in approximately 50 countries and in over 300 cities.

We do business in more than 35 currencies and 30 languages, which means we are uniquely positioned to help facilitate your project in the local environment, reduce risks and keep costs down.

YEAR FOUNDED 1957

APPROXIMATELY 8,700 EMPLOYEES

APPROXIMATELY 50 COUNTRIES

OVER 125,000 CUSTOMERS

APPROXIMATELY 320 WAREHOUSES/ BRANCHES

OVER $1.1 BILLION INVENTORY

IN OVER 300 CITIES

STOCK SYMBOL AXE

OVER 450,000 PRODUCTS

FORTUNE 500 COMPANY

About Anixter: anixter.com/aboutus/emea
Legal Statement: anixter.com/legalstatement

Anixter Inc. World Headquarters
2301 Patriot Boulevard
Glenview, Illinois 60026
1.224.521.8000

Anixter EMEA
Inspired, Easthampstead Road,
Bracknell, Berkshire
RG12 1YQ United Kingdom
+44 (0)1344 388100