The Future Data Centre

Trends in smart data centre design and construction

The Irish Advantage

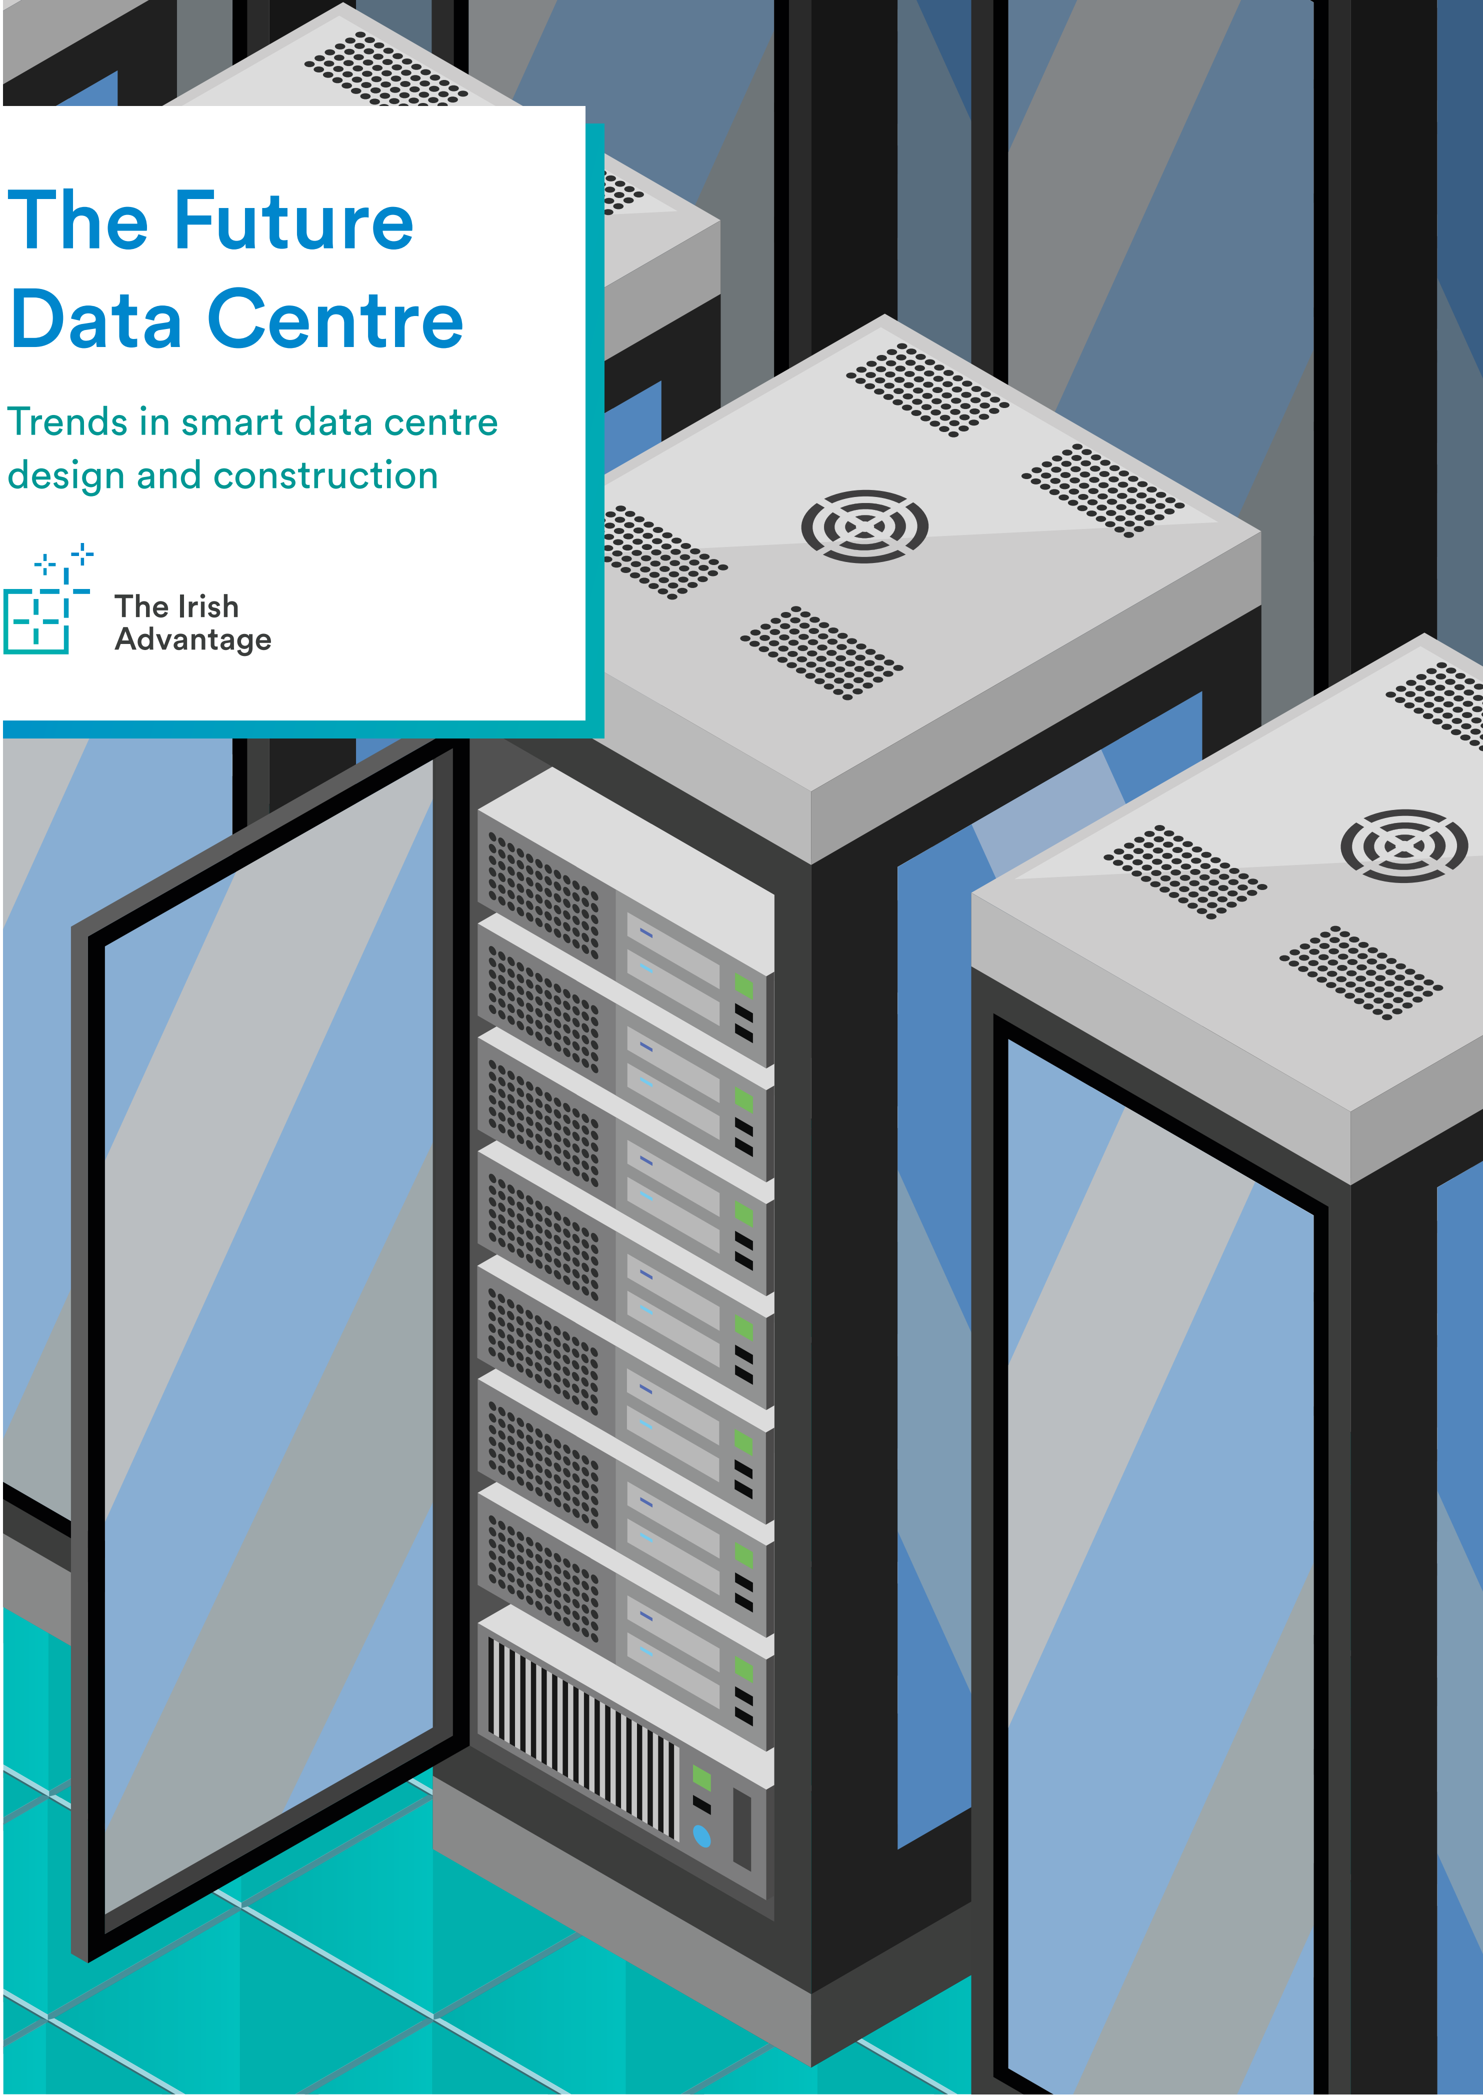

ENTERPRISE IRELAND

DCD

Contents

DatacenterDynamics (DCD) Comment

As with any construction process, the construction of a data centre is focused on the outcome and the blueprint set down by the design specifications. This paper indicates that the construction of data centres may be too outcome-focused with not enough attention paid to the process. This is particularly important given that the technological trends which are driving data centre change are creating facilities that are larger, denser, more networked, more converged and scalable. The increasing pace of technological change creates a need for environments where upgrade can be executed quickly.

This means decision processes which are more likely than in the 'legacy IT' era to be taken stage-by-stage, to be complex and to involve larger number of parties as the data centre is aligned with business objectives and as risk evaluation becomes more exhaustive. As such, data centres enable a range of consumer and business services. They have raised awareness of data centres among the public and legislators which together with the cost of resources puts an increased emphasis on energy efficiency, compliance and security.

While the base skills required for data centre construction will not change in the immediate future, and the data centre construction sector seems more advanced than the broader commercial construction sector in terms of best practice, the way forward will be marked by increased collaboration between the different parties and stages in what is sometimes a discontinuous build process and through the sharing of project knowledge specific to this form of construction.

Introduction

This white paper has been written on behalf of Enterprise Ireland in order to identify and analyse key trends in data centre procurement, design and construction.

The information presented in this report has been collected from a number of sources including interviews conducted specifically in relation to this project, together with analysis of data collected as part of the DCD annual Data Centre Census as well as surveys, media reports and analyses provided by independent experts. The DCD Global Census has collected over 2,000 responses annually from end user organisations and 500 from colocation, cloud and hosting providers.

Note: Quotes from the interviews conducted are shown through the body of the report in green.

Executive Summary

The Data Centre Landscape

The current data centre world is characterised by a number of factors, some well-established, some recent. The first is expense – data centres have always been expensive buildings to design, construct and run. Even before the moment when power reaches a rack for the first time, there will be considerable investment. Operational costs are high in comparison to most other commercial buildings, driven largely by the very high power consumption of the data centre.

There are also two major trends that are shaping the data centre market globally. The first is the move from in-house data centres into facilities and services provided outside the organisation ('outsourcing'). 'Outsourcing' can be broadly defined as the process of sourcing facilities, facility services, IT capacity or IT services from an external provider.

The second related trend is away from 'legacy' IT (physical servers) towards non-physical, 'dematerialised' infrastructure and components. This includes components and services that are created within and sourced from virtualised, cloud-based or software-defined environments.

These trends have created the requirement for larger facilities which can operate on the basis of the provision of standardised delivery models and services on a utility basis. Since the business model of cloud provision is based largely on volume and efficiency of delivery, the 'hyperscale' data centre has emerged. As many organisations with medium and smaller data centres have shut their own smaller facilities and migrated their workloads into cloud and colocation, so the profile of the world's data centres will continue to move towards larger facilities.

The costs of building a data centre today need to be set in the context of aversion to risk as data becomes an organisation's most important asset and needs to always be available and totally protected. The data centre is also subject to frequent technological change. Change is led by the evolution of the processing, storage and networking equipment and systems that are housed in the data centre and which represent the core purpose of its construction. Data centres may have a functional life of up to 20 years which, at an average rate of IT renewal of around three years, may mean six or seven upgrades.

The major responses of data centres to IT transformation over the past decade can be summarised as growing larger, denser, more focused on efficiency and demand, operated increasingly through software and software definition to meet a greater volume and diversity of IT workloads.

One of the enablers of demand/utility-based IT has been the development of modular units which can be built progressively into facilities as demand increases. While the debate between the benefits of modular versus more traditional construction methods has run in the industry for most of the past decade, modular will take an increasing share of global data centre build business, largely on the back of the previously mentioned mega- and hyper-facilities where staged construction programmes mean that capital costs can be matched more closely to incoming demand.

There are a number of further technological trends – the importance of networking internally and externally to the operation of colocation and cloud facilities, the evolution of converged and hyper-converged infrastructure to create a unified system that is more efficient to manage, thereby reducing TCO (total cost of ownership).

Changes to the design and build process

The design and build of data centres need to meet an increasing number of requirements to achieve an increasingly higher standard. The reasons given for investment in a data centre indicate that there is not one reason for data centre investment but a range of possible drivers such as the (perhaps obvious) need to increase IT capacity, to reduce operating costs, to improve security or simply to replace an old data centre.

These varied investment drivers mean the process of design and construction has to be smarter. The key changes to process include those concerning standards and legislative requirements, the technologies now available to facilitate key data centre functions, the exemplar provided by global cloud and colocation providers and radically changing decision making and procurement processes.

Standards and legislative requirements

As with any design and construction project, the process of designing and building a data centre needs to conform to local legislative requirements in terms of land use and planning permits, utility access and provisioning, environmental considerations, building codes as well as health, safety and labour requirements.

There are further requirements, some mandatory and some voluntary, when designing and building a data centre. These may include legislation designed to reduce energy consumption or wastage, industry standards to evaluate the levels of redundancy for which the data centre is designed and at which it is operated (such as TIA-942, the Uptime Institute Classification System and CENELEC EN 50600). The evolution of data centres that house data from across a wide geographic region as well as increasing concern about the privacy implications of data traffic have led to a raft of agreements and legislation to shore up data privacy and sovereignty practices. Further standards apply to security of data, both physical and cyber.

In terms of trends in legislation and standards, it is likely that both risk and energy consumption will become more controlled by both legislation and standards. Given the (almost) global political commitment to reducing carbon emissions, the challenge in most regions of the world is to maintain ICT growth while limiting carbon emissions from increased energy consumption.

What can (or should) be learned from hyperscale facilities?

To operate a viable business model, the hyperscale facility has changed the design, construction and fit-out practices that existed during the 'legacy' era. This is not surprising considering the scale of transactions that the companies that operate these facilities deal with. The business model of these companies is based on the capability of digital infrastructure to disrupt legacy business models. To achieve this requires facilities that can deliver enormous economies of scale, that are scalable and able to cope with peaks and troughs in demand, that are highly networked and sit at the centre of huge webs of connected users and devices, and which are able to analyse and model the massive amounts of data flowing through their systems to create insights into their users that can be used to generate further income streams.

So, what are the innovations that these companies have brought to the design and build of data centres?

In terms of location, many of these companies choose sites close to sources of renewable power and where the climate is cool enough to reduce the need for electrically-powered cooling. Some have also made equity investments in clean energy generation as a means of offsetting their carbon footprint.

The process of constructing these huge data centres is highly commoditised – the IT end facility package is viewed as a standard product, split out from the building of the facility. The process is modularised whereby common commercial module/parts are deployed in a standardised and flexible configuration, which will also support online expansion. Most components can be manufactured offsite and can be assembled, disassembled, renovated, moved easily to the site and replaced by 'hot-pluggable' components.

While information on the exact global size of these organisations is hard to come by, some do share information on their design and build practices. In 2011, Facebook launched its 'Open Compute' project. The project was intended initially to share information on hardware performance, to increase data centre efficiency and sustainability and the desire to 'demystify' data centres. This process has continued and expanded since, among global hyperscale companies.

There is no technological reason why smaller data centres cannot adopt the configurations and principles used by the major cloud providers and many have adopted software definition, Open Source data centre fabrics and migrated to faster network speeds. The issue is not one of technology but one of ROI – for smaller data centres, the costs of refreshing an on-prem data centre environment would probably not be justified, particularly if the benefits of such environments can be accessed through outsourcing.

The one set of facilities that stand to gain the most from adopting the technological innovations of the major cloud players are also those which stand to be most threatened by the continuing growth of such cloud players. These are colocation providers whose enterprise client base has migrated into cloud and who may only partially be able to redress the loss of that revenue stream through attracting

cloud and managed service providers. They have needed to evolve a business model based on connectivity inside and outside the facility and an interlinked ecosystem with internal communities and dense, modular, converged systems to meet the need for variable and scalable loads.

Changing technologies

As with any data centre of any criticality, issues of facility infrastructure design are dominated by power, cooling and resilience. Getting power and cooling right is necessary as the implications of not doing so carry threats at both extremes – under-provisioning will run the risk of shutdown while over-provisioning will negatively impact the bottom line. Energy efficiency is also assuming a new importance, so the legacy principle of resilience at any cost is slowly being abandoned and is a key design trend.

As server densities rise, the amount of heat generated is also increasing. New processors generate more than five times the heat of older processors, and new servers and switches could generate up to ten times the heat per square foot as those from ten years ago. More recently, data centres have started to explore a number of new cooling technologies and architectures to add further efficiency and cope with increasing rack densities. These include the introduction of vertical exhaust ducts, heat wheels, various close-coupled cooling approaches as well as chipset and server-based solutions such as processors that generate less heat, increasing the highest temperature at which they work reliably, improved heat transfer through changing the layout of components

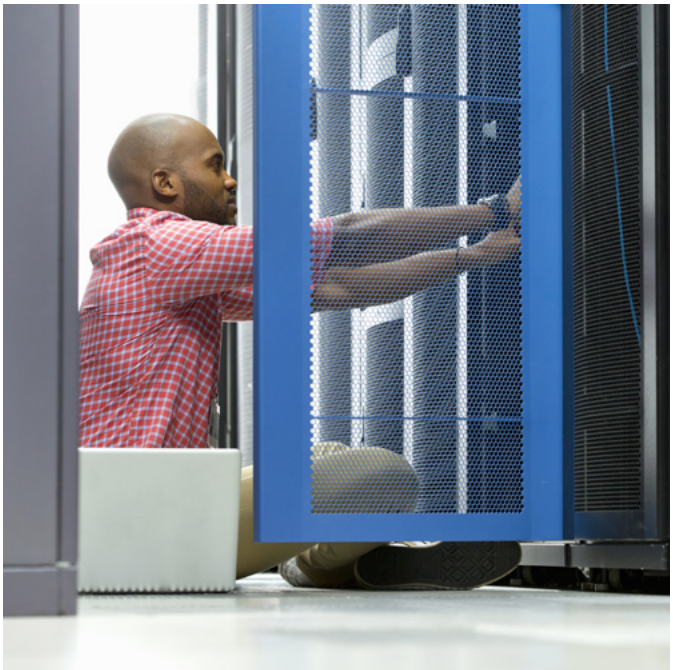

within the chassis and developing solutions which immerse the components in coolant.

Emphasis on greener and power-efficient technologies will only grow as part of a wider accountability trend based on concerns about power, emissions, excessive water consumption, noise and other kinds of environmental impact. Data centres already account for a substantial percentage of global energy consumption, and as digital media, Internet of Things (IoT) and Big Data requirements ramp up, their energy demands will only grow. Consumers and political bodies may simply intensify their scrutiny as this happens.

In terms of energy efficient power generation options, only a minority of organisations are looking at renewable or decentralised power generation options, most usually battery storage solutions. There is increased interest in cogeneration/combined heat and power solutions whereby data centres can use the ancillary benefits provided by their heat management solution to defray the environmental cost of the energy they use through the re-use of heat rejected from the data centre for heating local homes or community buildings.

As colocation facilities (or facilities that emerged from colocation) play a greater role in hosting and delivering enterprise IT, this increases the need for monitoring and management across the portfolio. Data centre and IT service providers are looking at operations that are intelligent and flexible, whereby clients can be charged on the basis of a number of business models. This means that a systematised and data-driven method of operation is required across different locations and different environments.

A majority of data centres across the world use management systems including Data Centre Infrastructure Management (DCIM) to manage their operations. The difference between the majority of systems and those in the process of development by global cloud and colocation operators is the degree to which machine learning and intelligence overlay their facility management.

This growth in space accounted for by colocation, hosting, IT service and cloud providers will also change the security focus among enterprise organisations that were previously guarding their data on-prem and who are now trusting it to an external provider. This includes focusing on the threat of disruption caused by malware, targeted Distributed Denial of Service (DDoS) and other electronic forms of assault as well as the threat of physical attacks on, or accidental damage to, premises and equipment.

Most published data indicates a wide range of financial impacts depending on the cause of the disruption, the company activities and workloads affected, the length of time of outage, the cost of restoration, and so on. The consequences may include loss of revenue, production and profitability, damage to brand and reputation, the costs of repair and restoration to equipment and systems, the impact on staff morale, future legal and insurance costs.

The future data centre of build

The success of a data centre build is judged on two main criteria – the first is the outcome – that the end data centre can meet the business objectives for which it was commissioned – and the second is the process needed to enable the outcome.

The means of building a data centre is changing to reflect the increased importance to the business and the technology available to do so. More people are involved. The number of parties (whether people or departments) involved in major data centre decisions (build, extension, major refit/refresh) has increased for an average range of 7-10 per project (2007 research) to 15-20 in 2016.

The process of decision making and procurement for data centre build has become more thorough and more accountable as data centres become more mission-critical and therefore more risk-averse. One consequence is that procurement has become more formalised, less reliant on open tender and more involved, down to items of lower value.

The people involved in decision making have also changed. Broadly, involvement in decisions now reaches far more widely across the organisation and is based on skill set and capability and will include external specialists within specially constituted project teams.

Partly as a result of the increased capabilities offered by major global suppliers, there is a trend towards a single provider of facility components, including enclosures, power distribution and protection, cooling, cabling and monitoring, rather than relying on different specialist providers.

In terms of design and build, most projects researched changed some of their design parameters as they progressed, normally as the result of different areas of expertise being introduced. This means that flexibility needs to be built into the process, usually around key 'milestone' decision points at which new contracts will be tendered or agreed.

So what is the future for data centre designers and builders? The construction industry overall has some major responsibilities globally since it accounts for around 7% of GDP (more in emerging markets with higher infrastructure requirements). It is also the sector that consumes the greatest amount of raw materials, and buildings are the largest source of carbon emission.

The construction of data centres will follow some projected overall construction sector trends but not all of them. A key difference between many construction projects and the data centre is the balance between the 'upfront' costs of construction and subsequent operational costs. While there is no such thing as an average data centre or an average commercial construction, the very high operating costs of the data centre mean that the construction sector will need to focus on 'whole of life' not just the initial construction costs.

A feature of the immediate future shared by the construction and engineering sectors is that of skills shortage. The data centre sector shares this, particularly outside the original IT and facility silos. Shortages have been measured in the skills required for digital transformation, in data analytics, in design and architecture, in coordinating between technology and business, in hybrid IT processes and in most other places where cloud and legacy intersect. The construction industry similarly will face skill shortages although these will continue to be localised and linked to the extent of construction activity in particular markets.

While the construction industry overall will focus increasingly on environmental performance in terms of materials, methods of construction and local environmental impact, data centre designers and constructors have for some time used some of the major technologies to reduce environmental impact by means such as modular construction and building components offsite to ship and fit locally. The enormous impact of the power bill on the TCO of a data centre means that ways of improving energy efficiency are designed in, and many large data centres are linked to renewable or sustainable power sources or, where possible, recycle heated air to warm local buildings. Thus the demands of the data centre industry have already prepared its design and construction sector well for the challenges that other parts of that industry are beginning to feel.

Section One

Building for Change

The costs of building a data centre today need to be set in the context of aversion to risk as data becomes an organisation's most important asset, and of frequent technological change.

Data centres are expensive buildings to design, construct and run. Even before power reaches a rack for the first time, there will be considerable investment. Actual costs per square metre vary according to location, the scale of the facility, the level of resilience and the required support investment into utilities and networks but range broadly from US$3,000/square metre to US$18,000/square metre. This means that, at the upper end, the data centre costs more per square metre than the estimated, if very hypothetical cost of rebuilding the Palace of Versailles today (US$7,500 to US$10,000/square metre) or the cost of building the world's largest occupied palace in Brunei (quoted at US$14,500 to US$16,500/square metre).

The costs of building a data centre today need to be set in the context of ever-increasing aversion to risk as data becomes an organisation's most important asset, and of frequent technological change. Change in the data centre is led by the evolution of the processing, storage and networking equipment and systems that are housed in the data centre and which represent the core purpose of its construction. Data centres may have a functional life of up to 20 years which, at an average rate of IT renewal of around three years, may mean six or seven upgrades.

The major responses of data centres to the evolution in IT equipment over the past decade can be summarised as growing larger, denser, more focused on efficiency and demand, operated increasingly through software and software definition to meet a greater volume and diversity of IT workloads. For each IT delivery trend, the design and construction of the data centre needs to adapt, for even as virtualisation and cloud take a greater role in the IT landscape, there will be a data centre housing the servers and support needed to offer those environments. Those shown on Figure 1 are intended as a very broad summary of some of the key responses.

There are two major trends that are shaping the data centre market globally. The first is the move from in-house data centres into facilities and services provided outside the organisation ('outsourcing'). The second related trend is away from 'legacy' IT (physical servers) towards nonphysical, 'dematerialised' infrastructure and components. This includes components and services that are created within and sourced from virtualised, cloud-based or software-defined environments.

"The future is clearly cloud for a lot of businesses but because we run a lot of infrastructure as part of our core business it is still relatively relevant to have on-prem systems for us. I think that this will gradually be replaced by cloud and colocation as we refine what services we want to keep in-house and which we want to place elsewhere. The biggest trend we are seeing is the ability to move services to Software as a Service (SaaS) delivery and this is typically our first option for a lot of small systems deployment for specialist roles. We have moved email and collaboration to SaaS delivery as this is more readily accessible to more of our workforce."

The practice of investing in outsourcing has grown prolifically over the past five years, and it will continue to do so.

Outsourcing providers include companies that provide data centres and/or data centre services such as colocation, cloud services and managed hosting to end users on a commercial basis, offering them an alternative to building and/or operating their own data centre. The practice of investing in outsourcing has grown prolifically over the past five years, and it will continue to do so. In the time period 2012 to 2020, investment in data centre outsourcing is expected to grow from US$25 billion to US$58 billion, at more than double the rate of investment in data centre facilities which increased from an initial amount of US$70 billion at 4.7% CAGR to US$101 billion. The latter investments include those made by colocation, cloud and other IT service providers in building the physical infrastructure necessary to deliver their services.

The proportion of IT infrastructure that is outsourced will increase at an average rate of 2.4% each year to 2020 while the rate at which IT infrastructure is transferred from ‘legacy’ into cloud environments will increase at an average rate of 3.6% each year. These figures indicate a gradual transition into both outsourcing and digitised infrastructure.

Figure 1:
Broad IT Delivery Trends and Design Responses

IT delivery trend	Facility design response	Drivers
Cloud	Increased size, density, efficiency, reduced latency	Business model based on scalability, margin based on control of operation and costs
Colocation, hosting, managed services	Increased size, density, efficiency; zonal structure/variable loads	As above and need to adapt to/compete against cloud
Miniaturisation (greater computing capacity in a smaller space)	Increased density, provisioning	IT delivery, workload and app requirements
Connection/ interconnection/network	Locational	Necessary for increasing level of internal and external data traffic
Virtualisation, software definition	Increased density, provisioning	More efficient form of infrastructure, more direct access into cloud
Demand-based compute	Modularity	Add/build infrastructure as required
Integration of processing, storage, networks	(Hyper) convergence	More integrated form of infrastructure, more direct access into cloud/virtualised environments

[Source: DCD 2017]

Figure 2:
The Growth of Outsourcing and Cloud

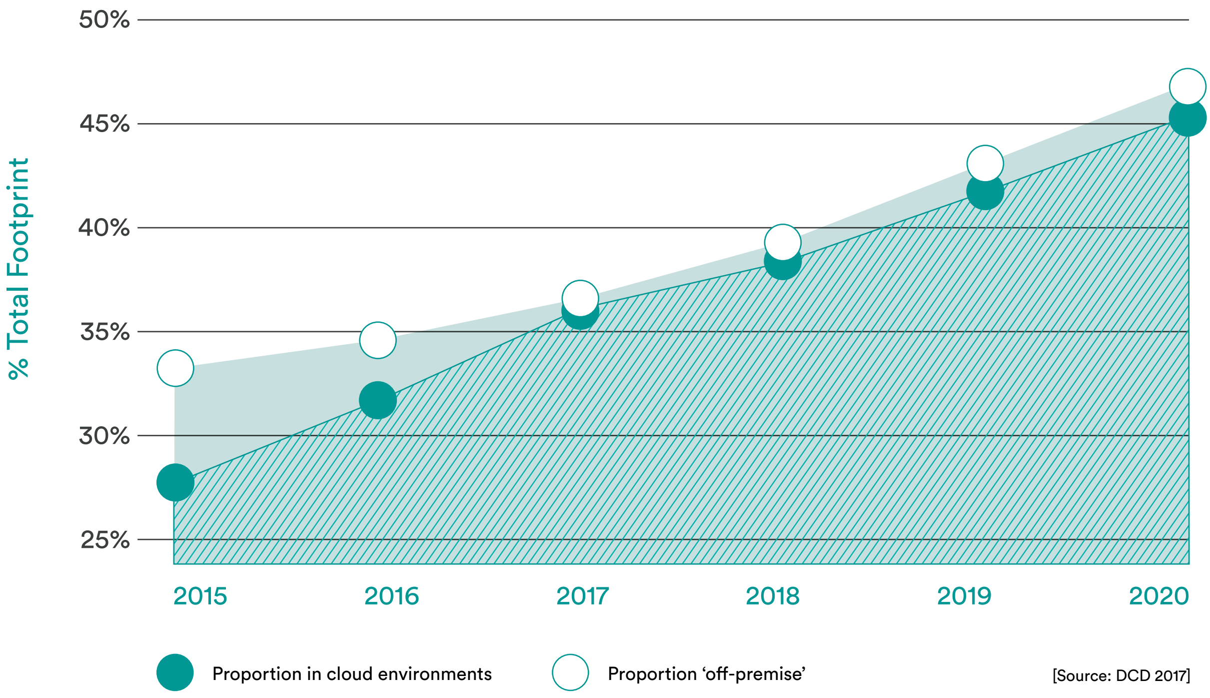

Figure 3:
The Growth of Large Data Centres (Aggregated GW Power Capacity)

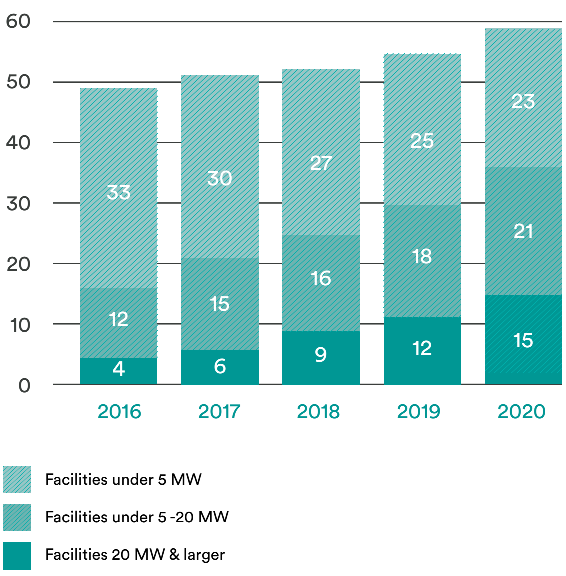

[Source: DCD 2017]

Figure 2: These trends have created the requirement for larger facilities which can operate on the basis of the provision of standardised IT delivery models and services on a utility basis. Since the business model of cloud provision is based largely on volume and efficiency of delivery so the 'hyperscale' data centre has emerged. As many organisations with medium and smaller data centres have migrated their workloads into cloud and colocation, so the profile of the world's data centres will continue to move towards larger facilities.

Figure 3: The increase in the power capacity held by large data centres is not just a function of increasing size but of increasing server densities as well. Over the thirteen years between 2001 and 2014, the space, power and load profile needed to execute the same amount of IT work have reduced by a factor of 2000 (for cabinet numbers), by a factor of over 900 for space and by over 800 for IT load power requirement. These figures indicate that the principle of Moore's Law (a doubling of the number of transistors in a dense integrated circuit every 24 months) has been exceeded and that trends in data

Figure 4:
Modular Represents an Increasing Proportion of Data Centre Investment

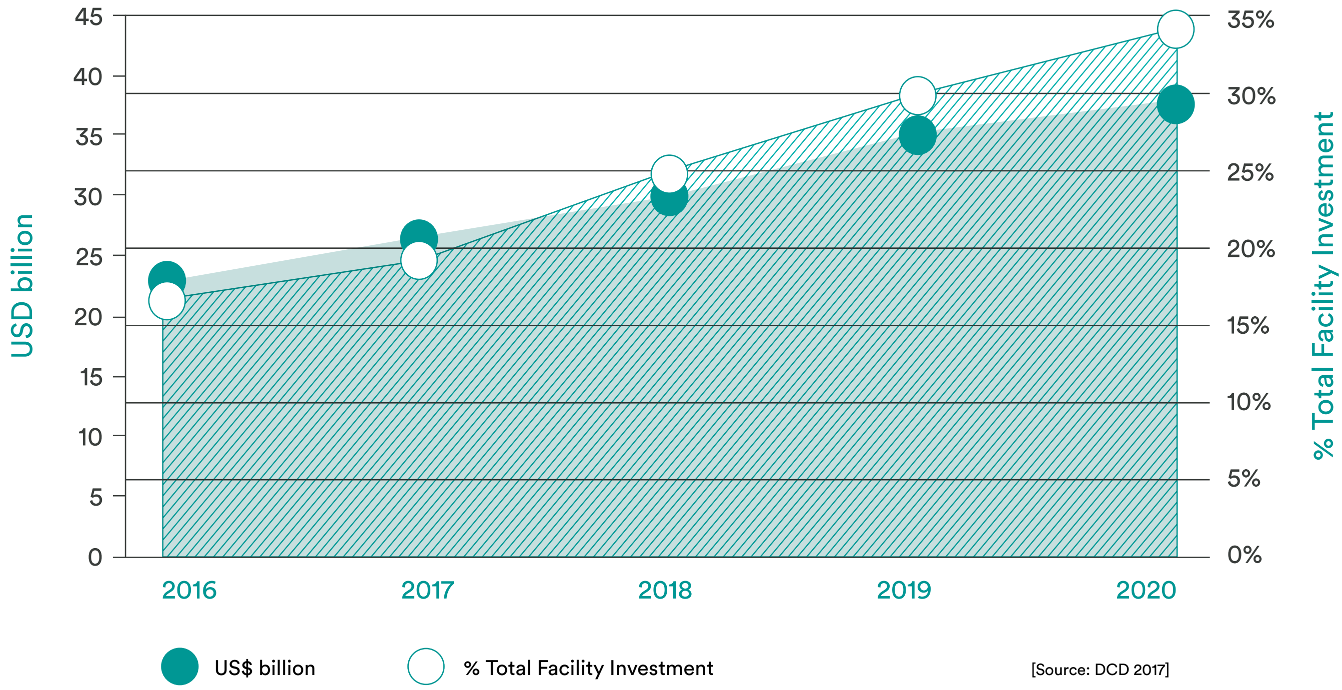

[Source: DCD 2017]

generation and traffic will continue exponentially. Demand continues to move faster than this increased capacity to supply, following the 'Jevons Paradox' whereby making supply more efficient does not reduce consumption but increases it.

Figure 4: One of the means for constructing a facility that grows with IT demand has been the development of modular units which can be built progressively into facilities as demand increases. While the debate between the benefits of modular versus more traditional construction methods has run in the industry for most of the past decade, modular will take an increasing share of global data centre build business, largely on the back of the previously mentioned mega- and hyper-facilities where staged construction programmes mean that capital costs can be matched more closely to incoming demand.

There are a number of further technological trends – the importance of networking internally and externally to the operation of colocation and cloud facilities, the evolution of converged and hyper-converged infrastructure to create a unified system that is more efficient to manage thereby reducing TCO (total cost of ownership).

The evolution of hosting, processing, network, storage and intelligence technologies together with the major move to cloud and towards edge processing means that the data centre needs to be designed across a number of generations of technology, some of which may make unpredictable demands on the physical constraints of the data centre.

Section Two

Changes to the Design Process

The design and build of data centres needs to meet an increasing number of requirements to an increasingly higher standard. The list of reasons given in research for investment in a data centre indicates that there is not one reason for data centre investment but a range of possible drivers most of which have increased in prominence between 2015 and 2016 (Figure 6).

To meet these investment drivers means the process of design and construction has to be smarter. The key changes to process include those concerning:

- Standards and legislative requirements
- Technologies for key data centre functions
- The exemplar provided by global cloud and colocation providers
- Decision making and procurement processes (discussed in the final section)

Figure 6:
Drivers for Data Centre Investment

Investment	2015	2016
Increased IT capacity requirements	43.5%	51.8%
To reduce operating costs	39.8%	51.6%
To improve security	33.5%	35.1%
End of facility life	32.3%	31.5%
To be ‘greener’ & more sustainable	30.5%	40.2%
To increase competitive differentiation	30.5%	38.2%
To increase power into facility	30.3%	40.6%
To improve network performance	29.7%	32.1%
To enable virtualisation/cloud computing, service development	29.6%	44.3%
To increase redundancy	29.3%	28.9%
To improve space use	27.0%	41.1%
Changing corporate & client requirements	26.5%	40.0%
To attract different client groups	25.9%	37.8%
To meet legislative or accreditation requirements	23.5%	27.5%
To access software-defined utilities & As-A-Service delivery models	22.8%	34.8%
To support the requirements of big data	16.5%	30.1%
Other reasons	8.8%	12.1%

[Source: DCD 2017]

Section Three

Standards & Legislative Requirements

As with any design and construction project, the process of designing and building a data centre needs to conform to local legislative requirements in terms of land use and planning permits, utility access and provisioning, environmental considerations, building codes as well as health, safety and labour requirements.

There are further requirements, some mandatory and some voluntary, when designing and building a data centre:

- As buildings that use considerably more energy than commercial office space, data centres may fall under legislation designed to reduce energy consumption, or specific steps may need to be taken in design and operation to avoid penalties for what is judged to be 'excessive' energy consumption. Increasingly, the consumption of energy from renewable and sustainable sources will be part of legislated targets. A range of metrics have been developed by the data centre industry to measure energy efficiency, utilisation and carbon emission which are quoted at the commissioning of a project.

- Since data centres are buildings which are required to offer continuous service availability, industry standards have also been developed to evaluate the levels of risk for which the data centre is designed and at which it is operated. Further standards have been established for best operational practice in terms of, for example, the temperature and humidity tolerances of IT equipment.

- The evolution of data centres that house data from across a wide geographic region, as well as increasing concern about the privacy implications of data traffic, have led to a raft of agreements and legislation to shore up data privacy and sovereignty practices by, for example, requiring that data is stored and processed within national borders. Since these agreements and laws impact most strongly on the management of cloud and networks rather than the physical construction of the data centre, they are not covered in this document.

The first set of data centre classification schemes focus on rating the reliability and redundancy of the data centre:

TIA-942

In the US, the Telecommunications Industry Association (accredited by the American National Standards Institute) defined the TIA-942: Data Centre Standards Overview in 2010. Like the Uptime Institute and CENELEC, TIA-942 specifies a four-tier classification for data centre infrastructure.

Figure 7:
Some Features of Different Tier Levels

Descriptor	Tier I	Tier II	Tier III	Tier IV
Redundant Components	N	N + 1	N + 1	2(N + 1)
Power & Cooling Delivery Paths	1 active	1 active	1 active + 1 passive	2 active
Single Points of Failure	Many + human error	Many + human error	Some + human error	None + human error
Site Availability	99.671%	99.749%	99.982%	99.995%
Site-caused Downtime per annum (actual)	28.8 hours	22 hours	1.6 hours	26.3 minutes
Power Outage Protection	NA	NA	72 hours	96 hours
Year First Deployed	1965	1970	1985	1995

[Source: DCD 2017 from various]

Uptime Institute Tier Classification System

The Uptime Institute's Tier Classification System first created in 1995 has become widely referenced in data centre construction and design. Tiers I-IV certification focuses on infrastructure performance, uptime, fault tolerance and disaster recovery, but a later standard dedicated to Operational Sustainability (2010) encourages consideration of specific risk exposures around site location – seismic, extreme weather, flooding, adjacent property uses, and physical security – either as a corporate policy or warranted by immediate surroundings.

CENELEC EN 50600

More recently, the European-based CENELEC EN 50600 data centre infrastructure series of standards have been formulated to deal with elements of site selection, construction, configuration and quality of construction. These will impact on data centres and buildings that house data centres. The series is separated into practices in terms of building construction, power distribution, environmental control, telecommunications cabling and security systems, as well as the management and operational systems that oversee all of these. It specifies a classification based around risks including availability, security and energy efficiency. Four classes have been developed running from Class 1 where there is no resilience and highest through to Class 4 which represents the lowest risk (but highest cost) proposition.

While these systems have been developed by established industry organisations and backed by official processes of inspection and accreditation, they are viewed also as generic indicators of data centre design and operation. The numbers of facilities that have been officially accredited by the Uptime Institute is far fewer than those stating in research that their facilities have been designed or operated to a certain standard. Tier III+ and, less commonly, Tier II+ have emerged to describe states of redundancy between the two Uptime Institute levels.

Figure 8:
Some Features of EN 50600 Standards

Infrastructure Component	Impact of Short Failure	Administrative Risk	Class 1		Class 2	Class 3	Class 4
Power Supply and Distribution EN 50600-2-2	High	Medium	No UPS	No Resilience (Single Path)	≥ N+1	Concurrently Maintainable (Active/Passive)	Fault Tolerant 2(N+1)
Environmental Control EN 50600-2-3	Low	Medium	No ECS	No Resilience (Single Path)	≥ N+1	Concurrently Maintainable (Active/Passive)	Fault Tolerant 2(N+1)
Telecoms Cabling EN 50600-2-4	Medium	High	No Resilience (Point to Point Cabling)		No Resilience (Fixed Cabling)	Concurrently Maintainable (Multi-Path)	Concurrently Maintainable (Diverse Routed)

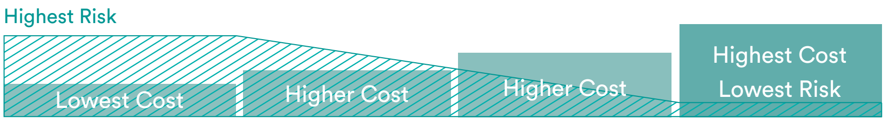

[Source: CENELEC]

More recently, a Tier V standard has been set down by Switch in the United States, adding a number of elements to established Tier models including:

- The capability to run 'forever without water'
- Protection against air pollutants
- Additional criteria for stored energy systems
- Critical distribution path and backbone resilience
- The available number of carrier services and the availability of D/DoS mitigation services
- Measures taken for the facility's physical security
- Protection against water intrusion
- Measures to ensure power sustainability and efficiency

It is too early to establish the extent to which these are intended to apply beyond Switch's Las Vegas 8 Data Centre but they have revitalised industry debate on the application of standards to the data centre.

A further trend in standards will reflect that not all IT loads within a data centre are of equal importance and that rather than taking a 'one-size-fits-all' approach, redundancy should be designed on the basis of prioritising loads or, in the case of multi-tenant data centres, in the willingness to pay extra for higher Tier levels.

Other data centre standards apply more specifically to issues of security and to energy efficiency:

ISO/IEC 27001

Introduced in 2005 and updated in 2013 by the International Organisation for Standardisation and the International Electrotechnical Commission ISO/IEC 27001 is a global voluntary certification framework for the legal, physical and technical controls applied to management systems within data centres and other IT infrastructure.

PCI-DSS

Dating back to 2004 the Global Payment Card Industry Data Security Standard (PCI-DSS) is a proprietary information security standard for organisations that process branded credit cards from the major card schemes including Visa, MasterCard, American Express, Discover, and JCB.

ISAE3402

The Statement on Standards for Attestation Engagements (SSAE 16) framework was defined by the American Institute of Certified Public Accounts (AICPA) in 2011. The audit leads to reporting under Service Organisational Control. SOC2 deals with security, confidentiality, privacy, availability and integrity. It requires that data centres operated by service organisations in the US audit their physical infrastructure and environmental safeguards.

BREEAM, CEEDA and LEED

Are accreditation schemes based on the improvement to efficient data centre operation? All have data centre-specific criteria which are largely design standards although BREEAM and LEED are based originally on specifications drawn up for a wider range of commercial buildings. National systems for rating and classifying buildings such as NABERS in Australia and Green Mark in Singapore are also common since individual markets will have particular national architectural challenges and climate-specific requirements.

In terms of trends in legislation and standards, it is likely that both risk and energy consumption will become more controlled by both of these methods. The pressure is getting stronger – non-compliance with the EU's General Data Protection Regulation can lead to fines of up to €20 million (around US$24 million) or 4% of global annual turnover. Simple math suggests that this may, in worst case scenarios, mean a shrinking company or closure.

Legislation to control data centre energy consumption has been debated in most markets where data centres have a footprint. Beijing, for example, issued a ban on the operation of data centres with a PUE less efficient than 1.5 within the city area. Legislation may be used also to control some of the other environmental impacts of data centres – the requirement to store diesel to fuel back-up generators has over the past years affected the operation of data centres both in Paris and in Virginia.

The context of ambitious climate change pledges made by the EU and across most other developed markets is likely to put pressure on governments to impose mandatory regulations in terms of energy efficiency and the reduction of carbon emissions. In the EU, energy efficiency legislation includes framework legislation for all the EU member states, as well as sectoral and procedural requirements and various product legislations. The framework indicates important changes in the legislation on both the EU and national levels, which are crucial for data centre operators, end users and technology providers with implications for their medium and long-term business strategies.

Methods of driving data centres towards improved energy efficiency have taken a number of forms, with evidence of success and failure for each. Methods include incentive for improvement, penalties such as increased taxation for not doing so, or through the promotion of energy efficiency and education as to the means of achieving it.

The challenge in most regions of the world is to maintain ICT growth while limiting carbon emissions from increased energy consumption. The EU on the whole, agreed to increase the renewable energy market share to 27% by 2020 and to improve energy efficiency by 27% by 2020 (as part of the 2030 Climate and Energy Framework). The International Energy Agency therefore recommended taking industry-specific measures to monitor and encourage energy efficiency which, for the data centre industry, might mean the introduction of mandatory regulations.

The key energy efficiency performance indicators will be critical in deciding how future EU regulations and taxation will evolve. The EU Code of Conduct for data centres is a well established set of guidelines and is likely to be a cornerstone to which energy efficiency initiatives will be linked.

The operation in the UK of a CCA for commercial data centres indicates that offering tax incentives may be a more efficient way to generate interest in energy efficiency rather than the promotion of the EU Data Centre Code of Conduct as a voluntary guideline for best practices leading to energy savings. The data centre industry in Europe is unlikely to reduce its carbon emissions entirely through self-regulation.

Section Four

What Can (or Should) Be Learned from Hyperscale Facilities?

In order to operate to a viable business model, the hyperscale facility has changed the design, construction and fit-out practices that existed during the 'legacy' era.

This is not surprising considering the scale of transactions that the companies that operate these facilities deal with – four billion searches a day (Google), 500 million tweets (Twitter), 254 million orders in a day (Alibaba), 350 million new photos per day (Facebook), 235 million messages per day (Tencent QQ).

The business model of these companies is based on the capability of digital infrastructure to disrupt legacy business models. To achieve this requires facilities that can deliver enormous economies in operation, that are scalable and able to cope with peaks and troughs in demand, that are highly networked and sit at the centre of huge webs of connected users and devices, and which are able to analyse and model the massive amounts of data flowing through their systems to create insights about their users that can be sold on.

So, what are the innovations that these companies have brought to the design and build of data centres?

In terms of location, many of these companies choose sites close to sources of renewable power and where the climate is cool enough to reduce the need for electrically-powered cooling. Such locations include Northern Europe, Ireland and remote locations in North America. Google, Facebook, Amazon and Yandex have all adopted this strategy.

Google, as one example, has also made equity investments in clean energy projects as a means of offsetting its carbon footprint. It has also entered into long-term power purchase agreements (PPAs) for renewable energy, in order to guarantee a long-term source of clean energy at stable prices, increase the amount of clean energy it uses and that is available in the grid, and help enable the construction of new renewable energy facilities.

The process of constructing these huge data centres is highly commoditised – the IT end facility package is viewed as a standard product, split out from the building of the facility. The process is modularised whereby common commercial module/parts are deployed in a standardised and flexible configuration, which will also support online expansion. Most components can be manufactured offsite and can be assembled, dissembled, renovated, moved easily and replaced by hot-pluggable components.

While information on the global size of these organisations is hard to come by, some are more than willing to share information on their design and build practices. In 2011, Facebook launched its 'Open Compute' project. The project was intended initially to share information on hardware performance, to increase data centre efficiency and sustainability and the desire to 'demystify' data centres. It is claimed that the initiatives generated by the project have created 38% more efficiency in terms of energy usage and 24% cost savings in terms of Facebook's Prineville site.

The OC programme has now achieved advances in Compute Servers (the latest in 2016), in Rack and Power hardware (2014), Storage (2017), GPU (2017), Network Switch (2015) and Network (2016). It is estimated that opening up its design specifications to members of its Open Computer Project may have saved it upwards of US$2 billion. Facebook also now uses an Open Compute database server and flash for storage.

The Scorpio Project, a customised rack server research initiative was pioneered by China's three largest internet companies (Baidu, Alibaba and Tencent) in October 2011. It has gained further interest from enterprises that provide cloud computing services and solutions – including carriers and equipment manufacturers. The Scorpio Project has also attracted interest in overseas markets because it is based on jointly-led open source, rather than by a single company.

The first two versions of Scorpio were mainly driven by Baidu, Alibaba and Tencent, along with the participation of the China Academy of Information and Communication Technology, China Mobile, China Telecom, Inspur, Huawei, Dell and NVIDIA, and with Intel as the technical consultant. The third version has involved more enterprise companies in an effort to promote innovations through the whole industry chain. As the adoption and use of cloud computing gained momentum in China, traditional servers became insufficient to meet the internet giants' requirements for massive development of IT resources in a rapid and flexible manner. In order to better cope with the problems posed by traditional servers including their huge energy consumption, a long deployment cycle, and the complexities involved in their operation and maintenance, the use of integrated and customised servers has become increasingly widespread.

The scale of these facilities means they have deployed innovative power and cooling solutions including the use of thermal energy storage systems, which cools water at night when temperatures are lower, stores it in insulated tanks where it retains its temperature, and then pumps it throughout the facility to cool servers during the day. In addition to being more environmentally friendly, the system reduces costs because power rates are typically lower at night when demand is lower. As a variation on the above, the air used for heating and cooling is sourced from outdoors and recycled. The plenum is used to hold hot air which is then used to dehumidify the data centre in hot conditions. The air moves through the data centre by pressurisation. The need to control cooling in order to improve energy efficiency is one shared with major colocation facilities:

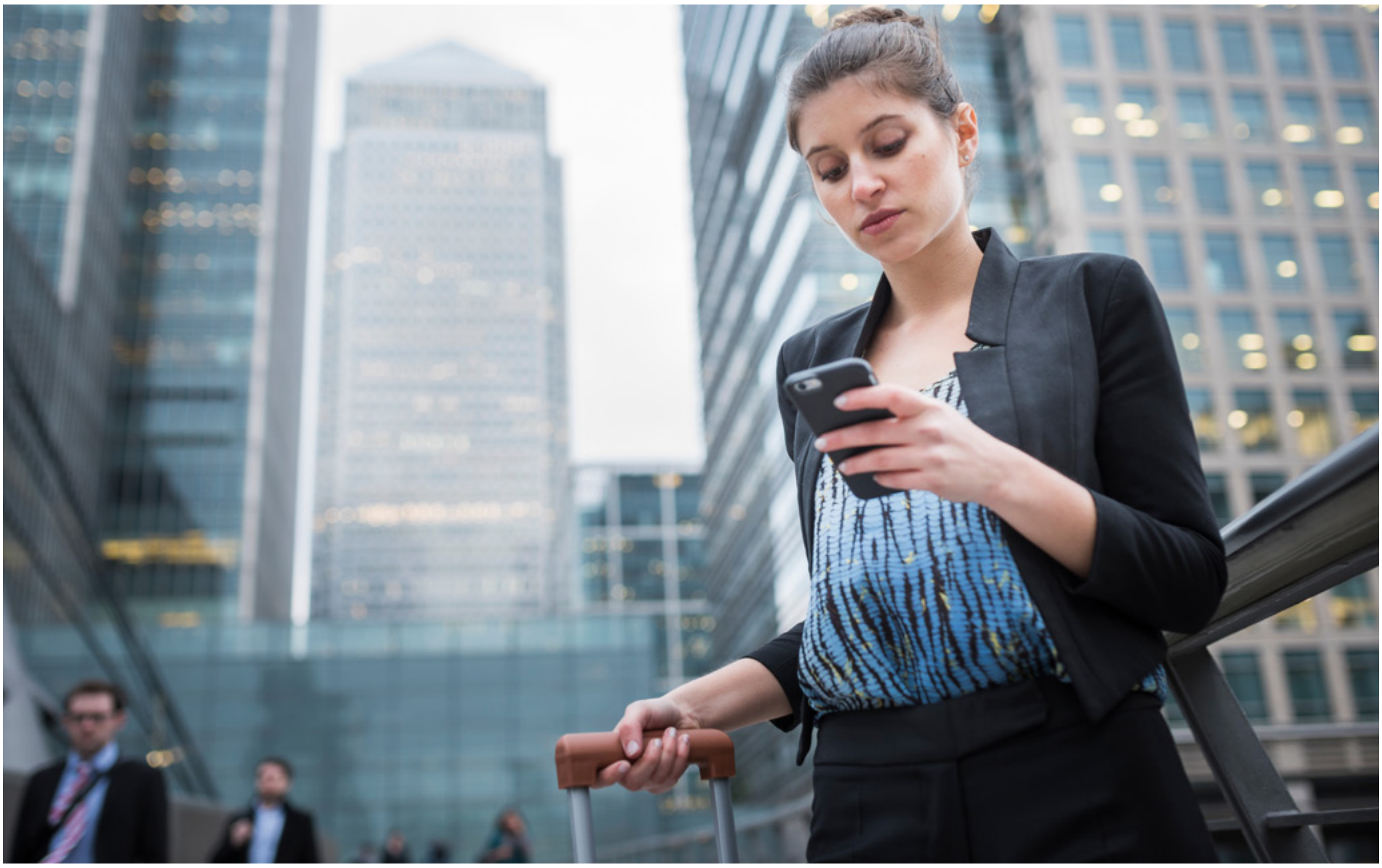

"PUE was 1.5 to 1.6, now we can get in the range 1.2 down to 1.15 based on true PUEs due to intelligent hot aisle containment and use of thermodynamics – evaporative cooling – we get to 1.12. That means 5kW goes to IT out of 5.5/6.0 overall. Previously we had to blow air under the floor. Outside it is 50°F at the moment. That's ideal but we need to warm the air up – take the natural air and mix it with heat from the DC we deliver at 74°F. We use overhead cooling so we drop the air into the space, the heat moves upwards towards the ceiling so we can pull it out – it reduces the energy used."

Power distribution and protection solutions also figure. Measures include putting a reserve bus in the power distribution system which switches when required between primary and reserve power flows. This cuts out the use of generators.

One component of the Open Compute data centre design is a high-efficiency electrical system. Facebook's Prineville, Oregon, location is the first implementation of these elements in a data centre. The facility utilises an electrical system with a 48VDC UPS system integrated with a 277VAC server power supply. Typical data centre power goes through transformers that lead to 11-17% loss of power before reaching the server. Facebook is using a configuration that is reported to lose only 2%.

The battery cabinet is a standalone independent cabinet that provides backup power at 48 volt DC nominal to a pair of triplet racks in the event of an AC outage in the data centre. The batteries are a sealed 12.5 volt DC nominal, high-rate discharge type with a ten year lifespan, commonly used in UPS systems, connected in a series of four elements for each group (called a string), for a nominal string voltage of 48VDC. There are five strings in parallel in the cabinet.

Traditional data centre components such as UPS devices may not be required by, for example, direct grid power to the server with HVDC backup. This increases energy efficiency. In-row cooling which shortens the distance air needs to travel also improves cooling efficiency.

Figure 9
The Facebook Data Centre Fabric

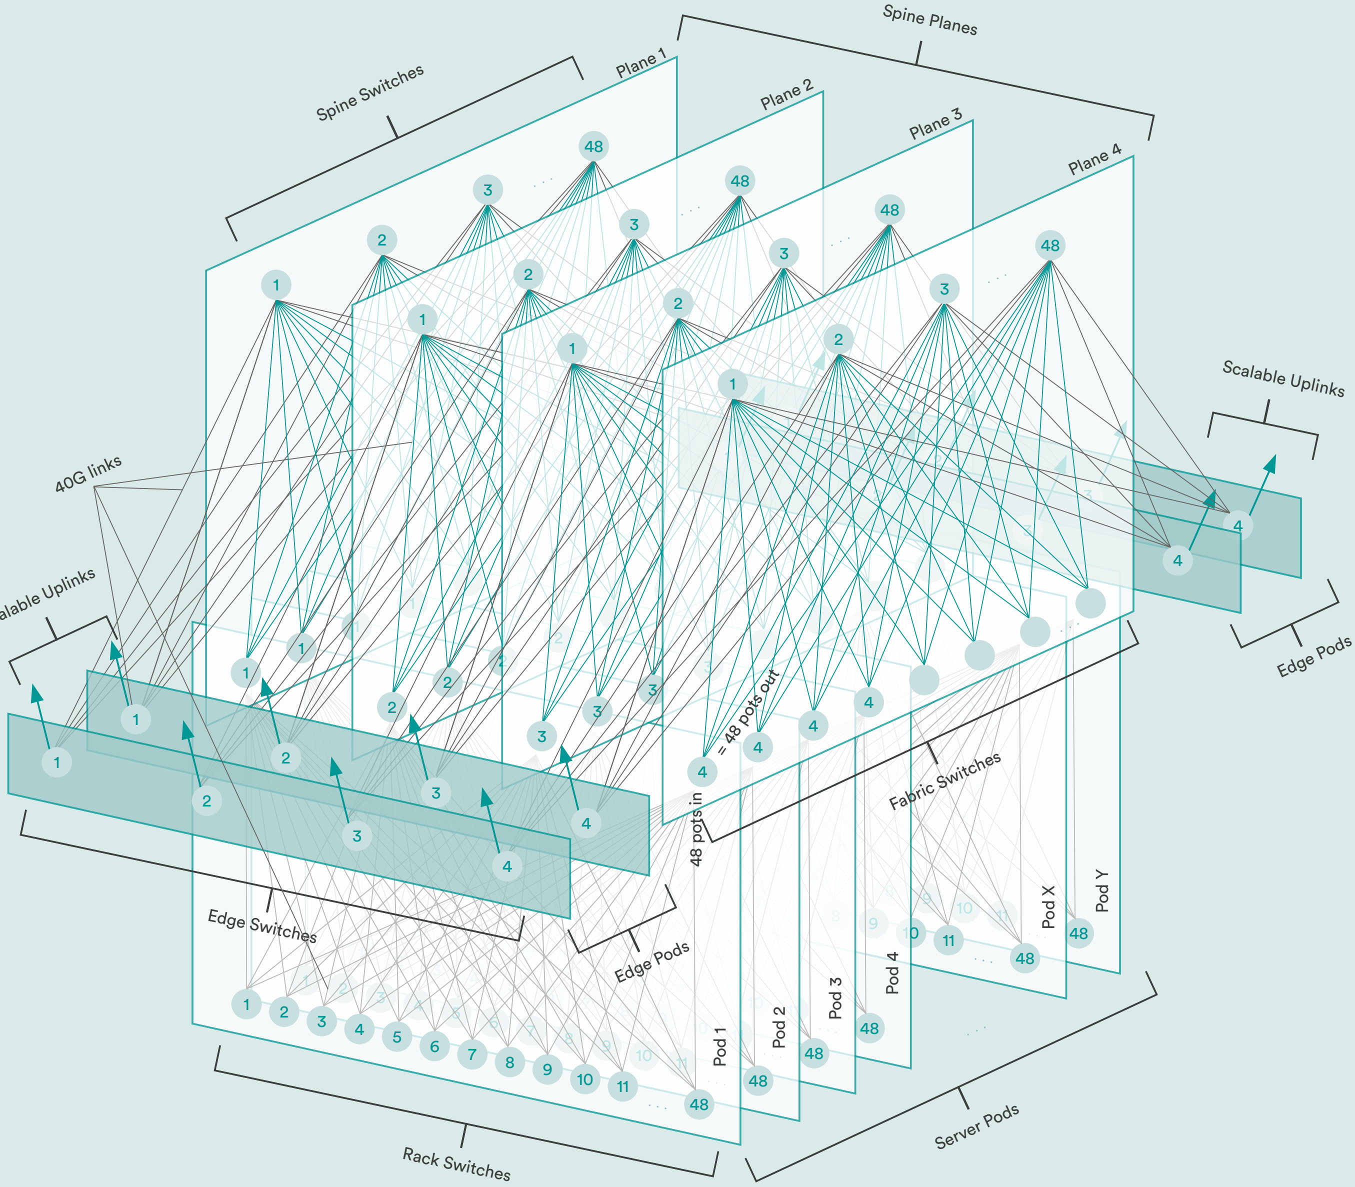

[Source: Facebook]

“Can you learn from the large cloud providers? I think design is done to death today – I think the next big thing is operational practice. Operational practice will become operational intelligence, the systems. These will become the next area and then this will move into preventative maintenance strategies.”

With regard to power supply, most of Tencent’s data centres do not use traditional UPSs but switch to the power supply model of “HVDC (50%) + Grid (50%)”. Tencent claims that, by adopting the new power supply model, the efficiency of the grid power can reach 100% as there is no conversion loss; and by using the energy-saving Sleep Mode, the efficiency of HVDC power can reach more than 94%, hence resulting in an overall power efficiency of 97%. Upgrades can also transition the system to one which uses the grid as the primary power source and HVDC as backup for achieving power efficiencies above 99%.

Baidu also uses a range of energy-saving methods and technologies, optimising power supply structure to reduce energy losses in conversion, using the mains supply and UPS simultaneously for power supply, enhancing supply/return water temperature to prolong free-cooling time, introducing variable frequency technology to data centre chillers and using CFD software for optimising server airflow. By doing all of this, Baidu can use natural air cooling 50% of the year and it is reported that its annual PUE is 1.37 on average and can be as low as 1.18 in optimum conditions.

In order to achieve scalability, the hyperscale data centre has moved away from the legacy hierarchical network topology to fabric network designs that offer greater capacity to deal with peak traffic, lower latency, higher redundancy in the event of switch breakdown and greater efficiency through, for example, allocating lower traffic items to less efficient areas of the data centre in order to reduce rotation.

Facebook’s data centre fabric used in its Altoona facility is a ‘pod and core’ configuration based around 48 rack clusters arranged in a leaf-spine topology through 40G downlinks and uplinks. The number of switches between fabric and spine creates a 3D structure where the extra dimension creates added layers of connectivity, redundancy and efficiency to the traffic process. It is easy to add capacity and edge pods allow scalability outside the core fabric.

For full operational efficiency, these data centres are automated and run by software defined hardware. Google uses machine learning to measure its data centre performance. It has developed a machine-learning algorithm that uses operational data to model plant performance, to predict PUE and to make decisions in relation to operations. This moves the operational decision making process towards machine learning and AI management.

What can the design and build of data centres take more generally from these hyperscale facilities? There is no technological reason why smaller data centres cannot adopt the configurations and design principles used by the major cloud providers (which are, after all, based on the concept of scalability) and many have adopted, for example, convergence, software definition, Open Source, data centre fabrics and have migrated to faster network speeds. The issue is not one of technology but one of ROI – for smaller data centres, the costs of refreshing an on-prem environment would probably not be justified, particularly if such environments can be accessed without significant capital cost through outsourcing.

"Can you learn from the large cloud providers? I think design is done to death today – I think the next big thing is operational practice. Operational practice will become operational intelligence, the systems. These will become the next area and then this will move into preventative maintenance strategies – as a rule of thumb for every dollar you spend on preventative maintenance you save five dollars on reactive maintenance spending. That's the next frontier as there is no more engineering you can do now – it's all about new technologies and product selection and integration into very common designs that are all very similar – it's all about the different types of chillers and the different types of generators and the different types of UPS – it has become a commodity, most data centre designers have 50%, 60% or 70% of the design already done. If you're linking designs for the future, then everything is about having the right structural loading, you can do the basic things like having high ceilings and strong floors to space around the building, then you can do anything inside."

"In terms of power we buy UPS efficient at low loads, down to 50%. We try to be more efficient with transformers but the greatest efficiency is with cooling. We try to be flexible like the mega cloud facilities but we can't do all of that. They have to eat their own dog food – we get smaller companies with older style computers which require 200 to 480 PDUs – we have tried to move towards 415 PDUs – that takes out a step. The hyperscale guys are into all kinds of crazy stuff but 96% of the industry is more normal."

The one set of facilities that stand to gain the most from adopting the technological innovations of the major cloud players are colocation providers. They are also the ones that stand to be most threatened by the continuing growth of such cloud players.

The enterprise client base has migrated into cloud and colocation providers may only partially be able to redress the loss of that revenue stream through attracting cloud and managed service providers. They have needed to evolve a business model based on connectivity inside and outside the facility and an interlinked ecosystem with internal communities and dense, modular, converged systems to meet the need for variable and scalable loads. This will drive the evolution of services as a matter of survival:

"Who is the market and what do they require? They are into hybrid so 40 or 50 rack clients are now ten to fifteen rack clients. It is about understanding their requirements and how we can facilitate those. Five years ago it was a fairly simple conversation around colocation based on a rack and power model, quite a simple offering. Now the model becomes harder to put together. Cloud is doubling apparently each year. For us, colocation is a short-term measure… that isn't to say that we don't make use of colocation services over an extended period but it tends to be because we can't easily predict the duration the service will be required or accurately identify what location will work best for us. Colocation is a useful tool to deliver rapid capability but it is increasingly being replaced by cloud that is instant or in-house that meets our IP related goals of privacy, control and accountability."

Section Five

Changing Technologies

The relationship between colocation, cloud and enterprise, sometimes competitive, sometimes collaborative within the data infrastructure portfolio, is complex and based on the deployment of technology.

It may be difficult for many enterprises to achieve all the opportunities presented by digital transformation if they rely entirely on in-house data centres. Outside the higher-spec converged data centre (which has pre-integrated and pre-tested components in which server, network and storage layers are able to operate together more effectively), many enterprise data centres have been designed and built for a previous era of IT and may not have the bandwidth, latency, capacity or agility to enable the development of a full digital footprint. Bringing such facilities up to the necessary specifications may be cost-prohibitive. The development of hybrid delivery options is seen as the way forward:

> "We operate a number of data centres that host a private cloud service that we use to deliver a number of services. The applications that demand the greatest scalability, such as CRM and general enterprise applications are hosted within these data centres. We also have a relationship with a number of providers of colocation services and cloud services – in particular SaaS vendors where we are gradually transitioning a lot of our traditional desktop environments. Many of our smaller global businesses will operate on a majority of colocation and third-party services. In the UK we tend to own the facilities we use. The profile is gradually changing to a cloud-first strategy. There is a need to reduce central IT costs and this is the simplest way to do this."

> "Colocation is the preferred choice for a lot of enterprise architecture. Cloud is the preference for desktop and SaaS type applications – email, mobility, collaboration and communications etc. Then there is a layer of internal systems that can't easily be migrated – call those legacy IT that we are formulating a strategy to deal with – they may be in colocation facilities or they may be on-prem."

> "High performance computing is currently all on-prem and is likely to stay that way – we can deliver that at a lower cost currently than we would achieve externally – if it was possible to buy from a cloud vendor at all. We have a real estate division and we effectively lease the space from them – but to every intent it is an internal operation."

As cloud provision has emerged as the default means of accessing IT services, it has changed also the paradigm whereby IT is consumed and paid for, towards a usage pay model. This trend presents a growing impact on how customers expect to use and buy IT and it means

Figure 10:
Power as the Key Operational Cost

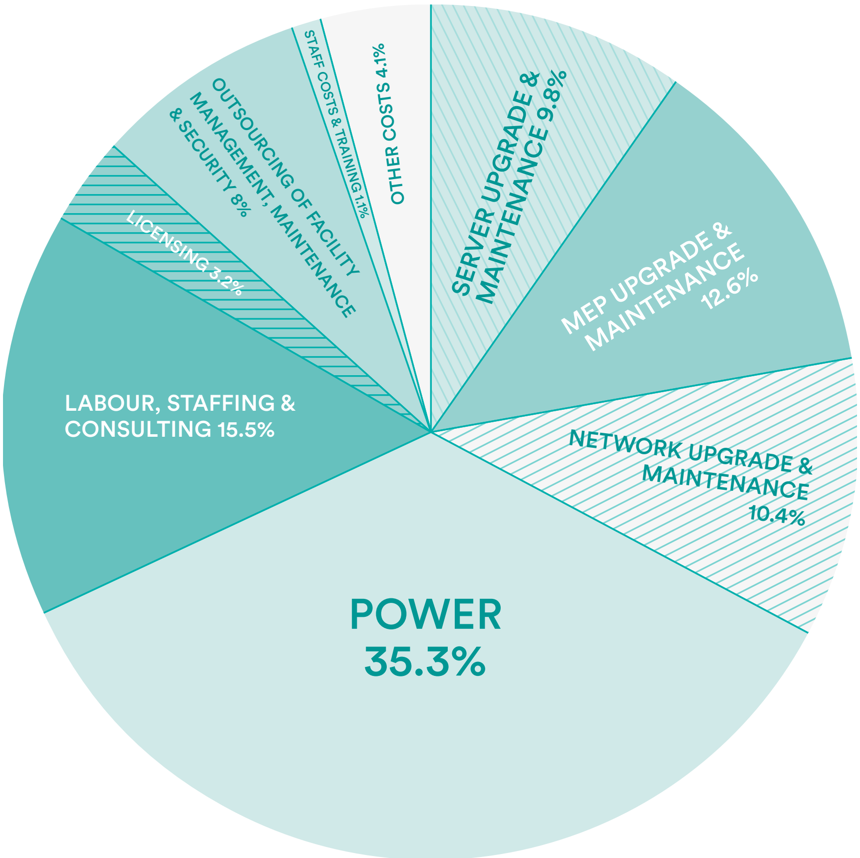

[Source: DCD]

also the capacity to predict demand and to build around that efficiently is increasingly critical. This evolution has changed the drivers for outsourcing, including colocation, from resource conservation and efficiency towards enabling the commercial opportunities that the fast-moving evolution of IT, technologies, architectures and networks can bring. In particular, those providers working from a hybrid part facility/part IT model need to develop the analytics required to optimise their space without under or over-provisioning:

> "What we are now competing against are the advances in technology particularly in server hardware. 1.5 inches of DC drive today can support up to 16TB worth of storage and you can put 32 of them into a single server that takes up one inch of physical space in one rack, one RU. And you go back three or four years using the old SAN technology that would take three or four racks."

Uncertainty is a result of the difficulty in reading upcoming technological change and, for colocation and data centre services, its impact upon the requirements for space and/or power as the unit of commerce:

> "In five years as much new space as now will still be needed for compute. Edge compute at a device level needs a network. I just don't know how quickly you get there. Corporate data centres have low fibre density. Ours are very high and as good as downtown which makes peering with AWS, Microsoft and Google easier. Our clients look for leverage from that as well as cheap power, high density, network density and more interconnectivity."

Much of the discussion about fitting out the data centre and operating it is based on strategic decisions and the balance that will be struck between on-prem, cloud and outsourcing. There are a number of general market trends that shape the design and operation of data centres.

The unpredictability of demand is a key reason for building a cluster of data centres on campuses, and for subdividing into halls and then zones within halls. For design and operational efficiency, modular and pod-based systems are used. These design principles are common to most except a minority still running facilities that pre-date the widespread deployment of modular solutions.

"We have a 20 year plan and we are using modular now to enable that. We have a zonal system for cloud based around densities. We look for branded cooling and power equipment but their prices are moving up. Energy will become a much more regulated industry like in Europe – there will be the development of an industry CCA that will arrange itself on the basis of collective best practice and licensing. There will be more green pressure on clients and lots more efficiency measures. They will need to run their IT hotter."

"We look to modular – we build in a block so we can put in capacity as we need it. We build in four quadrants with UPS and cooling, not containers. We design quadrants, data halls of 1800kW or 2700kW – the power and the cooling. There seems to be a trend to fewer racks but we don't care so long as we fill it – a couple or hundreds of racks at 5kW or 30kW – we just put in what we need."

The targeting of cloud clients and the offering of cloud services from within the facility means that the development programme for the facility needs to include the necessary speed of connection, infrastructure and level of skills that is scalable to meet variable demand patterns. For the retail colocation providers looking to develop their service profile, the key is very much seen to be in increasing connectivity. This is sometimes regarded as a matter of site and location, and sometimes of design.

"We will be launching cloud next year [2017] as we are aiming for content providers. Connectivity is a given, the more you have the better. We have access to dark fibre which has lowered the price by ten times. Connectivity will allow more local data centres and more data centre space close to the users. This will be important for edge computing – driverless cars cannot wait for a response and not all data will end up inside the biggest data centres."

As suggested in a previous chapter, energy efficiency is also assuming a new importance. Most critical to this process is the management of cooling, followed at a distance by UPS and power distribution. The more efficient use of space is part of this since a number of the measures increase the amount of space available for commercial use.

> "We want to get more eco-friendly, more green, more sensitive to what the future is in terms of technologies. If I'm going to cater to global customers then my footprint needs to be improved."

> "Cooling is the main change I have seen. There are now better designed and more robust cooling systems available – in Texas that is a big deal as we need to ensure the robustness of cooling. The impact this has on building design is one where the maintenance issues are eased and the cooling plant can now typically be housed further from equipment or in separate accommodation so that it can be isolated for maintenance and have no impact on the facility. A lot of these systems have also reduced in price so it is easier to justify having multiple redundant systems – especially for cooling – so again this impacts on design. Previously if a cooling system went down then we would reduce processing volumes to reduce heat build up. Now we would just utilise redundant cooling."

> "The change in design is relatively small in the ten years or so that I have been involved in data centres. The density of computing is obviously increasing but so too is efficiency so the cooling demands haven't changed hugely. The main thing that has changed is the serviceability of the cooling plant and power distribution systems. They are much smaller and much more efficient but they tend to be swapped out now rather than repaired on site. That means we need less space for maintenance. A lot of equipment is smaller with exception to power systems. Batteries and power management equipment is still relatively bulky."

Data centres must have an effective heat management strategy because servers generate heat while computing. In fact, for every kilowatt-hour of power consumed by a server, it produces an equivalent amount of heat. From 70°F onwards, every 18°F increase in the server's temperature results in a 50% reduction in reliability. Therefore it is critical that servers are kept at a steady temperature within a mandated range. Otherwise they run the risk of shutting down.

Servers are able to operate at higher temperature than in the past, with ASHRAE's guidelines in 2016 listing 80.6°F as the high end of acceptable server inlet temperatures. However, only 5% of data centre operators maintained their facilities at this temperature in 2016, suggesting a considerable degree of 'overcooling' of facilities and therefore of wasted energy. Between one-third and half of data centre power spending is on cooling, and power is the single largest operating expense for most data centres. Larger facilities might save hundreds of thousands of dollars in OPEX by adopting a more efficient cooling strategy.

As server densities rise, the amount of heat generated is also increasing. New processors generate more than five times the heat of older processors, and new servers and switches could generate up to ten times the heat per square foot as those from ten years ago. Racks supporting densities of 10kW and higher have increased from 13.7% of all racks worldwide in 2011, to 21.1% in 2017. Apart from new servers running hotter, data centre managers also face the challenge of their distribution across the data centre floor. Data centres have a mix of old and new equipment usually, meaning that some racks run considerably hotter than others. This creates hot spots, which then have to be managed, otherwise blocks of servers could go down even if the data centre as a whole is run at an acceptable ambient temperature.

Data centres have started to explore a number of new cooling technologies and architectures to add further efficiency and cope with increasing rack densities. These include the introduction of such cooling technologies as:

- Vertical Exhaust Ducts which are chimneys that extend from the rear of data centre racks to the hot air plenum in the ceiling. They provide a closed channel for hot exhaust air to be returned to the CRAC. As such, they eliminate the hot aisle, keeping the data centre temperature at an acceptable level.

- Heat Wheels. These absorb heat from the data centre and release it into an outside air plenum as they rotate. They are claimed as highly efficient, with some vendors claiming heat transfer efficiency of almost 90% while using 75- 90% less power than alternative approaches.

- Close-coupled cooling is the principle of directly cooling into the row or even rack where it is most needed. Approaches include in-row, in-rack (including rear door heat exchangers) and above-rack. Liquid cooling is often utilised for close cooling since water is seven times more efficient at absorbing heat than air. Increasingly, refrigerants are being used rather than water as they are even more heat efficient, and they evaporate so liquid leaks are less of a threat than is water to the IT equipment. According to various vendors and end users, close-coupled solutions can improve PUE by up to 20-30%, and reduce power usage substantially.

- Chipset and server-based solutions where the cooling is located at the chip.

- Intelligent control, monitoring and targeting technologies to improve cooling efficiency. Other initiatives include moving from raised floor to hot aisle containment, the greater use of natural air and running the IT equipment at higher temperatures.

Cooling is considered as the most critical infrastructure function to control efficiently:

> "The IT equipment is resistant to temperatures up to 45° so it won't be cooled in the same way and reliability will be different as there will be a network of sites that back each other up, so N+1 reliability will still be there but it will be there in a much diminished form. I could go to China and show you a data centre with no air conditioning at all. It runs every day in the hot season and in the cold season without any air conditioning and it is using all the same IT kit as we are in state-of-the art facilities here, or in Europe or the Americas and you have to scratch your head. The reality is the paradigm is now they can do it without any degradation of performance or any downtime. They have not been brought up so much on ASHRAE since the internal temperature inside the CPU can handle 45° so why can't I do that? The laws of thermodynamics are very simple yet we have these standards like Uptime, CRA and ASHRAE and instead of those documents being guidelines and then applied intelligently they have become almost like de facto reference bibles. People are scared to move."

Figure 11:
Natural Cooling Methods Currently Used

N=Unweighted Global	2012	2016
Air cooled chillers with free cooling	26.7%	43.2%
Direct air cooling (fresh air from outside)	26.5%	46.6%
Water-side economisers (water cooled CRAC)	26.3%	36.5%
Water cooled chillers with free cooling	26.1%	34.5%
Evaporative cooling	19.4%	26.6%
Indirect air economisers (thermal wheel/Kyoto cooling)	7.6%	11.3%
Other methods	12.2%	16.5%

[Source: DCD]

Figure 11: Data from the DCD Census in 2012 and 2016 indicates growth across a variety of natural/free air cooling methods including air cooled chillers with free cooling, direct air cooling (using fresh air from outside), water-side economisers (water cooled CRAC) and/or water cooled chillers with free cooling. There is commonality in adoption indicating that organisations that are innovative in their cooling practices and both able and willing to utilise 'free' cooling will look at different options. Adoption is, however, strongly regional, with adoption highest in North America and Europe.

Figure 12:
Natural Cooling Methods Planned

N=Unweighted Global	2012	2016
Direct air cooling (fresh air from outside)	28.3%	31.2%
Water cooled chillers with free cooling	22.2%	26.5%
Air cooled chillers with free cooling	21.0%	23.4%
Indirect air economisers (thermal wheel/Kyoto cooling)	20.3%	26.8%
Water-side economisers (water cooled CRAC)	19.8%	23.4%
Evaporative cooling	14.7%	26.6%
Other methods	12.5%	14.3%

[Source: DCD]

Figure 13:
Renewable or Decentralised Power Generation Options Under Consideration

N=Unweighted Global	2012	2016
Battery storage solutions	16.5%	38.9%
Solar/photovoltaic	11.1%	30.2%
Smart Grid	9.7%	19.6%
CHP/co-gen	8.3%	19.5%
CCHP/Tri-gen	7.9%	18.4%
Wind	5.3%	18.4%
Biogas/fuel cell	5.1%	16.7%
Biomass/co-gen	2.5%	16.0%
Other methods	8.1%	11.2%

[Source: DCD]

Figure 12: There are highest levels both in 2012 and 2016 of consideration for direct air cooling (using fresh air from outside), followed by water cooled chillers with free cooling, air cooled chillers with free cooling, indirect air economisers (thermal wheel/Kyoto cooling) and water-side economisers (water cooled CRAC).

Figure 13: In terms of energy efficient power generation options, only a minority of these organisations are looking at renewable or decentralised power generation options, most usually battery storage solutions. There has been a stronger growth in consideration for the four most popular power generation options than for the cooling options covered above. This is because cooling efficiencies have been a far more established concern of data centre operators as a leading opportunity for cost saving. The evolution of power generation technologies between 2012 and 2016 has been considerable.

The reasons for the increase in interest in cogeneration/ combined heat and power solutions is because data centres can use the ancillary benefits provided by their heat management solution to defray the environmental cost of the energy they use. District heating, i.e., the ability to re-use heat rejected from the data centre for heating local homes or community buildings, could be an important benefit. In some cases, energy can also be sold back to the grid, improving the business case and relieving operational expenditures. However, the cost of distributing district heating can be quite substantial and often is a limiting factor.

Despite reservations about the metric, PUE targets are mentioned as a key component still in the drive towards greater efficiency:

> "We started DC 1.0 in 2005/6, for our current generation we have developed DC 2.0. 1.0 was a raised floor room with Compact Room Air Conditioning Units (CRAC) units on the floor, looking at maybe 75W/ square foot and 35W/square foot for telcos. Now we have air handlers outside so there's more space for servers. We have moved from raised floor to hot aisle containment via chimney or enclosure and can get

to 150W/sq ft – we can get a higher density with an average cabinet density of 5 to 8 kW/rack. That's 80% of what we offer – some are between 12 and 20kW/rack, some as high as 25, that's IT load only. Bitcoin, for example, looks for 40kW/rack."

The theme of a multi-profile facility environment impacts also the provision of resilience. Most of these providers who are running campuses and larger facilities have evolved a mixed Tier system:

"I think things have changed but other things have stayed the same. For electrical design we rely on 15MW on each string of UPS loaded at 45% at 100Kva. For levels of redundancies we use N+1 moved to N+2 and the same with gensets. For basic sizing we favour 2MW gensets."

"We have a wholesale data centre of 20,800 square feet. It is Tier III+ except for electrical where we are 2N. We have both high and low voltage but we are not fault tolerant."

There is mention also, particularly in established markets, of clients turning their back on more traditional methods of defining resilience and replacing facility-based classifications with network-based load shifting in the event of any threat to availability:

"Data centres have not changed that much, requirements are getting smaller. It used to be 2-5 MW now it's maybe 250-1000 kW. There's less resilience required – very little call for Tier IV or even Tier III. They want concurrent maintainability and to use their networks for redundancy – so now it's maybe Tier II or I."

The pressures on these providers are acknowledged as they need to deal with an industry that is fragmenting into different sets of client requirements and therefore making different demands on the facility. In order to manage this state of affairs, these providers are looking at operations that are intelligent and flexible, whereby clients can be charged on the basis of a number of business models. This means that a systematised and data-driven method of operation is required across different locations and different environments.

As colocation and data centre services facilities play a greater role in hosting and delivering enterprise IT, so this increases the need for monitoring and management across the portfolio. The replacement of more traditional facility-based Tier classification systems by load-shifting across networks, the growth in requirements for data portability and the dawning of a new era of data sovereignty legislation (where that legislation is updated to recognise the current realities of data traffic) will all stretch the management of provider portfolios. This therefore represents another driver for 'intelligent' data centre management:

"Strategic foresight involves looking at where the market is going. We are not looking at the multi-MW business, more at 500kW to 1MW. Our approach will be based more on services – a portal that allows people to look at their power infrastructure and utilisation and we offer some degree of flexibility – if you don't use all the power you don't have to pay for it. We don't work it so you get to the end of a contract and hike your prices 80% because we know you can't go anywhere else. We have gone the other way – have a rack for a day."

A number of the respondents mention partnerships with management software/DCIM providers in order to help develop a more intelligent facility. Retail colocation is, however, seen to present some barriers to DCIM as the IT and the facility are in totally different hands, and outsourcing into colocation may be only a short-term arrangement. In operational terms, there is high mention of the need for greater control, better capacity planning, and better disaster recovery:

"The next five years will see a drive on efficiency at all levels through power distribution, UPS, raised floor. Security demands a lot of focus – we look at BMS systems, we look at all the sensors and determine how the ops team can actually take action on this with a few clicks to a dashboard. We will have more software and automation on the operations side. We're not some old DC with an Ops room and a guy coming out when there's a problem. It's more like network ops now and we're focused on that. We will drive this across the global platform."

Figure 14:
Management Systems Currently in Use

N=Unweighted Global	2012	2016
Threat management systems	65.3%	83.7%
Security command & control systems	64.3%	90.9%
Network & IT capacity management	63.7%	91.7%
Asset tracking	61.3%	78.3%
Maintenance tracking	60.9%	76.5%
Infrastructural change	59.4%	84.1%
Tracking power availability	59.2%	76.6%
Building & energy management	58.0%	64.6%
Tracking energy costs (electrical & thermal)	44.9%	61.6%
Energy optimisation & carbon reduction	29.1%	43.6%

[Source: DCD]

Figure 14: A majority of data centres across the world use management systems including DCIM to manage their operations. The difference between the majority of systems deployed below and those in the process of development by global cloud and colocation operators is the degree to which machine learning and intelligence overlay their facility management.

This growth in space accounted for by colocation, hosting, IT service and cloud providers will also change the security focus among organisations that were previously guarding their data on-prem and who are now trusting it to an external provider. The market shift to colocation, hosting and cloud-hosted services is also likely to put more pressure on physical security requirements, especially in multi-tenant facilities that see many different service providers and their customers house equipment side-by-side. The increasing connectivity of data centres, the automation of facilities and the use of IoT and analytics for management and operation will also increase the risk of cyberattacks on the infrastructure stack.

Section Six

The Future of Data Centre Build

The success of a data centre build is judged on two main criteria – the first is the outcome that the end data centre can meet the business objectives for which it was commissioned and the second is the process needed to enable the outcome.

In data centre literature, the second criterion is fairly much secondary – for every 100 papers, presentations or articles on technological, operational and design changes to the data centre, there may be a handful on the process of actually building. Possibly it is assumed that the means and the end are seamless.

The means of building a data centre is changing to reflect the increased importance to the business and the technology available to do so. The number of parties (whether people or departments) involved in major data centre decisions (build, extension, major refit/refresh) has increased for an average range of 7-10 per project (2007 research) to 15-20 in 2016. This reflects the increased numbers of technologies and disciplines that the data centre now encompasses, the increased connection with the business, increased risk aversion and legislative requirements.

In emerging regions of the world where new build is high, the requirements that the process of building a data centre makes on local skills are in part responsible for local shortages. The major global players with established build programmes have therefore made it a priority to ensure they have recruited or developed the skills base necessary to execute those future build programmes.

The process of decision making and procurement for data centre build has become more thorough and more accountable as data centres become more mission-critical and therefore more protected against risk. As part of this, procurement has become more formalised, less reliant on open tender and more involved down to items of lower value (as the principle that any minor change can affect the operation of the data centre). Therefore, formal policies, service-level agreements, preferred supplier and tendering and/or preferred supplier process have become normal procedures.

This may change the framework of procuring data centre builders on a construction project, and this reflects the broader industry as (some) buyers and suppliers look for greater flexibility and truer 'partnership' than is usually the case with the established and linear design – bid – build model. This allows the customer to maintain control over the project while the separation of responsibilities creates a degree of testing through the process. This means of build puts the onus on the client to manage the process and, historically, this process within the broader construction industry has been fraught with 'variation' and contractual revision as issues are passed from one stage to the next.

Figure 16:
Role in Facility Decision Making within Client Organisations (2015)

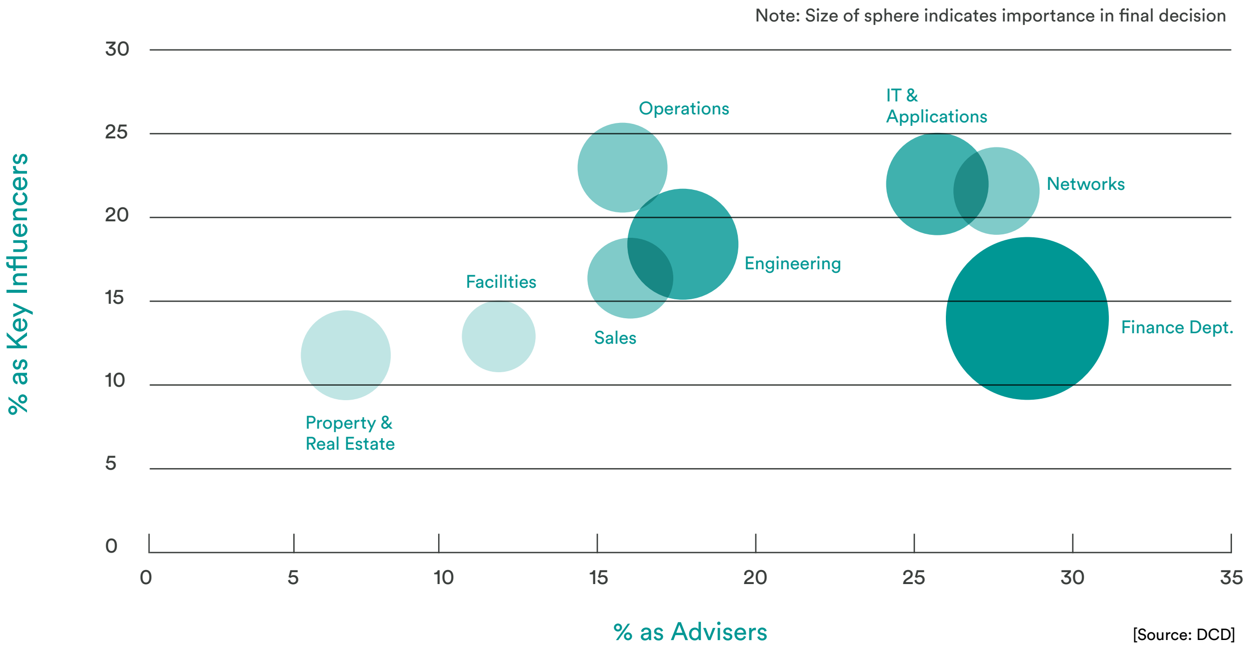

"The only issue was getting the various trades to be available at the optimum time for the project. It was a challenge to ensure that momentum was maintained and when a delay was encountered it would impact on subsequent stages. That is only an issue if you manage the project internally."

"The requirement for most projects is born out of a growing need for more data handling capacity and the planning horizon sits at around three years at a strategic level and probably just two years for knowing exactly where we need a facility and the capacity of it. The business is starting to identify sites for dedicated data centres that can be built with a five-fold size increase to allow for future consolidation of data centres should that become practical and necessary."

"To create the perfect data centre for us involves absolute certainty that all the critical systems will work as planned. We used an engineering firm to check things like adequacy of ventilation and cooling systems and to forecast power consumption. We also tested the various redundant systems, from cooling to power to connectivity. Probably nothing unusual in that but we went to considerable lengths for some systems to make those tests as realistic as possible."

The interviews suggest that a number of the companies which have relied on this means of procurement have looked away from it, particularly for larger and more complex projects where the values of integration are more important:

"It moved forward but then kept looping as issues that we thought had been resolved reappeared and had to be solved again – mainly technical issues but also some issues of provision and utilities."

"The main issue I have had over the years is the total lack of understanding of the relationship between the contractors and the equipment suppliers. I am concerned that decisions are not made in our best interest and instead based on the best deal they can make with suppliers – for this reason we try to do as much of the procurement as possible to ensure that there is no hidden economic driver for a particular brand selection."

"There is no perfect process – the main thing I would look for is flexibility. There is no way to accurately forecast computing needs for a data centre as the technology can change slightly and that can mean a change is requested at short notice. Having a vendor

Figure 17:
Final Authority in Facility Decision Making within Client Organisations (2015)

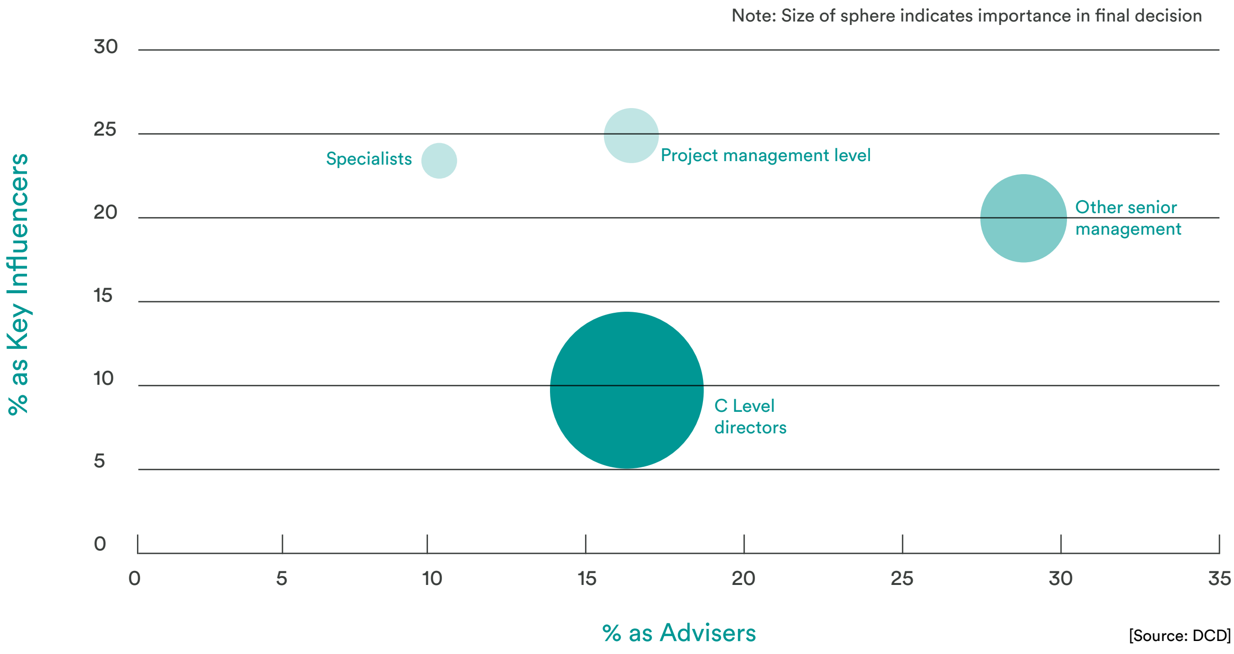

that is flexible is a huge benefit. We are more likely to subcontract the whole build than to manage individual aspects of the build as past experience suggests that it can reduce the time to market. For us having the asset working as soon as possible recovers any likely cost savings from managing the project ourselves. The main benefits of the main contractor route is the ability to de-risk and the ability to work the asset sooner and achieve a more rapid ROI."

This may lead to concentrating more of the project in fewer hands, thereby creating a greater project synergy between stages.

Figures 16 & 17: The people involved in decision making for major data centre projects have also changed. In the 'legacy' era, the structure would have been based more on personnel directly involved with IT and facilities within the corporate hierarchy. Now C-Level and senior managers with broader corporate responsibilities are more likely to be involved. Broadly, involvement in decisions now reaches far more widely across the organisation and is based more on skill set and capability and will include external specialists within specially constituted project teams.

The development and deployment of virtualised, cloud-based and software-defined data centre provision has been based on a number of promises of speed and simplicity. Because the outsourcing of space and services is the major growth area in data infrastructure, the construction of new enterprise facilities is more carefully thought about. It is no longer in most cases the default option.

Facilities management, IT and apps management and executive leadership are equally important in major decisions relating to data centre investment while engineering and operations have secondary roles.

A majority of C-level executives have final responsibility for making major facility decisions. They are supported by technical and engineering specialists, project managers and other senior managers in terms of influencing outcomes, and advice from senior managers.

Figure 18:
Role of Suppliers and Partners in Data Centre Investment Decision Making Process

Percentage working with clients on…	Technical advice	Evaluate	Recommend	Implement	Design, develop	Strategy
Equipment vendors	35%	39%	55%	34%	46%	39%
M&E designers and consultants	43%	48%	74%	40%	48%	39%
Architects	48%	48%	88%	41%	48%	40%
Integrators	28%	62%	84%	42%	46%	60%
Telco and network providers	41%	48%	88%	28%	40%	46%
Other consulting	42%	48%	68%	41%	40%	40%
Providers of finance	28%	62%	81%	62%	41%	81%
Colocation providers	21%	42%	81%	41%	46%	48%
Cloud providers	24%	48%	80%	46%	46%	48%

[Source: DCD]

Figure 18: External suppliers and consultants play a greater role in the decision making process, and according to DCD Census data their role is increasing and they are being viewed more and more as partners rather than just suppliers. This is in part because the larger companies who originally specialised in supplying facility equipment (power, cooling, UPS etc), IT equipment (servers) or network equipment have sought to move into professional and IT services in order to maintain their competitive and financial position.

Broadly, external suppliers and consultants play the greatest role in making recommendations to their clients, in design and development specific to a client's requirements and in providing technical advice and evaluations.

IT-system integrators play a particularly strong role in evaluation, making recommendations, design and assisting with strategy. Organisations providing financial advice on data centres have a similar profile, indicating the increasing requirement to invest more actively in the sector. Colocation and cloud providers play a large role in making recommendations, indicating that these options might be new and less well understood by some clients.

Equipment vendors play a particular role in making recommendations to clients.

Partly as a result of the increased capabilities offered by major global suppliers and trends towards convergence and turnkey solutions, there is a trend toward a single provider of facility components, including enclosures, power distribution and protection, cooling, cabling and monitoring, rather than relying on different specialist providers. This may also indicate some requirement to make the procurement process (and the maintenance/upkeep process afterwards) more streamlined.

The interviews conducted with companies which had built a data centre over the past two years indicate the following expectations of the pre-commissioning process.

The people undertaking different parts of the project are specialists, where these are not available in-house. At the most basic, specialists are looked for in structural/'shell' work, Heating Ventilation and Air Conditioning (HVAC), power distribution, equipment fit-out and infrastructure. Further specialists will be needed if there are particular needs, for example to increase the connectivity to and from the site, particular requirements of utilities or government.

All of these projects changed some of their design parameters as they progressed, normally as the result of different areas of expertise being introduced, this means that flexibility needs to be built into the process.

All of these projects changed some of their design parameters as they progressed, normally as the result of different areas of expertise being introduced. This means that flexibility needs to be built into the process, usually around key 'milestone' decision points at which new contracts will be tendered or agreed.

The 'ideal' pre-commissioning stage rarely happens but this means that broadly clients view the design element as more important than the builder involvement.

> "Although it isn't necessary for us to be involved in the detail of design I feel that it would be better to be informed of decisions as they are being made. There is potential for a problem to occur if expectations diverge from the solution being developed."

> "I think the design process needs to involve more people at the very early stage to get ideas and understand current limitations. The engineering team didn't know anything about our build until it was a done deal and that hasn't been a problem but I do wonder what opportunities have been missed as a result."

> "So an ideal pre-commissioning all of these variables would be fixed and there would be no ambiguity and the internal client (in our case three different R&D customers) would know the specification they required. We had a smaller than ideal space, less than ideal working conditions and were still negotiating with other business units over neighbouring space that we wanted to use for data centre space."

> "A more client-centric approach would have been better – so the cabinets and infrastructure was installed for 60% of the centre and there wasn't a clear plan on priority – so the improvement might be to design not just the finished layout but the layout at various stages of the data centre lifecycle."

> "Proven track record in data centre design would be useful… but the fabric of the building is not particularly different – so the builder is less critical than the designer. The designer would need to understand how the space is used, what access is required and for what size of cabinet/equipment. How that equipment needs to be serviced by cooling or power and what space those systems require."

There is mixed tolerance of schedule over-runs so long as their outcome will lead to improved project outcomes and any cost implications have been agreed in advance. Most recent projects have suffered over-runs – what causes friction is when this leads to claims for variation:

> "I think the speed at which a data centre can be built is a critical element – I have seen examples in the past where the build was so protracted that the workload was housed in a colocation facility and then never got transferred – the business case was a mess and

the data centre was very nearly disposed of shortly after completion due to a lack of need. It was only pressed into service after another business division required space. So the speed to market is an essential element – an ability to settle design details while the construction work is underway."

"Timescales for build were probably longer than had we outsourced the entire build and we had a handful of changes to design to manage – that led to further delays and complexities in the management of the project."

"The complexity and cost of even minor build means that best practice on such projects relies equally on coordinated teamwork and on a senior provider, sometimes a formal project manager, sometimes not, who can achieve an overview of the whole project."

"I would say that first and foremost good designers and builders are flexible and can work with each other to solve the problems that you present along the way – in our case we added more space and that seems on the surface a good problem to have but it led to constrictions in cooling and power elsewhere in the building and the impact of those changes meant several other issues needed to be solved."

"A good project manager can see the whole project – even if they are not responsible for elements of it. One example of this is the project manager we used for the internal work was thorough with the structural design checks before undertaking the work and noticed a number of issues and potential areas for improvement. The PM then liaised directly with our structural team to have them implement changes that ultimately saved us time and long-term maintenance costs."

"The other issue I have had is trying to work down to a budget – it can be difficult to identify where cost savings can be made when a data centre build looks over-budget. Scaling back some areas of spend is difficult. The process needs to be more collaborative between real estate agents, designers, builders and technologists. To get to the end to find the final project is over-budget, each agency suggests making the cost saving elsewhere and that is a time-consuming process to identify where there is unnecessary spend."

Other requirements include skill in modifying within existing buildings. While most headline data centres are green field construction projects, the vast majority of data centres, especially enterprise data centres, are located in buildings they share with other corporate activities:

"For us they would need to work within an existing building… it is unlikely that we would stray from our current strategy to host critical system within our office building and to migrate the rest to colocation. So that would be the first expectation – a data centre builder would need experience of modifying an existing building. There are not a lot of differences that I can identify – floating floors, equipment lifts are sometimes used and that sort of thing, but some product-specific knowledge – especially when managing a complete build – would be essential. They would need to know their way around power management services and ventilation on an appropriate scale."

"Everything is an order of magnitude more complex when converting typical office space or high-rise, to accommodate the needs of a data centre and high-density areas are difficult for noise/air pollution when stand-by-power is a consideration. The dedicated purpose-built data centre has better vehicle access, easier maintenance and better designed location of HVAC, power distribution and power generation services."

Reaction to design software is mixed, with the basis of the division of opinion being the skill of the consultants using the software:

"I think there is an over-reliance on design software and some more creative thinking could solve some of the design challenges we have encountered. There is a need for tools to aid in design but the pace of change possibly makes these difficult to maintain. Design capabilities have advanced a lot. The internal survey of our space was done electronically and with a huge degree of precision and then 3D modelling was used to manipulate layouts – the design process was far more iterative than previously due to the ease with which a layout could be changed."

"The best designers I have worked with have the ability to computer model the interior space and use 3D laser

Figure 19:
The Cumulative Cost of Data Centre Ownership

[Source: DCD]

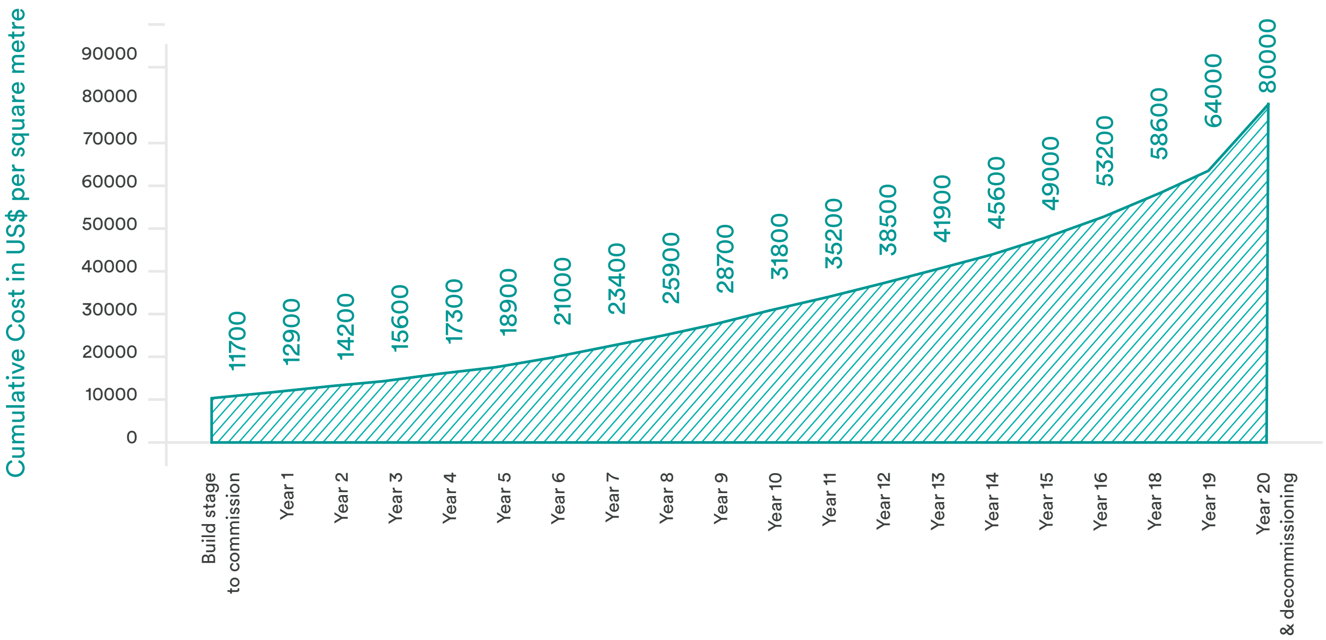

scanning techniques to get accurate dimensions from existing buildings. The process aids rapid decision making and some software can model heat dissipation requirements. Designers should be able to advise on temperature for operation of different equipment and how to best locate equipment so that cooler running equipment is grouped away from higher-temperature equipment wherever possible."

The criticality of the data centre and the aversion to risk means that cutting-edge technology is judged more conservatively:

"Advances in technology tend to come in waves – there can be relatively long periods where there is seemingly little innovation in cooling – and then there will be a steep change in the density of equipment of heat dissipation characteristics of equipment that shifts the design requirement. There is nearly always an issue with all of these steep changes in capability – no one will risk designing these new characteristics into a centre when they are unproven. There is a risk that if the technology fails to perform then a major redesign may be required. So even when there is a steep change in capabilities you normally design to at least one build version earlier."

"There is a new scepticism within our business of some of the cooling systems that are sold with green credentials. They put efficiency ahead of performance and while that has worked well for the majority of time it could have been an issue if we did not have the redundant systems so easily available."

So what is the future for designers and builders?

The construction industry overall has some major responsibilities globally since it accounts for around 7% of GDP (more in emerging markets with higher infrastructure requirements). It is also the sector that consumes the greatest amount of raw materials, and buildings are the largest source of carbon emission. There is considerable awareness of environmental issues among organisations which have commissioned data centres:

"It is increasingly evident that green systems (more energy efficient) are adaptable to run on renewables, is a key battleground for the future. It is a senseless activity when energy consumption with the logistics chain itself is so much greater – there are far better places where we might save energy but the data centre seems to get more press and is thus a soft target for energy saving. I expect therefore to see more pressure to reduce energy consumption.

Increasingly, the design process must pay attention to environmental concerns and we are now at a point where solar power can actually be more economical than running utility supply so a good design and build partner should know this side of the data centre project well. I think this is an area that a number of vendors could improve on."

The construction sector in 2016 accounted for around 10% to 15% of the value of data centre investment and expenditure (which was around US$140 billion in total). Design and other pre-commissioning requirements accounted for another 2% to 3% of the total. The data centre construction sector is more complex than just being a series of building companies. In major markets there are companies that specialise in the construction of data centres (and some in other types of mission-critical buildings) while in smaller markets there are no specialists, just business units within larger construction groups.

Since most dedicated data centre buildings are constructions where the outside 'shell' is relatively simple in comparison to the complex IT, infrastructure and network fit-out inside. This has enabled other classes of suppliers – equipment vendors, system integrators – to become builders. As with other construction markets there is also a constructor/developer sector, organised increasingly as Real Estate Investment Trusts (REITs).

The construction of data centres will follow some projected overall sector trends but not all of them.

A key difference between many construction projects and the data centre is the balance between the 'upfront' costs of construction and subsequent operational costs. Across the industry as a whole, the upfront costs are estimated (on commercial projects) to be between 25% and 45%. For data centres, the proportion is far less – Figure 19 calculated in 2015, based on respondent survey data, suggests that the proportion for a greenfield Tier III 500 square metre traditional build would be around 15% upfront on a 20 year TCO. While there is no such thing as an average data centre or an average commercial construction, the very high operating costs of the data centre mean that the construction sector will need to focus on 'whole of life' not just the initial construction costs.

Based on a number of reports and analyses, the construction and engineering sector is characterised as 'conservative' when it comes to process, by emphasising the final product rather than the process of delivering it. This conservatism can be indicated by information from established markets that suggests that the sector makes less use of IT beyond administrative functions than other sectors. This in turn can be linked to the fact that while IT is a necessary driver of service businesses, it is seen as more optional in production and construction sectors. It reflects also the considerable numbers of small and single-operator businesses across the sector, and their skills base.

Particular areas of improvement indicated by analysts indicate the coordination of skills groups through a project where they may become separated through a sequential delivery process, gaps in the transfer of knowledge including in the monitoring of projects.

The construction of data centres shares some of the characteristics of conservatism noted above and the areas of improvement are mentioned by companies building data centres. However, a major factor in this is the risk profile of data centres which will in itself encourage conservatism. The unpredictability of technological change remarked on by a number of respondents means that the initial construction and fit-out is not seen as the only version that the physical structure will see in its lifetime. Therefore, adoption of cutting-edge options may be deferred until later in the data centre lifecycle.

Despite the headlines given to innovation and technological development particularly in the IT stack, innovations take time to percolate through. Both cloud computing and edge processing have been around in nascent forms for over 20 years – the initiatives of Amazon and others have driven cloud into the mass market, while the future of edge is still taking shape. The enduring division of the industry between 'facility' and 'IT' possibly fosters conservatism in order to obtain collaboration for shared decision making.

A feature of the immediate future shared by the construction and engineering sectors is that of skills shortage. The data centre sector shares this, particularly outside the original IT and facility skill sets. Shortages have been measured in the skills required for digital transformation, in data analytics, in design and architecture, in coordinating between technology and business, in hybrid IT processes and in most other places where cloud and legacy intersect. The construction industry similarly will face skill shortages although these will continue to be localised and linked to the extent of construction activity in particular markets. In both industries, the profile of the key engineering function is similar and in established markets they will tend to be older and male. In the United States and Western Europe, the profile of the engineering sector has reached the stage where the numbers coming in are not replacing the numbers retiring. To some extent, the data centre sector has diversified its demographics through the larger role that IT is playing in the industry.

While the construction industry overall will focus increasingly on environmental performance in terms of materials, methods of construction and local environmental impact, data centre designers and constructors have for some time used some of the major techniques to reduce environmental impact such as modular construction and building components offsite to ship and fit locally. The enormous impact of the power bill on the TCO of a data centre means that ways of improving energy efficiency are designed in, and many large data centres are linked to renewable or sustainable power sources or, where possible, recycle heated air to warm local buildings.

Other areas of future direction include the greater use of BIM systems and the development of contracting models away from the traditional client/supplier relationship. The complex series of relationships that characterise data centre construction where the same company can adopt different roles in the supply chain on different projects means that on many projects the degree of collaboration has already evolved.

Appendix

Ireland – A World-Class Exporter of Data Centre Services

Ireland's strength as a location for data centres has led to the development of a world cluster of companies with an unparalleled competency in data centre design, build and fit-out.

Unequalled track record

The country's infrastructure, climate, skill base and business environment has made it one of the principal data centre hubs for many of the world's technology giants, including Microsoft, Amazon, Google, Dell EMC, Yahoo, IBM, HP, Facebook, Equinix and Digital Realty. Today, it continues to attract international data centre operators and colocation leaders.

Irish companies have developed successful partnerships with these global technology companies and continue to win new projects as the domestic data centre market prospers. Services provided by Irish companies include design and project management services, general construction, mechanical and electrical services and fitting out data centres with energy and security solutions.

With guidance and support from Enterprise Ireland's global network of offices, Irish contractors, construction services and engineering companies have exported this data centre expertise internationally. These companies are now delivering major projects across the UK, Europe, the Nordics, North America, Russia and the Middle East.

Enterprise Ireland client companies have grown into market leaders in the global high-tech construction sector. A particular area of expertise for these companies is owner-occupied, hyperscale data centres, in particular supporting the EMEA growth of the global technology companies that prioritise sustainable energy, low energy consumption and intelligent systems in their design and build process.

Flexible, partnership approach

To meet this client demand, Irish companies have demonstrated that they can mobilise construction and engineering professionals and a world-class supply chain in a timely and efficient manner for major overseas projects. As a result, they have earned an international reputation for their technical skills and performance in the execution, management and cost engineering of the most complex data centre projects.

Repeat business has become the norm for Irish companies in the high-tech building sector largely because they have proven their ability to deliver on the most demanding client requirements and schedules – even when requirements change and evolve over the course of the design and build process. The extent of their experience means there are few challenges that Irish companies have not previously encountered and have consequently developed the skill set to solve complex problems throughout the construction process.

It is common among Irish construction and engineering companies that staff retention is high. This means trusted working relationships between partners can be developed over time and there is minimal disruption for clients when transitioning from one major project to the next. Strong and stable relationships are of particular importance, where a complex mix of high-tech suppliers work in partnership with principal contractors, global technology companies, developers and asset managers.

Innovation leaders

Driven by the demand to build ever more complex, large-scale projects in shorter and shorter timeframes, Irish construction services and engineering companies have been at the forefront of introducing a range of innovative technologies and work practices.

In recent years, the Irish sector has been ahead of competitors in adopting digital technologies such as Building Information Modelling (BIM). This intelligent digital process gives architectural, engineering and construction professionals the insight and tools to more efficiently plan, design, construct and manage buildings and infrastructure. Irish companies have begun extending BIM to incorporate 4D and 5D modelling, which allows real-time visualisation of a project's progress and costs to be attached to different elements. BIM is increasingly being integrated with Data Centre Infrastructure Management (DCIM) models to provide a sophisticated management tool across the full life cycle of the data centre.

According to the Building Information Modelling (BIM) 2017 report by Construction IT Alliance (CiTA) in association with Enterprise Ireland, Ireland has made remarkable progress in recent years in building BIM capability. It states that Ireland is as mature as any country in the developed world in relation to BIM proficiency and use in high-tech construction projects.

The development of lean processes across the supply chain is contributing to a steep change in quality, productivity and health and safety in Ireland's construction industry. Cost and time savings for clients are accrued through the elimination of wasted time and delays on-site, the reduction of material waste, duplication, errors, absenteeism and reduced snagging. Offsite fabrication is another innovation driving this efficiency. Manufacturing modular components offsite means there is less opportunity for costly mistakes on the construction site itself and allows considerable work to have been completed before the site is made ready for building.

Overview of Irish data centre services:

- Consulting (including site selection)
- Design & Build
- Mechanical & Electrical Services and Equipment
- Construction
- Architects/Building Engineers
- Integration/Fit-Out
- Data Centre Change Programmes
- Service & Maintenance/Property Management
- Security
- Energy Solutions
- Data Centre Digital Construction & Facility Management Information
- Data Centre Operation

About Enterprise Ireland

Enterprise Ireland is Ireland's national export agency. With over 30 offices worldwide, our local sectoral specialists work with customers to understand their unique challenges and match them with Irish suppliers that can deliver leading edge products and services.

Our goal is to simplify and enhance the purchase process for customers and build successful, long-term business relationships between international companies and Irish suppliers.

Address
Enterprise Ireland
The Plaza
East Point Business Park
Dublin 3
D03 E5R6

Contact
+353 (1) 727 2000
irishadvantage.com
irishadvantage@enterprise-ireland.com

ENTERPRISE IRELAND

IrishAdvantage.com