



Straight talk: Colo Data Centers as a part of your DR strategy

By Bob Graney

Database clusters?	✓	Fully redundant SAN?	✓
Fully Meshed network?	✓	Redundant power & cooling systems?	Not checked.

What to do if your highly redundant network, servers and storage are supported by non-redundant data center infrastructure? What are the options?

A “Colo,” short for colocation facility, is a data center hosted by a vendor where multiple customers locate network, server and storage gear and connect to a variety of network service provider(s). Typically it is a hardened facility with multiple layers of security, and highly redundant power and cooling systems. A “pure colo” relationship is where the vendor *only* provides power, cooling, data center space and the related infrastructure typically found in enterprise data centers (security systems, fire detection/prevention, etc.) Most colo vendors also offer a suite of additional services that can be layered on top of the pure colo services which we will get into later.

The colo market is a growth market and many organizations are moving technical resources to a vendor’s data center for Disaster Recovery (DR), production or both. This paper will take a look at what the drivers are for organizations that are “going colo.” We will then take a look at shifts in the IT landscape that have helped “enable” colo migrations, followed by a few case studies and finally we will take a look at the critical considerations in your colo vendor selection.

Drivers for considering a colo facility in your data center strategy.

Most companies already have one or more production data centers and many organizations also have a DR site. So, why is there a big shift today to colo for either your production or DR data center, or both? Below is a list of some of the business drivers companies have cited in the past 12-18 months. If any of these factors apply to you, then read on to find out more on how you can learn about and leverage colo facilities.

- *Your current DR capability is “slim to none” and the business need has now grown to the point where you need a real DR data center*
- *Your production data center has space, power and/or cooling constraints which is restricting application growth.*
- *Your production data center no longer meets redundancy requirements for power/cooling and there are one or more critical single points of failure (i.e. there is no generator) which keeps you up at night.*

Straight talk: Colo Data Centers as a part of your DR strategy - Graney

Continuity e-Guide “Meet the Experts”

June 23, 2010, Issue 340

- *The hot site “shared” model for DR no longer will meet the projected requirements. In other words, the recovery strategy from a few years ago where you were given 2-4 days to recover is now less than 24 hours.*
- *Your production data center is sitting in the same building where most of your people sit and the thought of losing the data center AND all of your staff workarea at the same time is making you queasy.*

These are by no means all of the factors in considering adding colo to your data center portfolio. In the past these factors would have driven organizations to secure the capital funding required to add on data center space or even build a new data center. In today’s market such a huge capital expenditure is a difficult proposal for executive management to support. A colo solution moves the capital required for data center infrastructure (e.g. raised floor, UPS, generators, chillers and CRAC units) to an operating expense. This infrastructure is now provided to your firm as a service.

Enablers for the move to colo

There are a few shifts in the corporate IT landscape which help enable this move to colo space by allowing data centers to be remote and many miles away from corporate locations. This opens the door to the idea of using external third party data centers to house mission critical applications, servers and storage.

- *Wide Area Network (WAN) connectivity options are both cheaper and easier. The costs for bandwidth are going down and Metro Ethernet services reduce the equipment required to get connected. Find out what vendors are “lit” in your building and may be able to get a 10MB Ethernet handoff for less than an existing T1 costs now.*
- *WAN Accelerators are network appliances that can dramatically reduce bandwidth requirements by increasing transmission speeds anywhere from 5 to 10 times faster than non-optimized lines. WAN optimization has moved into the mainstream of corporate IT and has greatly enhanced/expanded WAN usage.*
- *There is a concentration of network capabilities in a colo facility that serve as a magnet for content providers, video distribution and exchanges forming communities of interest. Major colo facilities with scores or hundreds of network POPS provide the lowest prices and most choices for finding vendors that you need to reach. This can not be recreated in an enterprise data center.*
- *The number of IT staff that needs to “touch the server” is notably less than it was several years ago. Server administrators remotely make configuration changes or install software that in the past required a visit to the data center. Even the morning ritual of swapping tapes is less of a factor today with data replication and disk based backup strategies moving data between data centers. In larger organizations today with multiple data centers, most DBAs do not know or even care where the actual box sits.*

Case Studies:

Below are three snapshot profiles of typical companies which are considering or have recently made the move to colo facilities.

Manufacturing

The IT team for an international manufacturing firm performed an analysis for site selection of a DR data center. Backup tapes were going offsite and the IT team knew that a real DR solution was needed, but were having difficulty with selling DR to the business. The project team developed the network, server and storage requirements and looked at internal versus external sites for building out a DR data center. Critical business applications had a 4 hour Recovery Time Objective (RTO), so a hot site solution quickly moved down the list of options as a DR site for the more critical applications.

All of the internal sites required upgrades in raised floor infrastructure (power/cooling); however it quickly became apparent that even with those upgrades an internal data center would not approach the power/cooling redundancy levels found in most colo facilities. The primary iSeries platform was already outsourced to a vendor, so they had an existing relationship with a vendor that offered colo space. The financials, as well as other soft factors like staffing and training, clearly supported the colo option and the decision was made.

Once that decision was behind them, the IT team quickly realized that the DR site was an enterprise class facility with dual power feeds, redundant UPS systems and redundant generators. Meanwhile the production data center back at corporate headquarters had only one generator and one UPS. In addition, the production data center had cooling issues and was in the basement so flooding was a concern. Clearly, the priorities for resiliency were reversed! While the immediate goal was to build out DR in the colo, it was expected that at least some production systems would start moving to the colo within 18 months in order to take advantage of the higher levels of resiliency in the colo data center.

Financial Trading

An international financial trading firm was developing a data center strategy for its US operations. Acquisitions had resulted in multiple data centers and the primary production facility was maxed out on space, cooling and power. Moving development and QA environments to another floor did buy the firm 18 months, but a long term strategy was needed. The IT team looked at expanding the primary internal facilities versus moving either production or DR out to a colo facility. All options were put on the table, but the analysis showed (a) management was not keen on the required capital expenditure needed to expand the primary data center and (b) production should be placed in the more resilient data center, the colo. The resulting strategy was to move production out to the colo and use the existing production data center for DR. Combined with a server virtualization strategy, it is anticipated that development and QA can also be housed in the current production facility.

Medical Services

A medical services institution has two urban data centers within close proximity of each other. Both data centers have a combination of cooling, power and/or space constraints, but IT demand is growing exponentially. In addition, each of the data centers had significant gaps in power/cooling redundancy which more and more represented too much risk. The IT team developed a DR site strategy within the context of a long term data center strategy for the hospital. The options considered were (a) expanding the existing data centers, but this was a non-starter as it did not deal with the close proximity issue of the two facilities, (b) another internal location on campus, (c) the use of a hot site vendor for DR, but this was not preferred as the recovery time objective, or “RTO,” for mission critical applications was less than 24 hours and (d) moving one or both data centers to colo space. While still in the early stages, the long term strategy calls for a new internal production data center to be included in

construction of a new building on campus. As this is on a three year horizon, the IT team is also planning on securing colo space within 12 months to handle expansion of DR services.

Important considerations in picking a colo vendor

There are many factors in the consideration of vendor selection for a colo data center which will be specific to your industry and your organization. For example, certain financial segments are super sensitive about latency between themselves and customers and/or brokers. These folks look at 40 milliseconds as equivalent to using the Pony Express and now talk in terms of transaction roundtrips in microseconds. Other industries are primarily focused on Internet response time because that is how their customers get to them. So with the caveat that the following list needs to be adjusted to your industry and situation; here is the list of factors to consider for a colo vendor.

Power

- By far the largest cost component of a data center Total Cost of Ownership (TCO) is power and your organization may want to evaluate the true cost of power. One thing to be aware of is that most colo vendors charge for power allocated and your monthly bill will be the same for an empty rack as if it was fully populated.

Resiliency

- Do you need a Tier 3 or Tier 4 data center or will Tier 2 suffice? The Uptime Institute methodology is considered an industry standard in rating the level of infrastructure redundancy in a data center. A simplified view of that methodology is shown in the table below and one basic question for any potential vendor is “What Tier is this facility?” Please go to www.uptimeinstitute.org for details on data center infrastructure.

Tier	Level of redundancy	Planned Outages (Hrs/Yr)	Unplanned Outages (Hrs/Yr)
Tier One	No component redundancy	24	4.8
Tier Two	N + 1 component redundancy in a single delivery path/system	22	4
Tier Three	Component redundancy with one active, one passive delivery path/system. Zero downtime for infrastructure maintenance	0	1.6
Tier Four	Dual independent delivery path/systems with N+1 redundancy in each. No SPOFs. Zero downtime for infrastructure maintenance	0	0.8

Vendor Reputation & Financial Stability

- You want to make sure they will be around in 5 years.

Global Geographic Breadth

- an international company may prefer an international vendor, no?

Geographic DC Proximity

- Distance between production and DR must be considered, often required to be a minimum number of miles, balanced against replication and/or latency considerations which limit the distance between data centers.

Client Proximity

- Financial firms need to drive down latency between their systems and client systems as well as exchanges, trading partners and market data providers. Think black box talking to black box and the need for speed that wants your server as close as possible to the other server.

Carrier Neutral

- Consider your network connectivity options and compare your list of network service providers against the providers that are “lit” at the colo being considered.

SAS 70 Type II Certification

- For most organizations this should be a given. Get the SAS 70 reports.

Competitive Price

- Price is king, right?

Additional data center services

The other factor in vendor selection may be the availability of enhanced services from the colo vendor. So far we have only discussed pure colocation vendors, but some vendors offer other important service areas, such as network services, managed services, hosting services and complete outsourcing or fully managed services. A breakdown of the different types of managed services is shown in the list below.

- ***Data Center Hosting Services*** – server and/or application hosting
- ***Client Proximity Hosting*** – connect to specific clients, markets or exchanges
- ***Managed Services*** – monitoring, data backup etc.
- ***Network Services*** - Internet bandwidth, private lines etc.
- ***Migration Planning & Professional Services***
- ***Disaster Recovery***

Cloud services is one of the new buzzwords in the IT industry which encompasses various pieces of the services offered above, traditional Application Service Providers (ASP)s and true “cloud” services with real time provisioning of virtual resources. Cloud services may be a factor in your colo decision, but the “cloud” is a whole other discussion that deserves the full treatment next time.

Straight talk: Colo Data Centers as a part of your DR strategy - Graney
Continuity e-Guide “Meet the Experts”
June 23, 2010, Issue 340

Conclusion

Colo vendors come in all shapes and sizes and customers can contract for one, ten or 1000 cabinets. There are a host of factors that should be considered and the priority of those factors change depending on your firm’s industry, size and exact requirements. Your company’s need for additional services will be another consideration when starting down the path of “going colo.” Good Luck!

About the Expert

Bob Graney is the Practice Director for Eagle Rock Alliance’s (www.eaglerockalliance.com) Technology Solutions Group. He has over twenty-five years experience in IT infrastructure and Disaster Recovery. He can be reached at bgraney@eaglerockalliance.com.

