

REGIONAL DATA CENTRE STRATEGY: AT A CROSSROADS?

As compute and storage demands have risen, the traditional data centre strategy of 'build and they will come" has served the industry well. However, the demands of enterprises – enabled by AI-driven automation – are becoming more diverse. A one-size-fits-all approach will no longer work: Data centre strategies must become more varied and more customer-oriented.

Chris Barraclough, CEO

October 2024

© STL Partners

The traditional data centre strategy: build and they will come

The data centre industry has enjoyed enormous growth over the last two decades as the Internet has spawned vast quantities of data and new 'as a service' business models in software and compute. Much of this growth has come in the form of hyperscale facilities from the major players – Google, AWS, Microsoft Azure, Meta, Alibaba and the like. Hyperscale data centres are more energy efficient, and their scale enables them to balance compute and storage demands effectively. Regional data centres – located in tier 2 and 3 cities – have provided colocation facilities for enterprises seeking local compute and storage that want to avoid the hassle and cost of running their own facilities.

With strong demand, the focus of data centre operators has been on solving supply-side issues: securing powered land, access to water for cooling, and planning permits. It is only a little simplistic to say that the data centre business has been a real estate play: securing sites, managing the build, and operating a cooled warehouse securely and efficiently. Filling up a data centre has been relatively straightforward: Hyperscalers can predict demand for their services accurately and so haven't 'overbuilt' and regional data centres have been supply-constrained so there has been plenty of demand from enterprises and governments (and Hyperscalers moving closer to the 'edge') for their facilities. But this situation is changing.

New forces driving data centre transformation

Until recently, managing compute workloads was relatively straightforward for enterprises. There were broadly three compute and storage options:

1. For applications requiring ultra-low latency – such as manufacturing systems – **on premises compute** was the answer. This was also true where close enterprise control was deemed critical – where an application or company data was too sensitive to manage in a third-party data centre.
2. **Colocation**, usually in a local data centre, was an option where data sovereignty remained important – for example where the public cloud was deemed a risk owing the potential for data to be moved beyond a specified jurisdiction.
3. For many other workloads – those used to run software as a service, social media, data analytics, disaster recovery, etc. – the flexibility and cost benefits associated with **public cloud** was the solution.

Now, there are alterations both on the demand side, from customers, and on the supply side – from investors, regulators, and from the changing nature of data centre competitors and the structure of the industry – which call into question a strategy focused predominantly on building facilities. For many regional data centres and those 'at the edge', this strategy will come under pressure in the near term. For larger data centres, such as many in the FLAPD market in Western Europe, the traditional business model is likely to continue to be successful for longer. Nevertheless, even here, it will eventually come under threat.

Let's explore these forces in a bit more detail.

Demand-side forces

Over the next decade or so, the range of local compute options available for companies will increase to meet their growing range of needs. The reasons for enterprise hybrid (cloud) compute solutions include:

- **Increases in automation generally, and in AI-driven activities specifically**, which require massive (hyperscale) compute to train Large Language Models (LLMs) and local low-latency compute for inferencing and acting quickly. AI is, therefore, driving compute growth at hyperscale compute facilities now and, as inferencing needs increase to meet a growing number of use cases including those involving AI agents, this will drive demand up for regional, network, and on-premises compute.
- **More stringent demands for sovereign compute solutions** from regulators meaning a growth in local public and private cloud solutions. We wrote recently about developments in the EU on data

© STL Partners

sovereignty and cybersecurity and the funding that the €1.2bn in funding that the EU has invested to encourage European companies to build a sovereign cloud and challenge the dominance of the US hyperscalers. We also touched on AWS's plans to invest €7.8bn in building a European Sovereign cloud. The is a further impetus for increased local compute infrastructure.

- The desire from many companies to **repatriate applications and data back from the public cloud** to other more local solutions. We covered the reasons for this in a recent article. These include compliance associated with data sovereignty mentioned in the bullet above, but also include cost and control concerns. Indeed, in a survey by Barclays Bank during 2024 83% of CIOs said they planned to move workloads back to the private cloud from the public cloud – up from 43% in 2020.
- **Concerns about sustainability** – both the carbon footprint associated with powering data centres and the water to cool them. Distributed compute – either at the far edge or in regional data centres – provides the opportunity to defray cooling issues by providing heat to residential, office, and recreational buildings (e.g. swimming pools).

Supply-side forces

Historically, by contrast to the consolidation inherent amongst the hyperscalers, the regional data centre industry has been relatively fragmented – particularly in Europe. There has been a handful of larger players, such as Equinix, offering internet exchanges globally but, generally, the market has comprised of lots of smaller local data centres meeting the compute and storage needs of local businesses. Asset-lite managed services providers have been responsible for 'knitting together' solutions from a range of data centres to provide national and regional offerings for larger firms.

But, again, this is starting to change. **As regional and edge compute is attracting more capital, data centre consolidation strategies are gaining traction, reshaping the market.** This is happening at a national level in Europe – for example Pulsant, in the UK, has grown to a portfolio of twelve data centres across the UK by snapping up individual sites. Consolidation is also starting to occur at a regional level – with German headquartered nLighten leading the charge to build a pan-European regional data centre 'platform' through acquisitions in Spain, France, Switzerland, Belgium, Netherlands, and the UK.

Consolidation is further driving greater diversity in local compute solutions because the price investors pay for an asset is not just driven by 'supply side' attributes (data centre capacity, expansion potential, power availability (which is particularly constrained in the FLAPD markets and likely to increase demand in tier 2 and 3 locations), sustainability concerns of investors etc.) but also by its commercial operations. Vendors can attract higher prices by differentiating their sites from competitors based on their portfolio of services and customers. Higher growth and margin solutions and customers will drive asset prices up. Colo and public cloud alone are unlikely to yield a premium price if market demand is shifting to highly interconnected private cloud!

New regional data centre strategy required

Hybrid compute options

At STL, we believe that these demand- and supply-side forces will lead enterprises to adopt a range of compute and storage solutions to meet their specific portfolio of use cases. We outline the five core infrastructure solutions that enterprises will adopt in their 'hybrid compute' strategies in the chart below. Four of these – colocation, single-tenant private cloud, multi-tenant private cloud, and public cloud are relevant to regional data centres. Private cloud will continue to grow rapidly and regional data centres have a strategic choice about whether they participate in this market. Investment will be needed, and success is not assured: for example, failing to provide a cloud management solution that works for managed service providers and enterprises will result in low take-up of private cloud. But not competing is self-limiting as the market for colo slows and hyperscalers continue to use their scale to squeeze margins in public cloud.

© STL Partners

Data centre strategy: core infrastructure solutions overview

DC core infrastructure options	Company own on-premises compute	Colocation	Single-tenant private cloud	Multi-tenant private cloud	Public cloud
What is it?	Company hosts its own compute and storage infrastructure on its own premises	Enterprise's own compute & storage infrastructure hosted in third-party data centre	Third-party data centre provides dedicated managed compute & storage service for the enterprise	Third-party data centre provides shared managed compute & storage service for the enterprise	Third-party (hyperscaler) provides shared compute and storage infrastructure for the enterprise
Applicable use cases	• Mission critical operations where need to scale is limited and the need for close, real-time control outweighs cost-savings, e.g. in-hospital patient monitoring, manufacturing control system	• Important applications where on-premises investment is not viable but demand could scale and company wants to retain control over compute, e.g. video analytics, disaster recovery	• Applications where cost management is important but responsiveness (inc. latency), customisation, & security/ sovereignty remain vital , e.g. manufacturing control system analytics, disaster recovery	• Similar to single-tenant but where perceived risks of deploying on shared infrastructure is lower e.g. cloud gaming, back-office IT systems, disaster recovery	• Business support where the need to reduce cost and manage scalability is compelling and (perceived) risk of data sovereignty is low, e.g. Chatbots, back-office IT systems, disaster recovery
Enterprise benefits	+ Company retains full control over policies and security + Ultra low latency + Data protection & sovereignty	+ Reduced premises capital and operating costs vs on-prem + Data protection & sovereignty	+ Lower compute infrastructure capital costs vs Colocation + Data protection & sovereignty + Savings on skills & people costs	+ Lower compute infrastructure costs vs Single-tenant private cloud + Data protection & sovereignty + Scalability – dial up and down	+ Lower compute infrastructure capital costs vs Private cloud + Access to hyperscaler tech + Scalability – dial up and down
Enterprise downsides	– High capital & operating costs – Difficult to future-proof: easy to be out of date – Tough to build in-house skills	– Compute and storage infrastructure costs remain high – Future-proofing compute & storage technology & building skills still tough	– Compute and storage infrastructure costs still high (dedicated infrastructure)	– Concerns about data security offered by smaller local players	– Reliance on hyperscaler business model & pricing – Data sovereignty risk
Regional DC benefits	+ Not applicable – no opportunity	+ Simple to sell and delivery – focus is on managing and maintaining facilities	+ Higher revenues – charge for compute as well as facilities	+ Higher revenues and opportunity for higher margins by managing multiple clients on shared servers	+ Utilisation: hyperscalers will take a large % of DC + Simple: one big client to manage
Regional DC downsides	– Regional DC sees no revenue – all kept in-house	– Likely to slow down and then decline in medium- to long-term as enterprises shift away from owning and managing servers	– Higher working capital – incur upfront capex costs which are amortised over contract	– Greater complexity – managing CPUs and GPUs, as well as multiple clients – Need for investment in cloud management platform	– Very stringent requirements from hyperscalers – lots of compliance – Low margin: prices squeezed

Source: STL Partners

© STL Partners

The need for a services-oriented data centre strategy

Many regional data centres currently don't offer private cloud, or it is a small part of their portfolio. Despite the risks of entering the private cloud market, ignoring these sectors and sticking to colo and public cloud will inherently limit revenue and margin growth. **Either way, regional data centres must consciously decide on their service portfolio in a way they have not had to consider before**. The compute world is becoming more complex for enterprises and this complexity requires a more nuanced commercial strategy from data centres. Management teams at regional data centres need to consider which customers they will serve and with what services. A regional data centre commercial strategy should seek to answer the following questions:

1. **What use cases to address?** For example, where and how to address the demand for AI applications?
2. **Which vertical industries to serve?** Horizontal strategies are likely to be replaced with a more verticalised approach as data centres need to show expertise which is relevant to the specific compute needs of their clients?
3. **What portfolio of services to offer?** Stick with colo and/or public cloud or embrace the growth (and complexity) of private cloud?
4. **How to optimise channel strategy?** – for example, how to work with managed IT service providers?
5. **How important is connectivity** For instance, a large amount of connectivity from multiple carriers is particularly important at the edge and for private cloud.
6. **Whether to venture beyond infrastructure?** Move into professional services and managing applications?
7. **What opportunities exist to acquire new assets with a complementary services portfolio?**

Chris Barraclough is the CEO at STL Partners specialising in strategy and finance in digital infrastructure. STL works with data centre clients and investors to drive organic and inorganic growth – including the questions outlined above.

Get in touch with the author to learn more chris.barraclough@stlpartners.com

Or visit STL Partners' Data Centre Hub https://stlpartners.com/data-centres/

October 2024

© STL Partners