INTACLOUD

# Best Practices for Selecting a Colocation Provider

## Summary

Third-party colocation facilities enable I&O leaders to increase data center space, minimize investment in on-premises data centers, and interconnect to digital business partners, cloud and service providers. I&O leaders must identify providers and facilities that align to their specific needs.

## Overview

### Key Challenges

- A growing proportion of enterprises are hesitant to commit long-term capital funding to data center facility investments, particularly when faced with the probability that workloads will move to the cloud.
- Organizations faced with the need for additional data center resources must choose whether to retrofit existing facilities, build new or utilize third-party space — called colocation.
- No two colocation data centers or providers are alike; thus, an in-depth understanding of selection criteria and differentiating factors is required for a proper evaluation in order to minimize the risk of poor choices.

### Recommendations

I&O leaders focused on infrastructure delivery strategies:

- Select colocation providers based on appropriate weightings of colocation differentiators using the filters and process described in Gartner's colocation selection process model.
- Avoid stranding capital or locking in to excess capacity in long-term contracts by properly sizing power requirements and factoring in technology trends such as hyperconverged integrated system (HCIS) and flash memory, as well as the expected widespread use of cloud in the future.
- Insist on flexible and innovative contractual terms, including the options of "ramping up" and "ramping down" power commitments over the life of the contract, to minimize the impacts of "overbuying" capacity.

- Streamline the selection process by establishing a clear view of requirements before beginning any market assessment. Frame basic requirements in terms of space (described by the number of racks or cabinets) and power (kW capacity at both a rack and an overall level), and frame advanced requirements in terms of interconnection capabilities and added value features (e.g., managed services).

# Introduction

Colocation suppliers own and operate data centers where they deliver space, power, cooling and services to many customers in the same facility, charging for services on a monthly basis. In the U.S. alone, nearly 2,000 multitenant data centers of various qualities provide services to customers, with some major metropolitan areas having as many as 40 or more providers competing for customers. We urge clients to engage with Gartner analysts to shortlist providers in selected geographies.

Buyers of colocation services are now commonly interested in far more than simply leasing space. In addition to traditional facility differentiators (e.g., power density and resilience), other services (such as interconnection capabilities and managed services) enable colocation to play a more strategic role in global infrastructure delivery and network design.

However, as no two data center facilities or providers are the same, I&O leaders must approach the colocation market using a structured approach that considers long-term data center strategy, opportunities for consolidation, and the use of a robust framework to determine appropriate suppliers and facilities.

This best practice research provides a three-step process and key considerations for selecting colocation facilities and vendors (see Figure 1).

Figure 1. Three-Step Process to Identify Suppliers Suitable for RFP

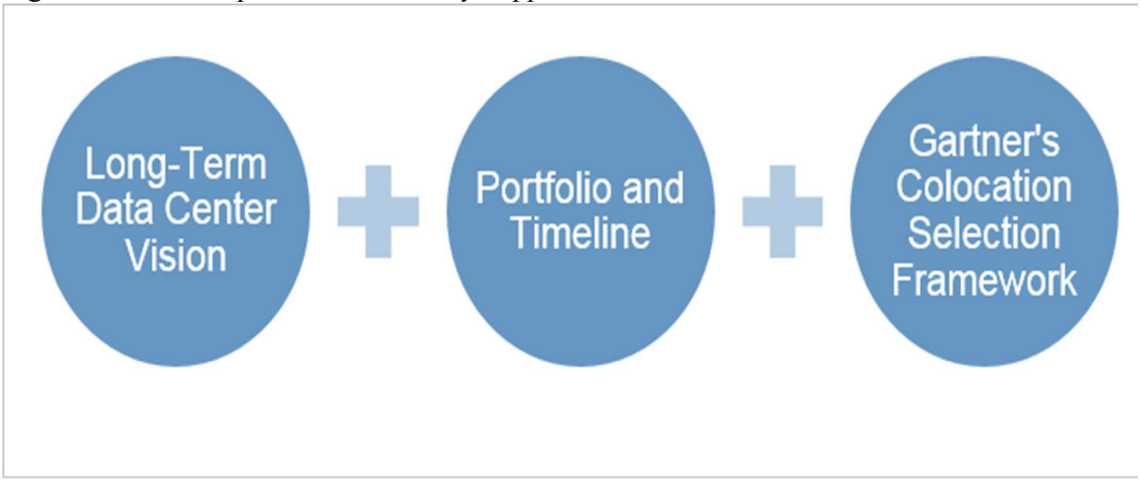


Source: Gartner (March 2017)


# Analysis

## Step 1: Clarify the Future Vision for Data Centers Within the Enterprise

Before selecting suppliers, build a clear picture of future intentions for data centers within the enterprise, so that colocation offerings can be contextualized within the "bigger picture" of business goals, and technology strategy. Tactical colocation sourcing may address obvious short-term needs, but it risks missing out on valuable long-term modernization opportunities needed to deliver IT strategy.

Examples of data center strategies that benefit from colocation include:

- Shifting mechanical and electric maintenance responsibilities for data center facilities to specialist providers.
- Minimizing commitment to customer-owned data centers as cloud strategies develop and on-premises needs fluctuate.
- Housing sensitive data outside public clouds for compliance purposes.
- Locating content and systems geographically closer to users for performance purposes.
- Improving availability of systems using purpose-built, highly resilient data center facilities.
- Changing data center costs to become operational expenses more closely related to actual consumption levels.
- Avoiding tying up capital funds in physical facilities where overprovisioning (and underutilizing) capacity is common.
- Locating customer-owned IT assets involved in hybrid cloud systems (e.g., data storage systems) close to cloud provider infrastructure for performance and compliance purposes.
- Enhancing enterprise networking capabilities — e.g., to improve performance through colocation-based low-latency connections to cloud providers and digital business partners.
- Accessing advanced provider services to support delivery of IT strategy, such as managed services and cloud-enablement capabilities.
- Improving disaster recovery capabilities by introducing data centers that meet business objectives for alternative geographies, supplier diversity and facility resilience.
- Creating "on-premises backup" of workloads in infrastructure as a service (IaaS) or SaaS.

## Step 2: Create a Data Center List and Timeline of Opportunities

When selecting potential suppliers and facilities, consider that consolidation enables opportunities for efficiency. Reducing the number of suppliers cuts down on activities such as contract, billing and supplier management. Consolidating facilities reduces overheads associated with deploying and maintaining multiple IT infrastructures.

Build a simple view of consolidation opportunities by creating a table of requirements per facility and per region. Include: Current data center locations and future geographies where data center services may be needed.

- A portfolio view of all current and future data centers, including qualitative metrics such as criticality; and quantitative metrics such as unit considerations (space footprint and cabinet count, density, depreciation cycle or contract end date) and price per unit considerations (price per cabinet, price per kW, volume discounts, annual escalators, etc.).
- Influencing factors leading to consolidation opportunities (e.g., capacity issues, lease ending, office move, multiple data centers in a single geography as a result of mergers and acquisitions).
- Key dates where colocation sourcing decisions are required (consider time for negotiating, commissioning and migration activities). Also track expiration and renewal dates to create a working portfolio of data center assets.
- Quantities of space (total number of racks) and power (total kW) required per facility.
- Interconnection needs to carriers, service providers, cloud providers and digital business partners.
- Managed services (basic [e.g., blended internet] and advanced [e.g., private cloud]).
- Growth trends and technology influences (e.g., anticipated cloud adoption) per facility, needed to build an overall projection of requirements for a three- to five-year period.

Supplier capabilities can be screened against the overall portfolio of opportunity in order to best understand where consolidation could be achieved.

## Step 3: Use Gartner's Colocation Selection Framework to Select Potential Suppliers

Gartner's colocation selection process (see Figure 2) guides clients through a series of steps, each of which is an important consideration when selecting suppliers, as explained below.

Figure 2. Colocation Selection Framework

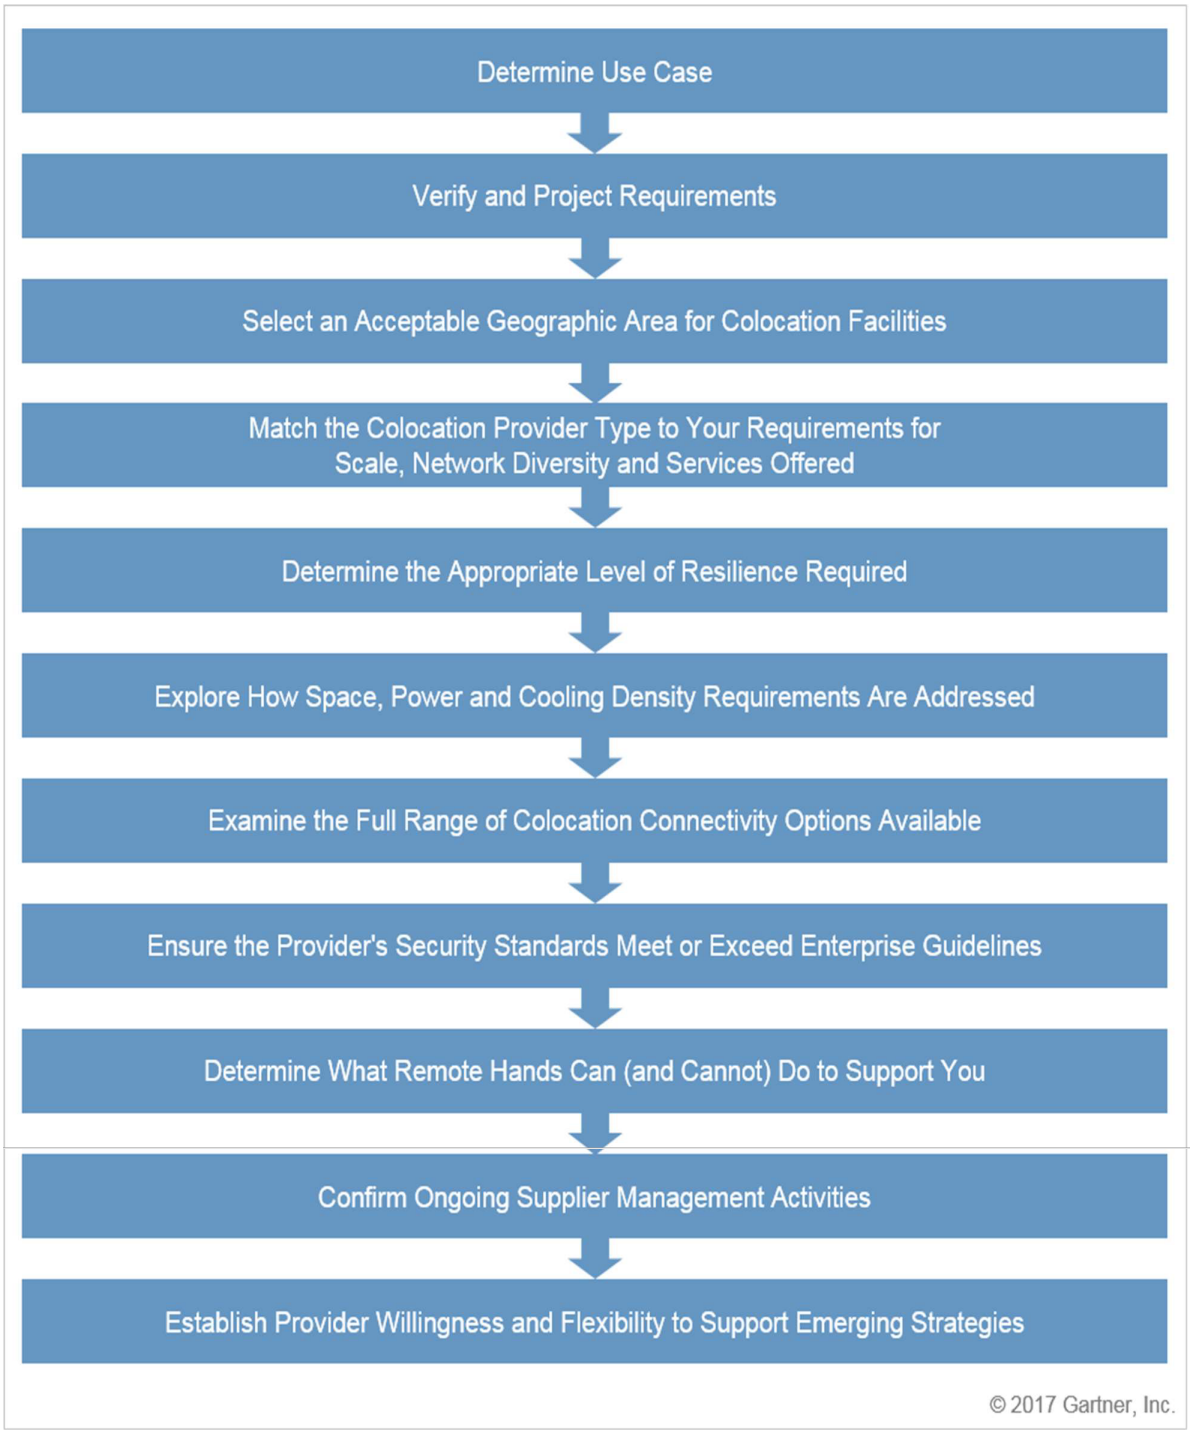


Source: Gartner (March 2017)


## Determine Your Use Case

One of the first things that need to be done when planning for or selecting colocation is to determine the intended use case or purpose of the colocation deployment. That is, "What services and capabilities are you looking for?" There are myriad differing service portfolios, ranging from simple space and power, to complex integrated solutions including space, power, managed services, private cloud and so on. A firm grasp of what purpose you intend the colocation deployment to serve — both nowand over the next three to five years — must be achieved first. At its simplest, we differentiate colocation use cases into those for space and power, or "server farms"; and those that are communications-enabling, such as interconnection sites, sites specializing and cloud onramps, etc. Colocation for server farms often caters to very large footprints, are priced on a per-kilowatt basis, and may offer free or low-cost cross-connects. Communications-enabling colocation sites, on the other hand, often cater to small, dense footprints (usually 10 cabinets or less); may be priced by the cabinet; and generally charge a premium for cross-connects. While there are many examples of providers offering some combinations of the services, it is important and useful to determine the primary goal of your colocation deployment in order to most effectively serve the most pressing needs. This determination is critical, as suppliers are generally good at one major use case or another, but not both.

## Verify and Project Basic Requirements for Power and Space

Core components of colocation contracts are total power capacity needed (measured in kW), floor space required (often estimated in number of racks or cabinets), and power densities required per rack (measured in kW).

Verify requirements by reviewing data center infrastructure management (DCIM) tools and/or onboard system management information, or measure the existing power draw for those systems on your data center floor using manual methods. If this is not possible — or if you are deploying new equipment to a colocation facility and are reliant on vendor specifications — many colocation providers suggest a general approach of adding up the maximum wattage value for all the equipment to be deployed in the rack, and then taking 50% to 65% of that value since most IT equipment doesn't run at "faceplate" power draw at all times.

Project how these requirements will likely grow or shrink over the lifetime of the agreement to determine the flexibility that will be needed in contract terms. Consider historic growth patterns, potential technology changes (e.g., increased use of virtualization, reduced power required by flash storage) and the influence of strategies such as cloud adoption.

## Select an Acceptable Geographic Area for Colocation Facilities

Establishing a wide range of acceptable locations for colocation facilities enables geographically linked opportunities, such as the ability to distance the data center from potential natural disasters. Colocation facilities can be distanced from users where applications are more tolerant of network latency, but also as a result of sophisticated remote hardware management capabilities.

However, each enterprise has unique needs, which will define acceptable geographies. Potential influences include:

- Tolerances for network latency between user communities, systems and partners.
- Cost of network links between user communities and colocation sites.
- The need to avoid natural and/or man-made hazards.
- The need for recurring staff access and the location of that staff.
- Ability to consolidate multiple data centers.
- Data sovereignty/compliance restrictions.
- Cross-border impacts — legal and commercial implications, language barriers and so on.
- Regional opportunities — less expensive power, climate difference (e.g., natural cooling) and labor rates.
- Desire to leverage locations in which facilities have network infrastructure of interest to the customer (e.g., cable landing stations with routes to countries of interest, containing a key point of network presence for the region, or having a particular ecosystem of trading partners contained within it).

Generally, the potential impact of the considerations above increases as the distance from users of the colocation site increases. However, this remoteness must be balanced against the potential benefits a large sourcing geography can bring in terms of a broader range of suppliers and facilities, and access to other regional opportunities.

Recommendations:

- Identify acceptable geographies based on your maximum tolerances, such that you can maintain a wide scope of potential suppliers.
- Examine the value of a supplier's regional knowledge and experience e.g., when expanding into new geographies.
- Analyze the practical implications of expanding into new geographies, such as the challenge of language barriers and the length of time to establish contractual arrangements. Leveraging the geographic reach of a supplier based in your home country might be preferable over a local provider.

## Match the Colocation Provider Type to Your Requirements for Scale, Network Diversity and Services Offered

The colocation market is largely segmented between retail and wholesale products, and further differentiated between carrier-based and carrier-neutral colocation offerings. Each model has distinctions that must be understood when determining the best fit for an organization's specific needs.

## CHOOSE BETWEEN RETAIL OR WHOLESALE COLOCATION

While the decision between retail and wholesale colocation has historically revolved around how much space and power the enterprise is looking to source, the lines between such models is blurring. A more relevant differentiator may instead be the level of support services required of the colocation provider. Retail providers may offer anywhere from single, locking cabinets all the way up to 3,000 or even 5,000 square feet, and up to one or two megawatts of power capacity. As part of the service, the provider will supply racks, deliver power circuits, build cross-connections, undertake cabling and include additional services such as remote hands. All of this is designed to make it simple for enterprises to install their equipment in racks, connect up to power and start using the service. Services beyond space, power and cooling are often focused on enhancing enterprise data center capabilities, such as offering rich connection options to partners and vendors (e.g., cloud suppliers) and layering on enterprise-focused managed services (which vary depending on the supplier's strengths).

Wholesale providers are a blend of commercial real estate and data center operators that have a portfolio of properties, most with proximity to larger technology and financial markets. They typically lease large quantities of empty "white space" to customers (more than the quantities above), with redundant power and cooling and little or nothing in the way of services. In some markets, wholesale providers have begun offering smaller footprints with additional services for shorter periods of time. In effect, they're offering "retail" space in an effort to monetize the "white space" or unused "chunks" of floor space that exist between the footprints of larger wholesale customers.

While the lines between retail and wholesale are blurring — with retailers looking for ever-larger deals and wholesalers looking to move downstream into the lucrative retail market and consequently offering additional retail-like services — the classifications and roles of market players remain largely the same. (For information about the blurring of lines between wholesale and retail markets, see "Four Dramatic Shifts in Colocation Offerings and Contracts." )

The majority (in total numbers of inquiries, not capacity required) of Gartner client inquiries regarding colocation relate to retail requirements, while wholesale build-outs often cater to complete outplacements of entire very large data centers, often for the largest enterprises or service providers, such as hosting providers or cloud providers.

Recommendations:

- Select wholesale providers for enterprise use when organizational needs exceed one megawatt of power, or when you require unique build-outs that retailers either can't or won't provide.
- Select wholesale providers for enterprise use in smaller deals, but be certain to compare the services offered with available retail provider alternatives.

## CHOOSE BETWEEN CARRIER-CENTRIC AND CARRIER-NEUTRAL COLOCATION

I&O leaders should seek competitive bids from a range of suppliers, including both carrier-neutral and carrier-centric variants. This helps obtain a view of overall pricing within the local market and an understanding of what providers consider their differentiating factors, based on where they are charging more.

The following sections explain essential differences between carrier-centric and carrier-neutral colocation.

### CARRIER-CENTRIC COLOCATION PROVIDERS

Carriers that own local and/or long-haul fiber assets constitute a shrinking portion of the colocation market. A key benefit of these offerings is the ability to connect to the carrier's services (e.g., MPLS network) through interconnection capabilities provided within the same facility, which eliminates the need for third-party WAN circuits. Common services provided include cloud computing, managed hosting and application hosting. Customers can enter into a basic colocation contract and eitherlayer on managed services when needed, or migrate to other offerings such as managed hosting services at a later date.

Enterprises reviewing carrier-centric offerings for retail colocation requirements should select this option when:

- The ability to connect within the facility to that specific carrier's network services is deemed advantageous (e.g., to reduce circuit requirements or to enable low-cost access to a range of additional services).
- Minimizing the number of suppliers for IT needs is an organizational preference (e.g., streamlined billing, single agreement, single supplier to manage, and so on).
- Local geographies are not rich with stable, alternative options (e.g., there is a lack of mature carrier-neutral facilities or it is a developing region).
- Leveraging the carrier is of interest for future regional opportunities (like delivering in-region network circuits, disaster recovery alternatives or services available through its local partnerships).
- The enterprise offers services that potentially complement the carrier's growth strategy or that of partners who are colocated within the same facility. Colocation could be used as a strategic tool when developing partnerships to deploy equipment and build direct connections (e.g., colocation becomes a "meeting point" to enable service provision).

When reviewing carrier colocation options, be aware that the quality of carrier-owned data centers can vary greatly. Some providers may have newer, purpose-built facilities designed for the needs of modern IT equipment, such as high densities of power and cooling; while others will have converted older telecom "central office" spaces into colocation facilities, which generally have much lower power and cooling density available. I&O leaders should also verify the availability of alternative carriers and service providers within these facilities, as these competitive offerings may prove expensive and limited in variety, especially when compared to a mature, carrier-neutral facility.

CARRIER-NEUTRAL COLOCATION PROVIDERS

Carrier-neutral colocation providers generally don't own local or long-haul fiber assets and, therefore, contract with their customers primarily for data center services — space, power, cooling and interconnections between the customer and service providers. A rich and often diverse range of carriers will reside within these facilities to provide internet and WAN connectivity to customers. Interconnectivity between organizations within the facility is provided by the colocation provider, by either routing a direct copper or fiber-optic connection between customers' secured spaces (if possible), or routing them through a Layer 2 switching fabric.

Enterprises are able to choose from a wide range of network carriers present in the facility. Some providers promote the ability for customers to directly connect with "communities-of-interest networks" within the same facility (or on the same campus), avoiding WAN circuit charges and minimizing network latency (which is often a critical differentiator when architecting for systems such as those involved in financial trading).

Differentiators of carrier-neutral offerings for retail colocation include:

- A rich range of carrier options (e.g., for diversity, large bandwidth requirements or for competitive pricing situations).
- High densities of power in highly resilient facilities.
- Access to ecosystems of beneficial partners within the facility.
- Direct connectivity to cloud providers and other digital business partners in the facility.
- An alternative to the supplier delivering their incumbent network (e.g., for disaster recovery purposes).
- A supplier for whom colocation is their primary product.
- Suppliers who do not focus on supplying large networks, but supply alternative services of interest to the enterprise (e.g., disaster recovery as a service [DRaaS] or managed hosting).

When considering carrier-neutral colocation, be aware that some carrier-neutral colocation providers are pure-play companies that only provide colocation services. Others may also offer internet connectivity services and their own managed hosting and cloud services in addition to colocation space — the maturity and scalability of such services should be thoroughly evaluated during RFP. In facilities with a high number of carriers present, colocation providers can often charge a significant premium over competitors — this is also seen when partner ecosystems are located within the facility.

## Determine the Appropriate Level of Resilience Required

Colocation providers may describe their facilities as having a "tier rating" of availability when discussing the redundancy they have built into their infrastructure. Two separate organizations maintain their own standards definitions — Uptime Institute and American National Standards Institute (ANSI). Each describes similar design standards and expected availability for a facility.

Gartner recommends clients consider the marketed tier rating (or tier equivalent terms) to be only a guide: Ask about the specific design, construction and operation elements that go into resilience claims. Relatively few facilities have been certified in comparison to the overall market, so keep in mind your own requirements for resilience to avoid limiting the range of providers to only those with certification status. Operators are usually well-prepared to discuss resilience design features with prospective clients.

While data center design and construction features are critical aspects of resilience, an in-depth understanding of the infrastructure in place and the policies and procedures of the data center may provide a better indication of likely resilience than a "Tier X-equivalent" marketing label. Gartner has seen proposals that positioned a provider as above Tier 4, while we also know of well-designed and -operated basic N+1 facilities that have coped with extreme weather events and regional disasters without outages. In addition to the discussion surrounding a facility's resilience is the broader issue of application and data resilience, such as applications where 100% uptime is not critical. Paying a premium for a 2N+2 facility where the use case doesn't demand it or where application resilience is handled in software across data centers may not be money well-spent.

Broadly, I&O leaders should expect pricing for highly resilient data centers (e.g., Tier 4 rating) to be above the average for the local market, especially in areas susceptible to weather events and natural hazards. However, opting for the very highest levels of resilience will usually result in a very limited number of facilities. Therefore, factor in the true level of resilience needed, as lower levels may still be entirely appropriate and available at a lower price point.

Recommendations

- Identify your own requirements for resilience (to be used in conjunction with official rating systems, or as an alternative where certification status has not been obtained).
- Factor in natural hazards (e.g., susceptibility to weather events) when deciding how to prioritize resiliency features.
- Look beyond marketing claims to determine the true history of facility reliability (e.g., the number of service issues and root cause). Review unexpected outages, near misses and externally influenced events, to obtain evidence of root cause analysis and completion of all recommended actions.
- Avoid overspecifying resilience when not required by the workload (e.g., test/development or research workloads).
- Assess whether disaster recovery capabilities and high-availability design built into applications (e.g., active/active split between data centers) make paying premium prices for extremely high levels of physical resilience redundant.
- Review previous and upcoming proactive maintenance plans for the facility to determine where resilience may be deprecated temporarily in the future.

INTACLOUD

# Explore How Space, Power and Cooling Density Requirements Are Addressed

I&O leaders must qualify the supplier's ability to support its per-rack requirements for power, space, cooling and overall equipment weight, in order to know that the facility design is capable of supporting their needs in an efficient manner both today and in the future.

Demanding high densities of power and cooling in facilities that were not designed to cater to such needs may result in customer equipment being distributed across more racks (and/or more power circuits), increasing the amount of supporting infrastructure (e.g., network equipment and cabling) required as well. Poorly planned oversupply of services to racks results in customers paying a premium for capabilities they simply do not leverage.

Data centers designed to deliver high densities of power and cooling across the facility are usually able to support 4 kW to 6 kW racks with ease, with 10 or more kW racks often possible. Low-density facilities may offer 2 kW per rack, with more power potentially available (by "borrowing" capacity from other areas to support customer needs) It is common for suppliers to match cooling requirements to the power density supplied to the rack, but specific cooling requirements can also be addressed. Temperature and humidity should be underpinned by a provider's SLA, and controlled within a range similar to that specified by ASHRAE for the operation of data processing environments (65 F to 80 F/18 C to 27 C, with relative humidity range of 30% to 60%).

The practical implications of expansion must be investigated to ensure that customer growth can be easily supported. Some expansions may be limited to space within very close proximity of existing equipment (e.g., when adding new disk shelves to storage arrays). HCIS or top of rack (ToR) switching may alleviate some overcrowding issues, but expansion into distant parts of the data center may be seen as unattractive due to the cost of deploying additional infrastructure (e.g., network equipment) or simply the inconvenience of managing multiple spaces. Providers are used to supporting growth requirements, and may offer "right of first refusal" options or other space reservation options.

Finally, the ability to measure and track data center utilization (e.g., power, space, temperature) directly affects your ability to operate in an efficient, cost-optimized manner. Poor understanding of utilization results in stranded (i.e., wasted) capacity and inefficiency. DCIM capabilities can be deployed by the customer, but such systems are also in use by colocation suppliers to manage the facility, and some of them may choose to make this information available to customers on a live or periodic basis.

Recommendations

- Ask the provider what standard power and cooling density is supported per rack across the entire facility, and compare this to your requirements.
- Request the provider specifies the limitations in weight, power density and cooling of cabinets, and whether there are "uplift" costs when specific tolerances are exceeded. Determine how nonstandard requirements would be addressed (e.g., multiple circuits, custom cooling).
- Confirm the provider matches cooling capacity to power capacity per rack.
- Verify what SLAs exist to support the availability and performance of all offerings (particularly power, space and cooling). But also review whether SLAs are available for installation and expansion services.
- Identify how suppliers are prepared to quickly accommodate requirements for future growth in a practical sense by focusing on space expansions, and the need to increase power and cooling without interrupting services.
- Clarify what data center rack styles are available (including height, width and depth), and any limitations on interrack cabling. Ask what restrictions are in place if customers need to bring their own racks or other custom equipment.
- Clarify what DCIM capabilities the provider is able to extend to you and how extensive the information is. Many are somewhat limited in their depth and breadth today, but one way or another, you will need to know how the provider can track utilization levels or at least power consumption on an ongoing basis.

## Examine the Full Range of Colocation Connectivity Options Available

The rich connectivity available in colocation facilities can enable new capabilities to drive transformation and enable digital business. Put another way, the concentration of network, compute and storage services possible in a colocation center, paired with the topology possibilities of engaging in multiple distributed colocation centers, provides the physical enablement of entirely new business models, such as edge computing. I&O leaders should ensure that facility selection criteria consider the possibilities in each of the three main facets of colocation networking:

1. Access
    - From enterprise users to colocation assets
    - From enterprise data center to colocation assets
    - From internet users to colocation assets
2. Intracolocation
    - Peering among enterprises and networks
    - Switching between enterprises and networks
    - Direct connect to cloud and hosting providers
    - Connection of enterprise assets to and between advanced business

- services based on SaaS or provider-owned appliances and cloud-based assets

3. Intercolocation
    - Colocation centers as interconnected nodes
    - Colocation assets to other colocation assets (e.g., disaster recovery [DR])
    - Enterprise-to-colocation-to-colocation-to-customer/partner/remote enterprise

Carrier-centric colocation providers, by the nature of their business, are highly capable of offering data center connectivity services. However, while alternative carriers are usually available in the facility (particularly in a facility that is mature), they may be more expensive than the carrier themselves.

Carrier-neutral facilities offer the enterprise the ability to invite competitive bids for service from multiple carriers, although the provider may only offer a very limited range of network products itself (such as blended internet offerings and connectivity to its other data centers within the metropolitan area).

Recommendations

- Review colocation interconnection possibilities to understand the potential of the three facets listed above (see "Colocation Networking: Connectivity Options to Drive Transformation and Enable Digital Business").
- Ask the provider about features or plans for enabling software-defined networking within the facility to enable dynamic interconnection with peers and services.
- Ask the provider about features or plans to support software-defined networking service providers external to the facility (e.g., Megaport, Console, Nuvem Networks) to provide flexibility in networking and cloud access.
- Identify connectivity requirements beyond pure access to the colocation facility, and evaluate the provider's offerings against your broader needs (e.g., cloud connectivity).
- Obtain quotes from multiple suppliers, and investigate costs and flexibility of adding additional bandwidth in the future.
- Determine physical resilience for network connectivity within the facility — e.g., do your preferred carriers maintain full diversity from the building entry pathways through to your secure area?
- Involve network architects in connectivity evaluation to ensure the capability, maturity, price and reliability of services and products are understood.


## Ensure the Provider's Security Standards Meet or Exceed Enterprise Guidelines

I&O leaders must determine what minimum security controls they require when selecting providers. Common offerings include single or multifactor access controls, 24/7 on-site security guards, CCTV, and advanced controls such as mantraps and biometrics.

Shipping and storage capabilities are important to understand (e.g., what loading dock facilities exist and how long equipment can be stored in goods-received areas), as not all providers offer online portals and electronic notifications, which prove helpful in simplifying these processes.

Providers may also comply with standards of interest to enterprise clients such as ISO 27001, ISO 9001, and Federal Risk and Authorization Management Program (FedRAMP).

Recommendations

- Ensure that security processes at a colocation facility can meet or exceed your requirements, and that suitable evidence is available in a timely manner to meet your compliance requirements.
- Define your expectations in the event of a security incident (e.g., CCTV reviewed within a 24-hour period) and ascertain how providers will adhere to these needs.
- Ask whether online portals are available to streamline routine access control tasks.
- Identify if compliance standards are required as part of your selection criteria.

## Determine What Remote Hands Can (and Cannot) Do to Support You

Colocation providers commonly offer "remote hands" services to prevent customers needing to send engineers to site for basic tasks, charging for such services via a monthly block of preallocated time, or via an hourly fee. Not all facilities offer remote hands, and the portfolio of services offered varies widely at those facilities that do depending on the extent of the data center staff's staffing levels and skills. Do not overestimate remote-hands capabilities — they are commonly supplemented with specialist vendor or third-party-based agreements for more complex break/fix and request fulfillment needs.

When discussing potential requirements for providers, differentiate between urgent reactive activities (i.e., break/fix) and regular request fulfillment (e.g., backup tape rotation). This can enable a better understanding of how activities are handled, as (for example) staffing levels may be reduced outside of core office hours, affecting which services are available.

For more advanced services, some colocation providers offer managed services for

infrastructure such as monitoring, network management and hardware break/fix. These services should be closely reviewed to understand the true extent of the service on offer and how it compares to other market offerings. The range, breadth and capabilities of such offerings are unique to each provider and extend beyond the scope of this research.

Recommendations

- Clarify the depth, breadth and availability of remote-hands capabilities, taking into account the implications if local engineering should be needed.
- Confirm how routine remote-hands tasks are handled as opposed to emergency incidents.
- Identify whether an SLA is available to support remote-hands requests, as attending to equipment can often be of the utmost urgency.
- Identify how remote hands are typically charged.

## Confirm Ongoing Supplier Management Activities

When selecting providers, it is important to consider the ongoing management of the supplier beyond just commissioning the service and establishing initial operations. To establish how easy and efficient services are to manage during ongoing operations, I&O leaders must evaluate elements such as:

- Billing — Does the supplier provide online invoice facilities, and the ability to obtain raw billing data for cost modeling and further analysis?
- Account management — Does the supplier regularly schedule service review meetings? What is the content of such meetings, and what reporting does the supplier bring to the review?
- Technical support and professional services — Does the supplier have on-staff expertise to assist with data center planning, efficiency reviews and technology integration?
- User-accessible portal — Is an online portal available within which all elements of the service can be measured, monitored and managed? Are any important interaction processes still handled outside of that portal (e.g., shipping)?

# Establish Provider Willingness and Flexibility to Support Emerging Strategies

As clients consider how future business and technology strategies may change their overall data center requirements, the supplier's approach to supporting the enterprise's changing demands must be ascertained. Historically, data center requirements rarely shrank; however, the reality for many enterprises is that technology strategies — such as virtualization, flash storage, hyperconvergence and cloud-first hosting — all introduce levels of uncertainty over demand for future data center needs. For many IT leaders, establishing multiyear contracts without the flexibility to contract or divert spend into alternative services may be inappropriate.

While colocation pricing options are not as dynamic as those of hyperscale cloud providers, contractual flexibility is a topic that is frequently raised by clients seeking assistance from Gartner analysts. This topic should be pursued with potential providers to ensure that colocation costs remain optimized as requirements change over time.

Providers are often willing to negotiate contractual terms with some level of flexibility in the makeup of commitment levels if they can predict and protect a reasonably sized proportion of the overall contract value. Enterprise clients often meet this need by identifying a predictable basic requirement for colocation from systems that are unlikely to be affected by new technology strategies in the short term, and by considering a staggered approach to introducing flexibility over the colocation agreement time frame.

Recommendations :

- Establish how changing technology strategies may introduce the requirement to contract additional services over time.
- Be prepared to establish minimum commitment levels of power, space and contract duration, on top of which more flexible contractual terms can be built.
- Explore what options suppliers propose for contracting requirements (e.g., diverting spend into alternative products, having minimum commit levels for services and so on).
- Investigate what suppliers can offer beyond pure colocation. What are they doing to support clients who are already traveling on similar journeys in terms of services, support and flexibility?
- Focus on pricing and minimum commitment levels during RFP stages to determine the cost of flexibility where highly elastic colocation models are proposed.
- Identify whether ramp-up periods are available to minimize unused capacity as agreements commence.