

**GLOBAL
TECHNOLOGY
SOLUTIONS
GROUP**

**COLOCATION AND
CONNECTIVITY IN
HYBRID CLOUD**



While the largest number of colocation customers remain focused on space and power, forward thinking entities recognize the connectivity model.

For the majority of organizations past even a modest level of scale, hybrid cloud will be the model for the foreseeable future. "Hybrid cloud" for the purpose of this paper is defined as the sourcing of workloads from, and requiring integration among, a complex configuration of multiple data center, colocation, cloud and edge environments.

For many of these organizations,

- the edge will be key – collecting data, processing and sending back to the core
- latency will play a key role in determining where processing data may reside

For all, cost is a key consideration.

While the largest number of colocation customers remain focused on space and power, forward thinking entities recognize the connectivity model, the performance and latency considerations which accompany -

- The processing of on-prem data in cloud- native applications;
- The movement of data from the edge to the core;
- The integration of multiple clouds.

SETTING THE STAGE: IT'S *STILL* ALL ABOUT THE WORKLOAD

Workload Placement Analysis drove 30% cost savings on the 84% of workloads that did not require a direct connection.

We recently visited with a colocation provider at their facility outside a major metropolitan area. They told us about a client who had a sourcing decision to make involving 50 distinct workloads. This provider was about to lose the business because their competitor had an established connectivity link which was appealing to the customer. The account executive persuaded the client to consider how many of these workloads actually required this connectivity... in other words, to do the analysis to **define their requirements**. The client went back and did the homework and found that only eight (8!) of the 50 workloads required the more expensive solution.

The price delta between the two? 30%.

The answer is always, always, always in the detailed analysis.

THE CHANGING DATA CENTER LANDSCAPE

Gartner tells us that more than 80% of large enterprises in North America will utilize colocation providers by 2023. This will not necessarily replace their on-prem DCs, or even become their primary data center, but may provide additional capacity, geographic expansion, or disaster recovery. The reduction in the number of data centers is concentrated at the smaller end. The number of both midsize DCs and racks/computer rooms will be down 10% over the next several years. However, large and enterprise DCs will grow: hyperscale providers are on a trajectory to increase their capacity by more than 50% from 2018-2023. And in the next two to five years, a large percentage - estimates vary, widely - of enterprise-generated data will be created and processed outside the data center or cloud. Current circumstances (Summer 2020), if anything, amplify all of these considerations.

FACTORS DRIVING THE CHANGE

CORE VS. CONTEXT

For years, and increasingly, some number of firms seek to “get out of the physical data center business” simply because it’s neither a marketplace differentiator, nor a core competency, and the industry has continued to provide better (and even more flexible) solutions.

CLOUD-DRIVEN UNCERTAINTY

More recently, as organizations have looked to the cloud, uncertainty has increased as to what the requirements will be five years from now for the enterprise-managed data center. The investment required to keep them current has raised questions from finance, and rightfully so.

CONNECTIVITY

The largest segment of the end-user marketplace is still focused on “colocation value proposition 1.0” – traditional space and power. However, forward thinking entities are focused on the connectivity model, on the performance and latency considerations which accompany the processing of on-prem data in cloud- native applications, bringing data from the edge to the core, and integrating multiple clouds.

In conversation with clients we’ve noticed that some conflate the terms “colocation” and “connectivity model” as they seek the right answer on how to manage and pay for the movement of data to and from various processing and data storage locations. It depends on the provider: only a few of the colocation providers even aspire to become the connectivity hub.

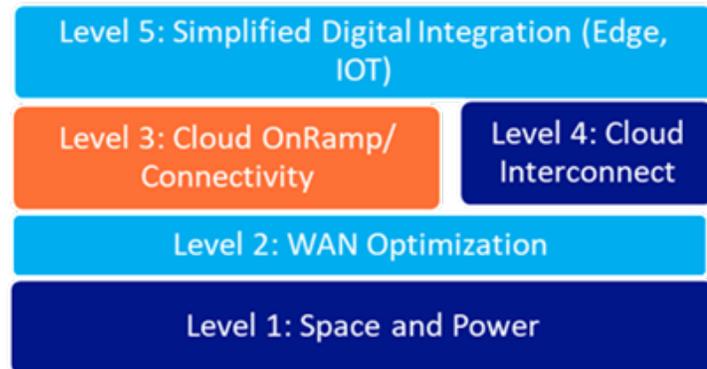
THE PROVIDER LANDSCAPE

At this point we'll note that GTSG has no interest, financial or otherwise, in promoting any provider: colocation, software, hardware or cloud. Our independence in making recommendations to our clients is more important to us than any commission we might earn from a third party.

Having said that, there is a difference in strategy among the providers: some have made substantially more investment than others in meeting the connectivity requirements their clients organizations will place on them going forward.

Even as the number of cloud workloads grows, and as the number of midsize and small data centers decline, some workload will need to be enterprise - managed for good reason (e.g., data residency, security, absence of a suitable cloud alternative). The chart lays out the levels of value which providers can bring to their clients:

There are stark differences in regional footprints, available connections, and pricing models. The answer lies in the analysis: old-school "rules of thumb" no longer apply in today's complex hybrid environment.



Within this model for looking at the market, there is an evolutionary path to the top tier. However, there are providers who don't even aspire to get past level one.

On-Prem vs Colocation

There are several compelling reasons to require on-prem or colocation in your service delivery portfolio, along with cloud: data-residency requirements, high security or stringent audit requirements, performance or latency, unsupported legacy equipment, or simply the absence of a business case to optimize applications for the cloud.

Once the decision is made to retain an enterprise- managed infrastructure, the “build (or, retain and reinvest) vs buy” decision must be made. Colocation offers a path from capital investment to a more predictable operating expense-based model, both for scalability, and in many cases a superior redundancy architecture.

One client hired us to do two projects concurrently.

- Help them construct their workload placement strategy- which workloads would be moving to the cloud over the next five years, and which the enterprise would need to continue to manage; and interdependently,
- Help them assess the level of investment required to keep their physical DC up to standard during the transition period.

On Latency

We have a client looking, in effect, to relocate their primary data center to one of the major cloud providers.

The “roadmap” distance is a bit over 200 miles. Of course, network paths don’t follow road maps: the number of hops, differences in link details and each intersection can impact perceived speed. Ping times are only part of the picture, bandwidth can’t solve every problem – and of course network is not always the issue in the first place. Chattiness can outweigh distance; database tuning can be crucial, and some data can’t be compressed.

On Pricing Considerations

Simplicity vs. transparency: An all-in price per KW is convenient to manage, but if a DR site sits largely idle, a metered approach may be better than a flat rate, making pricing transparency the goal.

Capacity Negotiations: Had historically been focused on the ability to add contiguous space – today, it’s often about the flexibility to downsize in the event significant workload should move to the cloud. Gartner reports that many of the clients they speak with are significantly overprovisioned.

The Workload Placement Analysis

A couple of years ago, Gartner wrote a research note (in which GTSG is cited) entitled “Developing a Practical Hybrid Workload Placement Strategy.” The premise of the note was that in order to construct a workload placement strategy worthy of the name, the analysis needed to understand the interdependencies among applications, databases and services just as well as we look to understand them before a major migration. This is necessary not only to construct least-risk move groups, as before the move event, but in order to understand where there is a risk of performance degradation – not only before the move, but before the decision.

Testing Center of Excellence

For many years, before major migrations, GTSG has utilized a “testing center of excellence” which requires a baseline of the performance of the workloads in question so that post-migration performance can be accurately compared. Today, we apply the same expertise to predicting the change in performance based on the new data center (whether on-prem, in cloud, or in colocation). Clearly, organizations should understand the level of performance they will be paying for well in advance of making commitments to providers.

Recall our story from the top of this paper: only 8 out of 50 workloads required the direct connect offered by the more expensive provider.

This is work GTSG does all day every day, for some of the biggest names with the lowest tolerance for latency in the United States.

As edge data becomes more important, the ability to predict latency will grow in importance as well. Our skills have been honed over decades.

The Colocation Provider Selection Process

If it's been a while since you've been through this process – most organizations don't have to select sites or partners very frequently – you'll want to consider some changes to the approach.

- Selection teams are expanded to represent applications, network, security and business continuity specialists.
- Resiliency planning must be part of the process from day one. GTSG performed a remediation engagement for a major healthcare provider after significant data center migration activity that had rendered their DR plans meaningless.
- You'll need to decide whether you want one or multiple providers
- You'll need to negotiate to balance the provider's need for commitment with your need for flexibility.
- Your candidate providers will need to demonstrate an ongoing commitment to investment; we'll discuss that in the success story just below.

A Living Strategy

One of the points we like to make is that there is no permanent answer. Jack Welch famously defined strategy as resource allocation: any cloud, hybrid workload placement, data center or infrastructure delivery strategy needs to allocate scarce resource to the effective delivery of compute power to the business. Who "owns" which components is irrelevant; the job is to provide, not to produce.

Getting to the Right Solution: A Success Story

Some time ago, GTSG was engaged to set the workload placement strategy for a client considering a long-overdue move from its primary and DR data centers, was at the outset of a cloud journey, and wanted a plan for what workload could move from the mainframe and what needed to remain there.

During the engagement, the client came to us to say that their RFP process for the Colo sites required help. Our team jumped in, brought our methods and templates, provided a list of candidate facilities in their required geographies, worked with the client through the process, and helped to evaluate and score the responses.

At the end of this process, the team took its recommendation to the CIO, who thanked them for their efforts- and then promptly directed them to a non-finalist. Stunned, the team went off to plan its site walkthrough.

When the team returned from the site visit, they didn't know what to say. The site was in disrepair and would not have satisfied this (or any other) executive.

GTSG helped them with yet another alternative provider- through our contacts in a marketplace in which we remain deeply engaged- and when we signed off, the client had moved ahead with a selection.

The Vision to See Where This is Headed

The previous story indicates a level of process definition, execution discipline, and marketplace experience and understanding required to make a good decision. Without the first two, we couldn't have finished on time – but without the experience and understanding, we could not have helped the client recover from the setback that the site visit turned out to be.

GTSG remains committed – permanently – to working with colocation providers, understanding the marketplace, and consulting with industry experts to bring our clients the best possible perspective going into a decision.

If you'd like to discuss how we can help, we'd be delighted. Please reach out to us at Partners@GTSG.com. Thank you.

If you'd like to discuss any of these topics further, please reach out to PARTNERS@GTSG.com.

GLOBAL TECHNOLOGY SOLUTIONS GROUP

T 877 467 9885
F 877 225 4084
W GTSGSERVICES.net
E Partners@GTSG.com

| | |
|---|--|
| <p>HYBRID CLOUD STRATEGY AND MIGRATION</p> <p>Strategic Approach</p> <ul style="list-style-type: none"> • Business case development • Transition planning • Technical modeling • Non-disruptive execution <p>Application Analysis Methodology and Tools</p> <ul style="list-style-type: none"> • Decomposition • Affinities • Wave planning <p>Project Leadership</p> <p>Implementation Subject Matter Expertise</p> | <p>INFRASTRUCTURE TRANSFORMATION</p> <p>Transition Services</p> <ul style="list-style-type: none"> • Insourcing/Outsourcing • Knowledge transfer and interim support • Application migration • Service management design <p>Disaster Recovery Design and Implementation</p> <p>High Availability Design and Implementation</p> <p>Application Assessment and Deployment</p> <ul style="list-style-type: none"> • Reference Architecture • Infrastructure Alternatives/Recommendations • Implementation/Migration |
| <p>INFRASTRUCTURE SUPPORT SERVICES</p> <p>Managed Services</p> <ul style="list-style-type: none"> • Multi-platform including DB & MW • Service-level based or FTE-based • Architecture, administration, programming, systems management • Remote or Onsite <p>Project Based Services</p> <ul style="list-style-type: none"> • Platform upgrades • Workload migrations • Implementation services • Consulting and Assessment (performance, DR, HA.) • Project Management | <p>INFRASTRUCTURE OPTIMIZATION</p> <p>Architecture Assessment and Design</p> <p>Server Virtualization/Consolidation</p> <p>Storage Optimization</p> <p>Data life-cycle management</p> <ul style="list-style-type: none"> • Tiering • Standardization/Automation <p>Application Decomposition Application</p> <p>Re-design/Remediation Performance</p> <p>Management and Tuning Latency</p> <p>Analysis and Consulting</p> |