



THE INFORMATION SOURCE FOR THE DATA CENTER INDUSTRY

Data Center Knowledge Guide to **Colocation Selection**

*Best Practices and Critical Considerations for Choosing the Right Data Center
Colocation Solution*

By Bill Kleyman,
Cloud and Virtualization Architect
October 2012

SUMMARY

Today's modern IT infrastructure is demanding more out of its resources, expansion strategies, delivery methods and disaster recovery methodologies. To help meet these new growing demands, many organizations are turning to colocation providers to help service their data center needs. But what makes up a good data center decision? Is it just the location? What about the bandwidth coming into the site? What about green technologies and the adoption of new unified technologies? The reality here is that there are several key consideration points which need to be made when choosing the right data center colocation solution. Now with cloud computing and a truly distributed, data-on-demand, environment, making the right choice has become more important than ever.

INTRODUCTION IT'S MORE THAN JUST LOCATION

Over the last decade, corporate IT infrastructure has played an increasing role in everyday business functionality. More organizations are relying on their IT systems to work harder, faster and more efficiently. This push towards an always-on, always-available environment created a growing market for data center colocation. Commercial and government entities are spending millions of dollars to ensure that their environments are capable of expanding directly in line with the needs of the organization. As the technology evolves, these companies have to make the decision as to which data center provider they wanted to work with. This is where many managers and administrators run into some design and planning challenges.

This guide will seek to identify the best practices and considerations which go in to making the right decision for your colocation provider needs.

Contents

SUMMARY.....	2
INTRODUCTION – It's more than just location	2
SECTION 1: Location – Safe, secure and accessible.....	3
The Physical Location and the IT Staff	3
Design, Build and Operate.....	3
Weather Patterns.....	4
Seismic Risk.....	6
Power Grid and Redundancy	7
Accessibility – Routes, roads and airports ...	8
SECTION 2: Facility – Critical infrastructure that grows with you.....	9
High-Density Metered Power	9
Cooling	10
Physical Security	10
Data Center Infrastructure Management (DCIM).....	11
Amenities.....	11
Green Technology.....	12
SECTION 3: Analyzing the Workload – What is being delivered and where?	13
Cloud and Big Data.....	13
Mission Critical vs. Good Enough	14
Network Bandwidth and Latency	14
Balancing the Workload, Continuity and Disaster Recovery	15
SECTION 4: Environment Control and Management – Good SLAs make good neighbors.....	16
Working with a Service Level Agreement... ..	16
Maintenance and Testing.....	16
Having a DR-ready Contract	17
Using Management Tools.....	17
Issue Resolution and Communication.....	18
SECTION 5: Colocation Decision Checklist and Best Practices.....	19
Avoid "Cookie-Cutter" Data Centers.....	20
Customer Experience is King	20
Know What it Costs Should Your Environment Go Down	20
Educate the IT Team.....	20
Understand the Uptime Institute Data Center Tier Model.....	21
CONCLUSION.....	22

SECTION 1

LOCATION – Safe, secure and accessible

As one of the more important sections in this guide, examining the requirements for the location is vital to the selection of a good colocation data center.

Making decisions based on distance alone is no longer enough — especially when choosing such an important part of an organization's infrastructure. As discussed earlier, over the past few years, more variables have entered into the decision making process. This section will analyze these variables and where IT managers and administrators should focus their attention in deciding the right location and facility for their IT environment.

The Physical Location and the IT Staff

When selecting a location for your colocation data center, an important planning criteria will be addressing your IT staff including application developers, data base administrators, and your CIO — not just the data center team. In some cases, you'll want to consider staff response time when choosing a location.

- **Local IT support.** Many times organizations will select a data center without actually considering the proximity to what is known as a "Response Team." Uptime within a datacenter infrastructure is always important so the proximity of the response team is very important. Also important is access to the facility. For example, your staff may be near the facility, but if there is only one access road your systems could be at risk. Depending on the contract, some organizations choose to manage their own hardware and IT workloads within their colocation environment. This support structure means that there has to be reliable IT staff available within 'X' miles to provide immediate support.

- **Outsourced ("Remote Hands") IT support.** In some situations, many colocation data centers offer a "remote hands" service, where the provider maintains a staff of trained IT technicians who are available for maintenance items and operational support. While most colocation customers will still want to maintain an on-site staff member or two, using these services evens the playing field concerning the location factor in the colocation decision process.
- **Application developers and data base administrators.** In today's IT-centric companies, the application is the business. And as data becomes mission critical, DBAs (data base administrators) are joining the front line of IT operations. At times these professionals will need access to the cage and overall they will need secured, remote access to the systems.
- **CIOs and the C-suite.** In some cases, data centers are becoming the "factories" of .com businesses and service organizations. Expect that the C-suite, not just the CIO, will want to conduct site visits and have planning meetings on-site at the data center.

Design, Build and Operate

Some data center colocation companies are operators only. They outsource the design and build phases of the facility. Look for colo providers that design, build, and operate their own facilities. This combination of skills and resources can reduce overall costs by avoiding 3rd party markups and improve availability by knowing how each part of the facility works.

Weather Patterns

According to the National Oceanic and Atmospheric Administration, the total damage due to severe weather in 2011 climbed to \$1 billion over the course of 12-14 major events.

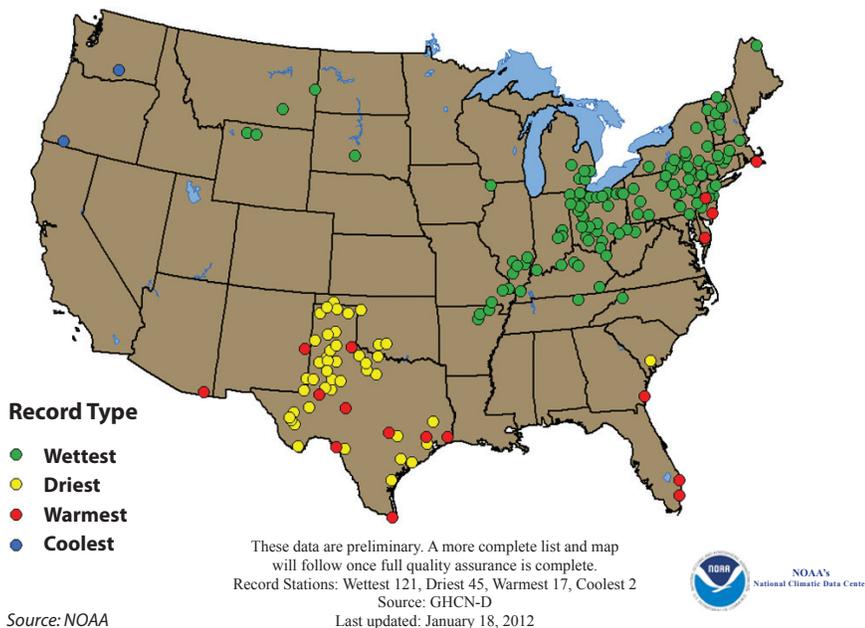
Weather Patterns



Source: NOAA

Many times a data center will have everything that a potential buyer may require — but there is a failure to analyze the weather patterns of a specific region. To detail the importance of researching a location with the appropriate climate, consider the following images.

Selected Annual Climate Records for 2011



Source: NOAA

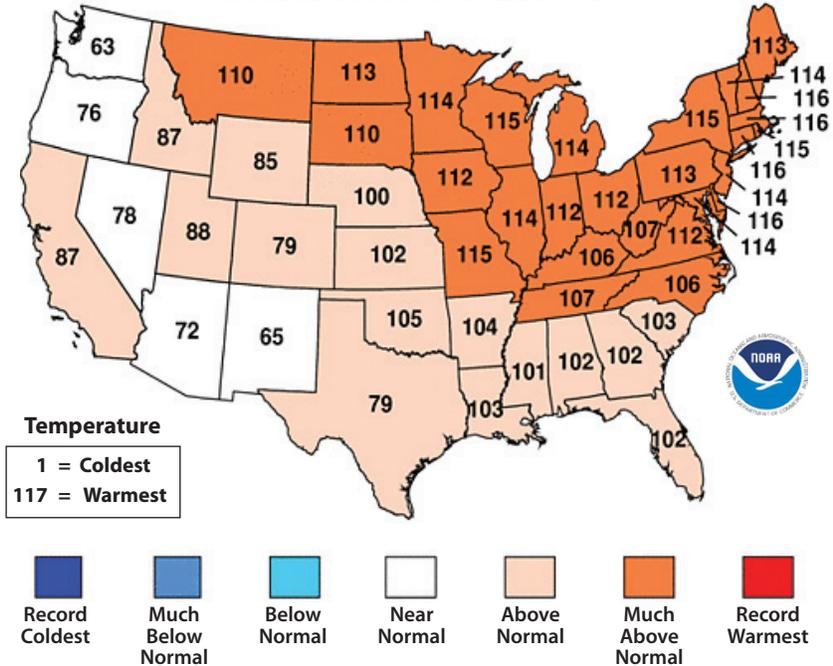
In selecting the right data center, an important decision variable will be the weather of a given region. Even while staying in-state, just a few miles can make all the difference. The decision around weather is always difficult because weather can be unpredictable — but it can still be measured. Plan around the following:

- **Rain fall and precipitation amounts**
- **Average highs/lows**
- **Wind gusts and averages**
- **Flooding**
- **Major temperature variances**

Examining averages and conducting regional weather research will help in the decision making process. Never focus on just one element; it's important to evaluate weather patterns as a whole since various elements can affect a data center differently. For example, in making a decision, a region with slightly warmer than desired temperatures may be desired over one that is more prone to heavy winds or flooding.

Dec 2011 – Feb 2012 Statewide Ranks

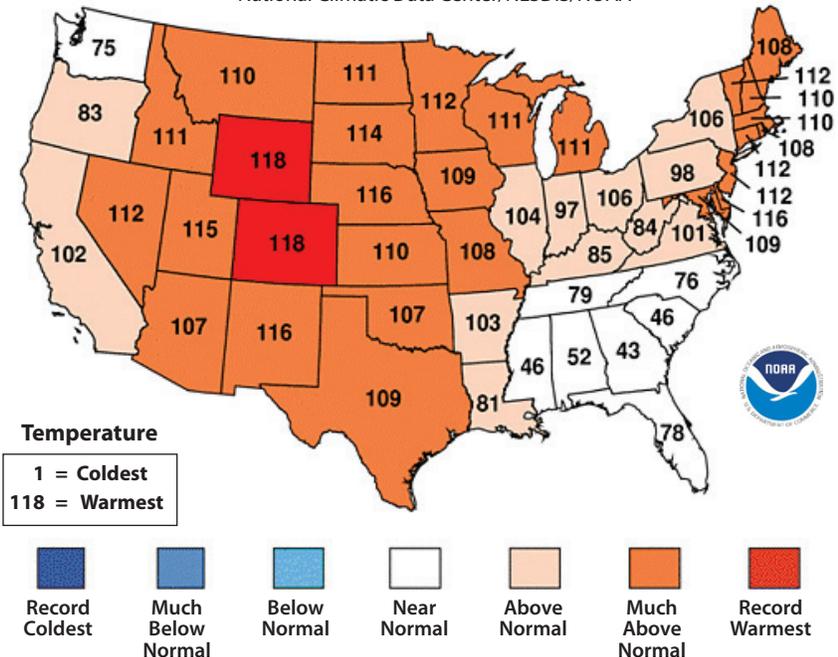
National Climatic Data Center/NESDIS/NOAA



Source: NOAA

June – August 2012 Statewide Ranks

National Climatic Data Center/NESDIS/NOAA



Source: NOAA

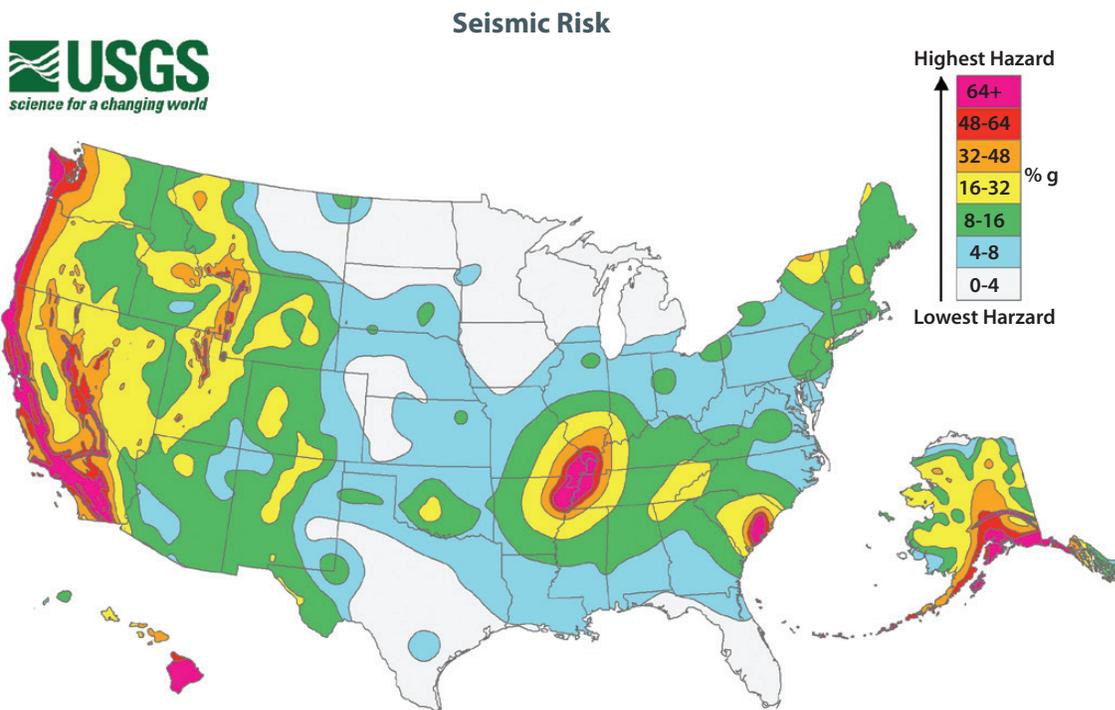
Seismic Risk

In conjunction with weather patterns, make sure to evaluate the seismic history of a given region to ensure the lowest seismic activity risk. Many times this variable is overlooked because of price, data center availability, or in hasty colocation decision-making. There are regions in the USA which are much more prone to seismic activity than others. Some data centers place more of an emphasis on equipment seismic bracing and/or improving their building's seismic stability. Remember, even a small seismic event can have serious repercussions on a live data center.

In the United States, there are a few major seismic hotspots that should be considered. Seismic activity measured within these hotspots can actually occur frequently.

As a whole, California experiences a number of small quakes with magnitudes ranging from 2-5 on the Richter scale. In Northern California, for example, recent data published in Data Center Knowledge says that there is a 63% chance of a magnitude 6.7 earthquake occurring in the next 30 years, a 67% chance in Southern California, and a 99.7% chance in California as a whole. However, not all of California is at risk to earthquakes. For example, some data center providers are locating in the Sacramento area, where seismic risk is quite low yet the location is within driving distance of San Francisco and Silicon Valley.

Seismic risk exists broadly in the United States as well. According to the Federal Emergency Management Agency, 38 US states have regions of moderate to high seismic hazard and 60% of the U.S. population lives in an area of moderate to high seismic risk.



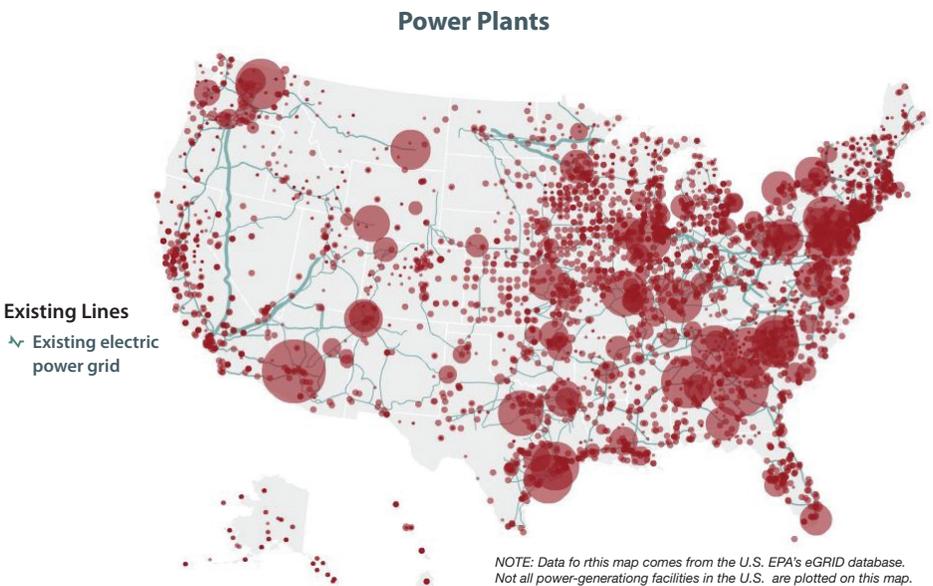
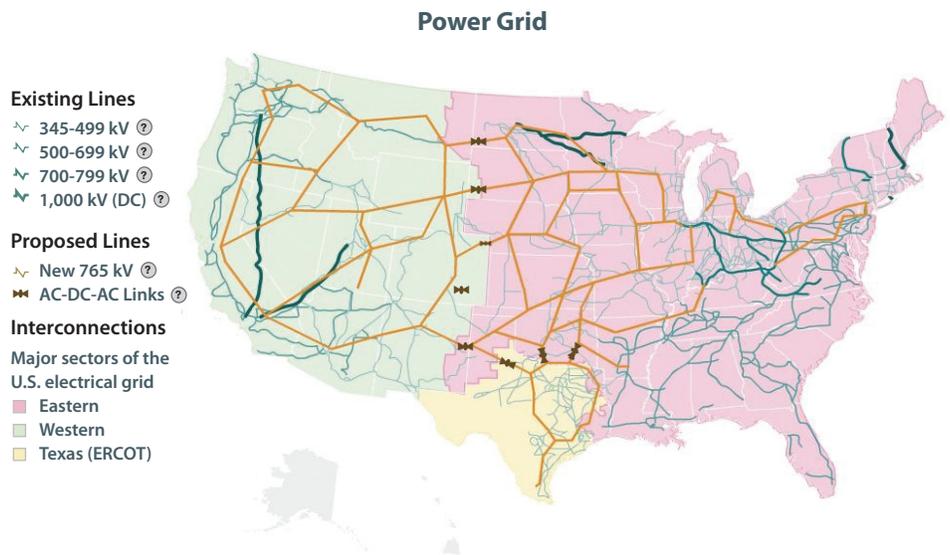
Power Grid and Redundancy

Another variable in the colocation decision process is the regional power grid infrastructure. Just because a geographic area has access to a good amount of power doesn't mean that the power is delivered to the data center. Look for the location of power stations, substations, and feeds to the facility as well as redundancy throughout the delivery system. Also, research recent area outages to understand the time-to-repair for the utility provider. When working with a colocation provider, it's important for them to understand the security and redundancy metrics of their local power grid. Many times, this can span at a national level. Take the time to understand the local power utilities, their capabilities, and how that ties into the colocation provider you are selecting. Remember, in emergency situations, you need to plan for redundancy and the availability of power.

Weather patterns also play a big role in how power is distributed. Areas which require more cooling due to heat constraints may have a stressed power infrastructure. Make sure that the data center you choose has multiple power sources in a location capable of handling those power needs. There are numerous locations where power is sufficient and will meet the demands of the average data center consumer.

Regional power considerations must be taken into account when selecting the right colocation data center.

Check where the power is coming from and where there are available redundancies. Ensure that there are no major power constraints in the area to allow for maximum data center operation.



Source: American Electric Power, American Wind Energy Association, Center for American Progress, Department of Energy, Edison Electric Institute, Energy Information Administration, Electric Power Research Institute, Federal Energy Regulatory Commission, National Renewable Energy Laboratory, U.S. Environmental Protection Agency, Western Resource Advocates. Credit: Producer: Andrew Prince; Designer: Alyson Hurt; Editors: Avie Schneider and Vikki Valentine; Supervising Editors: Anne Gudenkauf and Quinn O'Toole; Additional Research: Jenny Gold; Database and GIS Analysis: Robert Benincasa

Source: NPR

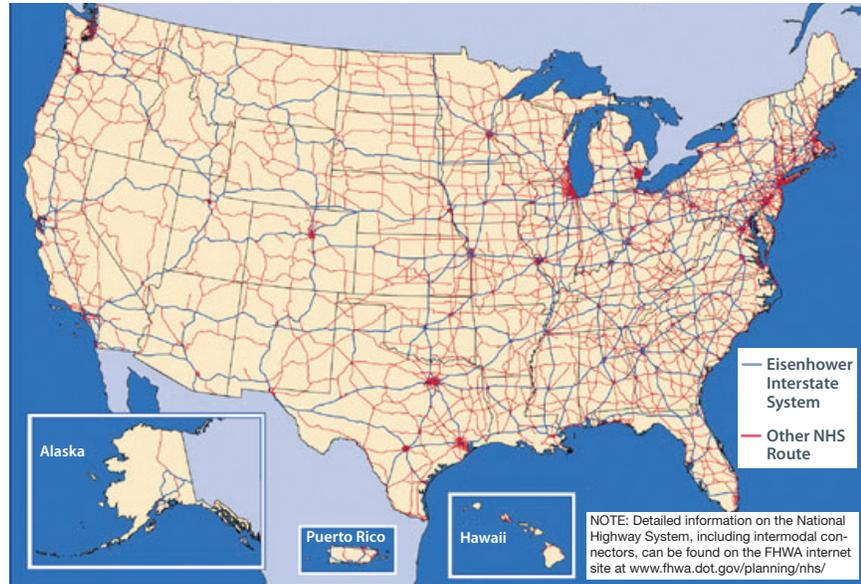
Accessibility – Routes, roads and airports

According to the National Highway System (NHS), highways in the United States represent only about 4% of the nation’s total public road miles and 6.6% of its lane miles, but carry about 50% of the travel.

In selecting a data center colocation provider, take the time to analyze a map of the region to ensure that there are easy ways in and out of the area. Look for routes in and out of the facility that do not require major roads or highways. Ensure that there are several ways to access the data center and that there are redundancies built in for easy access.

Another important factor is the availability of an airport. In some situations, equipment and personnel will have to be flown in for support. Ensuring that your data center has a readily accessible airport facility may be an important requirement. Furthermore, take the time to learn where international and regional airports reside. Since there isn’t always the need to be around a major airport, regional locations may serve the function as well.

The National Highway System



Source: Federal Highway Administration (FHWA)

USA Major Airports



Source: Maps of World (Updated August 2012)

SECTION 2

FACILITY – Critical infrastructure that grows with you

Now that there is a better understanding of the actual location of the data center, we have to look at the facility itself. There are several decision points which must be made when examining the right colocation datacenter for your organization. In selecting the right facility, the decision-making process should revolve around the following:

High-Density Metered Power

In data center colocation, power density and availability have emerged as critical requirements. Power densities found in data centers five years ago can no longer meet the needs of high capacity servers and storage devices. In addition, the availability architecture has changed dramatically. Traditionally, many data centers focused on the number of backup or spare devices in the power delivery architecture. You would see N designs where N is the number of devices needed. Today there are 2N designs which provide two active devices, N+1 with one spare and N+2 with two spares. The state of the art combines 2N feeds and N+2 infrastructure into a design known as 2N+2 which is described below. When working with a data center colocation, it's important to understand that today's demands have outgrown many existing data centers. Metered power allows you to pay for what you use today and grow your power draw incrementally over time. When making a data center decision, consider the following decision points:

- **Look for providers who are capable of scaling beyond 8 kW per rack or 200 watts per square foot.** Trends show that power consumption is on the rise and more will be required from today's modern colocation providers.
- **Access to a dedicated substation.** Some of the best data centers currently offer fully dedicated power substations as part of their delivery offering.

Power densities found in data centers five years ago can no longer meet the needs of high capacity servers and storage devices. In addition, the availability architecture has changed dramatically.

- **Look for 2N+2 electrical infrastructure redundancies.** As mentioned earlier, N+1 is no longer enough for a redundant environment. Truly progressive data center environments are now offering 2N+2 redundancy: two completely independent power paths from the utility to the server rack and an additional N+2 design for critical infrastructure such as generators and UPSs.
- **Usage based pricing with visibility.** Typical retail colocation providers charge for power by the circuit ("whip") or by the kilowatt ("all-in"). From a customer point of view, the challenge with these two pricing models is twofold — customers generally don't have a benchmark for how much power they actually need and customers don't usually use 100% of the power on day one — they ramp into their power requirements over time by migrating or growing their colocation footprint in phases. Leading colocation providers offer metered power where you pay for what you use and robust branch circuit monitoring so the customer has insight into their capacity requirements over time.

Cooling

The goal of any efficient data center colocation provider is to meet the cooling needs of the computing equipment and facility and drive down cooling costs and the power usage effectiveness rating or PUE. PUE shows how much overhead is associated with delivery of power to the rack. A measurement of 1 says there is no overhead. 1.2 would represent 20% overhead. In the most efficient scenario, customers would pay for the power they use (metered) multiplied by a PUE factor to account for additional power needed to cool the facility and keep the lights and other devices running. Look for a colocation provider who is thinking ahead and utilizing natural resources, such as free outside air, and managing to a low PUE.

- **Hybrid cooling technologies.** Look for providers using advanced, natural forms of cooling. Not only will this drive down costs, but it will also lessen the impact on the environment. Also look for a data center that integrates and optimizes these “free cooling” technologies into a hybrid cooling plant, capable of “mixing” cooling methods to be most efficient. Many such progressive colocation data centers are EPA ENERGY STAR certified.
- **Look for 100% availability.** Many data centers boast 99% uptime. In today’s high-demand, always-on information economy, that one percent can cost millions of dollars. One percent downtime means over 7 hours of outages per month and over 3.5 days of outages per year.
- **Expect thorough cooling redundancy.** Like the power requirements, with increased density, data centers are susceptible to overheating if an inadequate cooling system is in place. State-of-the-art colocation data centers offer an N+2 chiller plant and N+2 or greater redundancy on air handling units (CRAC or CRAH).

Physical Security

For organizations looking for truly secure facilities, insist on in-house security teams. In analyzing a good security model, consider the following:

- **In-house security staff.** Having an in-house security team (not outsourced) ensures that those employees have the data center’s security needs in mind. Armed guards and a full security staff should be a consideration in the decision process.
- **Multi-factor identification and authorization.** Ensuring the safety of millions of dollars worth of equipment will require ID checks, biometrics, and other forms of identification measures.
- **Layered security zones.** Layered security zones ensure that there is redundancy in the security policy as well. Entry points, floors, and access to customer cages all represent layers of security. Some data centers have gone so far as to build a building within a building for maximum security.
- **Camera and security systems monitor the 360-degree data center picture.** Truly secure environments will fully prohibit any public access. Furthermore, environments which are not 24x7x365 secure should be pushed down on the consideration list. Look for advanced security measures including state-of-the-art camera systems, bollards, fencing, and security all the way from the roof to the parking lots.
- **Advanced security certifications.** Some colocation providers are taking the next step in securing their infrastructure by obtaining advanced certification and audit metrics. For a security-minded organization, look for these certifications:
 - **PCI DSS 2.0 provider**
 - **SSAE 16 audited**
 - **ISO 27002**

Data Center Infrastructure Management (DCIM)

Don't be left in the dark, DCIM is much more than the latest buzz acronym to hit the data center industry. Data centers have always been purpose built facilities with lots of complex technology. Managing those technologies has been problematic. At best, devices would have management software but individual software systems could not work together. At worst, blinking lights needed to be monitored in person. The result was a highly tuned system that when it breaks, breaks ugly. DCIM changes that. Look for a data center colocation provider with DCIM built into their infrastructure.

- **Everything is connected.** Critical infrastructure must work as a system. Devices such as power and cooling delivery need to be connected to a common network to allow for seamless monitoring.
- **Sense and respond software.** Management software needs to do more than tell you what just happened. It needs to provide trend analysis and intelligence that can identify problems before they happen and solve them.
- **Innovation through integration.** Look for a data center provider that has in-house software development resources that can integrate disparate software management tools into a robust system.
- **Secured portal and reporting.** You need to have an online view of your critical infrastructure and the ability to produce reports and trending.

Amenities

Many times, IT professionals may spend days at a data center colocation working on racks, servers, applications or tuning the infrastructure. Good accommodations and a comfortable environment are very important decision points when considering a colocation data center.

- **Look for comfortable working areas.** No one enjoys working within a bland, enclosed space. A data center with a “campus” mentality will have good areas for engineers to work.
- **Office space.** Data centers are becoming a key part of the business operation. Look for data centers that offer onsite office space that you can rent for your staff.
- **Relaxation rooms.** IT can be a very stressful place. Large organizations should look for colocations which provide an area for downtime for their engineers.
- **Available workstations and conference rooms.** This should be a standard in the decision making process. Good workstations and conference areas increase productivity and employee morale.
- **Easily accessible network and Wi-Fi availability.** When working from a data center, engineers and customers must have access to a healthy network and wireless environment. Good providers will have capacity built in to ensure availability of a good network.
- **Fresh coffee and drinks.** As any IT person will tell you, coffee can oftentimes be an engineer's best friend. Furthermore, having a comfortable environment offering such amenities will keep engineers happier and in good spirits — this can be especially helpful when working on serious issues.

Green Technology

According to a recent Pike Research report, power and cooling infrastructure solutions will be the largest portion of the green data center market opportunity, representing 46% of revenue over the next five years. This means that data centers are focusing on green technologies — more specifically, efficiencies built around those green technologies.

- **Look for the EPA ENERGY STAR certification.** Very few data centers currently hold this certification – even fewer in the multi-tenant colocation data center space. This certification means that a colocation provider is actively seeking ways to improve their environment and continue to be efficient. Furthermore, these certified data centers are working with advanced metering technologies to monitor power being utilized by both the customer’s IT equipment as well as facility cooling and electrical equipment.
- **Opportunity for airside economization.** In most regions of the United States, air from the outside (“free air”) can be used during some periods of the day or year when the air is cold enough to maintain optimal data room temperature and humidity. When the data center is using free outside air, the only energy required is for running the fans to draw the air inside the building. Look for a colocation provider that utilizes this environmentally (and PUE) friendly cooling method.
- **Cooling and power efficiencies.** Given the current focus on the environmental effects of data centers in today’s “green” culture, many data centers providers are taking a closer look at ways to improve their cooling and power efficiency. Data centers are replacing constant speed pumps and fans in their cooling plants with variable frequency motors that can more accurately match cooling demand to supply. They’re also investing in smart, automated ways to configure and operate their cooling plants in response to data floor and outside temperatures and humidity.

The EPA ENERGY STAR Certification means that a colocation provider is actively seeking ways to improve their environment and continue to be efficient.

- **Look for water conservation efforts.** Using reclaimed water in data center cooling methods is catching on amongst many top colocation providers. There are two drivers for using reclaimed water: first, it is usually cheaper per gallon than traditional potable water (yet it is treated to almost the same quality standard), and second — it reduces the impact that data centers have on the environment. Favor colocation providers that integrate (or have plans to) reclaimed water use into their cooling plants.
- **Utility power.** You should be aware of the sources of energy used by the power utility. Are they investing in renewable sources of energy, such as hydro, wind, and solar? Do they have targets in place regarding how much of their energy should come from renewable sources? Is the utility actively working to reduce its carbon footprint? Every power utility is challenged in this area, but a number are making significant progress.

In selecting the right facility and location, working with the above considerations will help with data center efficiency for the current organizational demand — and the future requirements as well.

SECTION 3

ANALYZING THE WORKLOAD – What is being delivered and where?

In selecting the right type of data center colocation, administrators must thoroughly plan out their deployment and strategies. This means involving more than just facilities teams in the planning stages. The process to select a good data center has to involve not only the physical elements of the data center, but the workload to be delivered as well.

The current business landscape has created what is known as a “data-on-demand” generation where information is needed immediately on any device, anywhere and at any time.

Are you working with web applications? Or, are you delivering virtual desktops to users across the nation? There are several key considerations around the type of data or applications an organization is trying to deliver via the data center.

In making these preparations, engineers and facilities personnel should consider the following important decision points.

Cloud and Big Data

Cloud computing will continue to evolve and change how organizations distribute their data. The idea behind cloud computing is that information becomes distributed between multiple locations and can then be delivered more creatively. IT consumerization and BYOD (Bring Your Own Device) have driven the need for cloud computing as today’s IT environment continues to evolve. The current business landscape has created what is known as a “data-on-demand” generation where information is needed immediately on any device, anywhere and at any time. Because of this culture, it’s important to evaluate the following two components when selecting a colocation provider:

- **Data center distribution.** A national or global organization will have the need for data distribution. Many times application or database requirements indicate the need for the data to reside close to the actual workload. When looking for a data center partner, choose those that are close to the major telecommunications hubs and your critical user communities where centralized data is a requirement. This model will ensure lower latency and increase application/data performance.
- **Big data.** We are now dealing with terabytes — and even petabytes — worth of information in today’s information economy. This increase in data usage is the result of more information being processed within organizational databases as the need for distributed information rises. Organizations with heavy data demands should look for partners who are not only close to their data points, but also capable handling the bandwidth of that data. From a data center design and planning perspective, this means having the right type of network switches in place, and maintaining good connectivity within the data center itself.

Mission Critical vs. Good Enough

Take the time to understand the workloads, applications, and data instances which will be running within your data center. This means classifying the workloads based on their priority. In those situations, it's important to conduct a Business Impact Analysis (BIA) to truly understand the value of various data and infrastructure components. The results of a BIA are very descriptive and indicate which systems should run in high-availability Tier III data centers. Furthermore, it will help determine which systems must have higher uptime requirements versus those with a lower priority.

- **Mission critical.** In conducting a BIA, various components, data elements and hardware equipment may be classified as "mission critical." This means that specific infrastructure must remain up 100%, which requires a highly redundant supporting infrastructure. Recovery times for these components can range from just a few minutes to a few hours. The important thing to remember is that these components are the vital pieces which will keep your business running. During a BIA, recovery times are established based on the impact of those systems going down. By establishing recovery times and planning around them, managers can select a provider which is capable of meeting those demands with their internal processes, infrastructure, and support. Additionally, a good colocation provider will hold concurrent engineering sessions with your IT staff to fully understand the priorities in your IT deployment. They should be able and willing to strategize power circuit deployment to maximize redundancy through the use of multiple UPS systems (greater than the standard A/B power supply mentality).
- **Good enough.** Similar to mission critical components, the 'good enough' model contains infrastructure pieces which may have longer recovery time objectives. This doesn't necessarily mean they are any less important, rather, it means that these components can take longer to recover. These environment pieces are also established during a BIA and their recovery times are identified as well.

Network Bandwidth and Latency

With the increase of traffic moving through the Internet, there is a greater demand for more bandwidth and less latency. As discussed earlier, it's important to have your data reside closer to your users as well as the applications or workloads which are being accessed. Where data may have not fluctuated too much in the past, current demands are much different.

- **Bandwidth bursts.** Many providers now offer something known as bandwidth bursts. This allows the administrator to temporarily increase the amount of bandwidth available to the environment based on immediate demand. This is useful for seasonal or highly cyclical industries. There will come a time when for a period of business operation, more bandwidth is required to help deliver the data. In those cases, look for partners who can dynamically increase that amount and then de-provision those resources when they are no longer being used.
- **Network testing.** Always test your network and the network of the provider. Examine their internal speeds and see how your data will act on that network. This also means taking a look at the various ISP and connectivity providers being offered by the colocation provider. Many times a poor networking infrastructure won't be able to handle a large organization's 'Big Data' needs despite potentially having a fast Internet connection. Without good QoS (Quality of Service) and ISP segmentation, some data centers can actually become saturated. Look for partners with good, established connections providing guaranteed speeds.
- **Know your applications.** One of the best ways to gauge data requirements is to know and understand the underlying application or workload. Deployment best practices dictate that there must be a clear understanding of how an application functions, the resources it requires and how well it operates on a given platform. By designing the needs around the application, there is less chance that improper resources are assigned to that workload.

Balancing the Workload, Continuity and Disaster Recovery

Selecting a colocation provider goes well beyond just choosing their internal features and offerings. Companies looking to move to a provider platform must know what they are deploying, the continuity metrics of their infrastructure and incorporate disaster recovery into their planning.

- **Workload balancing.** When working with a data center provider, design your infrastructure around a well-balanced workload model. This means that no one server is over-provisioned and that each physical host is capable of handling the workload of another host should an event occur. Good workload balancing will ensure that no one system is ever over-burdened. This is where a good colocation partner can help. Many times monitoring tools can be used to see inside the workload to ensure that the physical server running that application is operating optimally. Sometimes features offering dynamic workload balancing are available. If that is a requirement, make sure to have that conversation with your colocation partner.
- **Business continuity.** In a business continuity model, the idea is to keep operations running optimally without disruptions in the general infrastructure. One of the best ways to understand business continuity metrics is to, again, conduct a BIA. By having documentation available showing which workload or server is most critical, measures can be taken to ensure maximum uptime.
- **Disaster recovery.** A core function of many colocation providers is their ability to act as a major disaster recovery component. In working with a partner, select a design which is capable of handling a major failure, while still recovering systems quickly. There is really no way of telling which components are more critical than others without conducting a BIA. Without this type of assessment, an organization can miss some vital

When working with a data center provider, design your infrastructure around a well-balanced workload model. This means that no one server is over-provisioned and that each physical host is capable of handling the workload of another host should an event occur.

pieces and severely lessen the effectiveness of a DR plan. Once the DR components are established, an organization can work with a colocation provider to develop a plan to ensure maximum uptime for those pieces. This is where clear communication and good DR documentation can really help. The idea here is to understand that a major event occurred and recover from that event as quickly and efficiently as possible. A good DR plan will have a price associated with it — but from a business outage perspective, it's worth it.

Various technologies can affect how well a data center performs. The distance the data has to travel and the amount of bandwidth provided by a colocation provider can mean the difference between a great user experience and a failed colocation deployment. Cloud computing has created a greater dependency on WAN technologies and virtualization has enabled significantly more powerful servers and more dense storage. With these new technologies come new considerations around how this type of data is being stored and delivered. When selecting the right colocation provider, make sure that their infrastructure is capable of growing with the needs of your organization.

SECTION 4

ENVIRONMENT CONTROL AND MANAGEMENT – Good SLAs make good neighbors

In selecting the right data center colocation, important considerations are the visibility and management tools available for the customer. Monitoring power consumption, having a DR-ready environment and truly partnering with the provider are vital consideration points in making the proper data center choice.

During the planning phases, things like contracts, expectations and management tools must be laid out to ensure that everyone is on the same page. When working with a colocation provider, there will be important planning points and ongoing considerations around a good data center rollout.

Working with a Service Level Agreement

When selecting the right colocation provider, creating or having a good SLA and establishing clear lines of demarcation are crucial. Many times, an SLA can be developed based on the needs of the organization and what is being hosted within the data center infrastructure. This means identifying key workloads, applications, servers and more. From there, an organization can develop base service agreements for uptime, issue resolution, response time and more. Creating a good SLA document can take some time — but it's important to do so carefully since this can govern the performance of your environment. Some very high-uptime environments will build in credits into their SLA. In these situations, for example, a colocation provider could issue credits if power is unavailable. Creating an SLA is a partnership between the data center provider and the customer. Expectations must be clearly laid out to ensure that all performance, recovery and other expectations are met. Surprises or encountering unknowns in a production, highly utilized environment can result in loss of productivity, time and dollars.

When you buy data center colocation you are buying a slice of critical infrastructure and ongoing maintenance. Without a robust maintenance program, technology will fail.

Maintenance and Testing

Don't forget, when you buy data center colocation you are buying a slice of critical infrastructure and ongoing maintenance. Without a robust maintenance program, technology will fail. Look for documented MOPs (method of procedure) and SOPs (standard operating procedure) that are used consistently and improved over time. Make sure your SLA does not exclude maintenance windows or emergency maintenance. Your colo provider should be able to show you their monthly, quarterly, and annual maintenance schedules for all critical elements of the mechanical and electrical systems including chillers, air handlers, generators, batteries, and UPSs. You should be able to observe and even participate in maintenance exercises. How are you notified about maintenance windows and procedures? Finally, ask the ultimate question, "Do you plan for and test a full utility outage?" Systems need to be designed with sufficient redundancy to allow for proper maintenance. Colocation providers are reluctant to maintain systems if it could potentially cause an outage. The industry best practice is to be able to "fix one and break one, concurrent with a utility outage."

Having a DR-ready Contract

For some organizations, moving to a colocation data center is the result of a disaster recovery plan. In these situations, it may very well be possible to integrate a DR contract into an SLA or as a stand-alone agreement. In this document, the organization and colocation provider establish which internal systems must be kept live and create a strategy to keep those systems running. When designing a contract around a disaster recovery initiative, consider some of the following points:

- **Use your BIA.** As mentioned earlier, a business impact analysis will outline the key components within a corporate environment which must be kept live or recovered quickly. This BIA can be shared with your colocation provider to ensure that they are able to meet those requirements. By having such a document, an organization can eliminate numerous variables in selecting a partner which is capable of meeting the company's DR needs.
- **Communicate clearly.** Good communication between the colocation partner and the organization is vital in any DR plan. A situation in which an unknown system or component (that was deemed as critical but not communicated) goes down will become a serious problem. Aside from bringing the piece back up – now there is the question of responsibility. By knowing who is responsible for which piece during a disaster event will greatly streamline the recovery process.
- **Understand and communicate the DR components.** Prior to moving to a colocation, establish your DR requirements, recovery time and future objectives. Much of this can be accomplished with a BIA, but planning for the future will involve conversations with the IT team and other business stakeholders. Once those needs are established, it's important to communicate them to the colocation provider. The idea here is to align thought processes to ensure a streamlined DR environment.

- **Onsite and offsite supplies.** If a disaster occurs, you need both onsite and offsite sources of key supplies. Are there onsite supplies of diesel fuel for generators and water for cooling systems? Are there established services in place for delivery of water and diesel fuel should onsite supplies be depleted? Does the colo provider conduct disaster recovery scenarios with key suppliers? For example, what if there is a power outage that causes the offsite diesel fuel supplier to be unable to fill their trucks — what happens?

Using Management Tools

One of the most important management aspects within any environment is the ability to have clear visibility. This means using not only native tools, but ones provided by the data center partner. Working with management and monitoring tools for the workload is very important. Also important is to have a good view into the physical infrastructure of the data center environment. Data and reporting from these monitoring tools should be made available through a secured portal.

- **Power monitoring.** Always monitor the power consumption rates of your environment. The idea here is not only to know how much power is being used, but to make the environment more efficient. Look for ways to save on power based on requirements. For example, certain power-heavy racks may need to be distributed more efficiently, thus saving on power consumption and costs.
- **Cooling monitoring.** Much like power, keeping an eye on cooling is important as well. This can be outlined as part of an SLA or an organization can manually monitor cooling as well. Ensuring that IT systems are running optimally will revolve around how well the environmentals are controlled.

- **Rack conditions and environmental.** Keeping track of the environment variables will help create a more efficient rack design. Some servers will generate more heat while others may need more power. By seeing what system is taking up which resources, administrators can better position their environment for optimal use. Work with a partner who can give you tools looking into the thermal and environmental performance of your infrastructure. Using those tools, an organization can make better decisions on how to build their rack infrastructure.
- **Uptime and status reports.** Regularly check individual system uptime reports and keep an eye on the status of various systems. Having an aggregate report will help administrators better understand how their environment is performing. Furthermore, managers can make efficiency modifications based on the status reports provided by a data center's reporting system.

Many times logs will show an issue arising before it becomes a major problem. Administrators can act on those log alerts and resolve problems at a much steadier pace rather than reacting to an emergency.

- **Logs.** A log monitoring platform is always very important. One recommendation is to have a log aggregation tool which collects various server, system and security logs for analysis. This not only creates an effective paper trail, it also helps with infrastructure efficiencies. A good log monitoring system is one of the first steps in designing a proactive, rather than a reactive environment. Many times logs will show an issue arising before it becomes a major problem. Administrators can act on those log alerts and resolve problems at a much steadier pace rather than reacting to an emergency.

Issue Resolution and Communication

A large part of having an effective environment will be the issue resolution practices and partner-to-customer communication. Although a lot of this can be outlined in the SLA; specific issue resolution matters need to be discussed. When designing an issue resolution conversation, it's important to identify core data center components and then communicate that to the colocation provider. For example, a hard drive within a particular system may take priority over another issue should a simultaneous event occur. In this scenario, the SLA and the BIA are used to create a clear plan for resolving issues quickly and in the right order. There will be instances when a specific problem takes precedence over others due to the nature of the occurrence. Without good communication, the colocation provider may not know which problem to fix first and assign a resource to a lesser important ticket. Share your BIA findings and clearly communicate which data center components need resolution first.

The process of selecting the right colocation provider should include planning around contract creation and ensuring that the right management tools are in place. A colocation data center is an important extension of any organizations and therefore must be properly managed. Good data center providers will oftentimes offer tools for direct visibility into an infrastructure. This means engineers will have a clear understanding of how their racks are being cooled, powered and monitored. These types of decisions make an infrastructure more efficient and much easier to manage in the long-term.

SECTION 5 COLOCATION DECISION CHECKLIST AND BEST PRACTICES

The colocation selection process never really revolves around just one metric. Instead, an organization has to look at various variables which can comprise a solution which will directly fit in with the needs of the business. One of the best ways to come to a decision is to outline what is important, create a check list, and follow that list during the selection process.

Remember, many times when the choice is made, a contract is signed for a pre-determined term. This means that once the gear is in the data center — it'll be time-consuming, complex, and expensive to move it to another site should an issue occur.

To get a solid understanding of which data center provider is the best option, working with the checklist on the right can really help. To use it, simply fill out the boxes as the example shows, multiply the Availability/Meets Requirements score times the Importance Rating score, and add up the results. The data center with the highest checklist score is probably the colocation provider that best fits your organization's direct needs both now and in the future.

The checklist to the right was designed from many of the already discussed variables in this guide. By understanding power demands, facility variables and all of the components which were discussed earlier, managers are able to make better colocation decisions.

The Colocation Decision Checklist

Feature/Requirement	Availability/Meets Requirements Yes = 1 No = 0	Importance Rating (1-10) – Low to High
<i>Example Feature Yes (1)</i>	<i>Yes (1)</i>	<i>8</i>
Physical Location (Regions)		
Redundant Power (At least N+2)		
Power Density		
Metered Power		
Efficient Cooling		
Green Technology Adoption/ Certifications		
Customer Service		
Access (Roads, Airports, Hotels)		
Office Space		
Loading Docks and Storage		
Engineering Expertise		
Response time ('x' minutes/hours)		
Contracts and SLAs		
Maintenance and Testing		
Customer Satisfaction		
Management and Monitoring Tools (DCIM)		
Bandwidth Availability		
Disaster Recovery Options		
Physical Security		
Expansion Capabilities		
Location Amenities (Lounge, workspaces, conference rooms, Wi-Fi)		
TOTALS		

There are some general best practices to follow when selecting the right provider. The choice to go with a data center must truly be unique since your organization's needs will certainly differ from others — even if you are in the same industry with like demands. Consider the following:

Avoid “Cookie-Cutter” Data Centers

There are many providers currently available on the market. Many of them offer similar features, benefits and services. An organization looking to truly partner with their provider should look for a more customized and personalized approach. This ensures that the needs of the individual organization are met and maintained. Work with the provider's account management team to ensure that your direct needs are met. This will not only help create a custom look and feel for your data center experience, but will also help with future growth and expansion.

Customer Experience is King

Traditionally data center colocation providers have focused primarily on infrastructure and cost. While those focus areas continue to be important, a new area is emerging known as customer experience. Look for providers that recognize the importance of customer experience to the services they deliver and seek to continually improve service levels. The best practice is a provider that measures customer satisfaction and loyalty through proven methodologies such as the Net Promoter Score and will share survey results and cure plans with you.

Know What it Costs Should Your Environment Go Down

One of the key benefits of moving towards a colocation solution is the potential for greater uptime and resiliency. Many times, one of the first things that upper management looks at is the cost of moving to a colocation data center. Although there is a price associated with such a move, take the time to understand what it costs for your business to be down for any period of time. By seeing a quantifiable number associated with a systems outage, it'll be significantly easier to justify the costs of moving to a fully redundant, high-uptime environment. Conducting a thorough BIA is probably one of the best ways to place a dollar figure on an environment in a down state.

Educate the IT Team

Moving to a colocation environment is a serious undertaking. One of the most important recommendations is to educate the IT team and keep them informed. Prior to even making the move, make sure that the internal IT team is familiar and comfortable with the various aspects of going to a colocation facility — and the colocation provider should be involved in that process. Management tools, new technologies and recovery procedures should all be discussed and documented. Furthermore, it's important to involve all of the IT stakeholders in the process. This means educating the security team on new data center policies, informing the networking team of new protocols and routing methods, and allowing the application and virtualization engineers an opportunity to learn the hardware that they will now be working with. The better informed the IT team, the quicker they can respond to upgrades, questions and issue resolutions.

Understand the Uptime Institute Data Center Tier Model

Developed by the Uptime Institute, the Tier classification examines site-specific infrastructure components required to maintain a certain level of data center operation. The idea here is to combine several subsystems into a cohesive unit to bring the best data center experience possible to the end-user. Individual metrics, such as power, UPS, cooling, are all measured as part of the integrated approach to delivering a resilient and efficient data center. In this model, there are 4 Tiers to understand:

- **Tier I:** Basic site infrastructure
- **Tier II:** Redundant site infrastructure capacity components
- **Tier III:** Concurrently maintainable site infrastructure
- **Tier IV:** Fault tolerant site infrastructure

One of the best ways to follow this Tier requirements model is to apply it to the respective data center and see how each subsystem fits into the tier model that's right for your organization.

Some data centers will have aspects of multiple tiers as part of their design. To understand this better, the following table from Uptime examines these requirements.

Remember, selecting a data center colocation provider is a longer term commitment. Most data center colocation contracts are 24-36 months long and it can be challenging to move or decommission gear. There are times when the planning process falls short and companies realize that their equipment is now residing at a provider who is unable to meet some core requirements. The secondary move to find the right provider can be extremely costly and can have seriously negative effects on both the infrastructure and the business model. Take the time to fill out the above checklist and discuss the results with the evaluation team to truly ensure that the vital IT gear being placed inside of a data center will be secure, safe and operational to the demands of the business.

Tier Requirements Summary

	Tier I	Tier II	Tier III	Tier IV
Active Capacity Components to Support the IT Load	N	N+1	N+1	N After any failure
Distribution Paths	1	1	1 Active & 1 Alternative	2 Simultaneously Active
Concurrently Maintainable	No	No	Yes	Yes
Fault Tolerance	No	No	No	Yes
Compartmentalization	No	No	No	Yes
Continuous Cooling	No	No	No	Yes

Source: Uptime Institute – Data Center Site Infrastructure Tier Standard: Topology

CONCLUSION

As more organizations move their environments into a colocation data center environment, the need for diverse and feature-rich offerings will continue to rise. High-density metered power, stable and accessible locations, carrier diversity, robust fiber, 100% availability, customer satisfaction, efficient cooling and 7x24 security are all variables which managers should seek from their providers. The selection process must be rigorous since the IT infrastructure is more important in today's business world than ever before. Companies are highly reliant on their IT systems and organizations should seek high-uptime, truly redundant colocation facilities capable of meeting some serious demands.

Remember to involve extended team members during the decision process. Managers from various departments may have differing requirements and it's important to align the IT infrastructure vision with already existing features available at a given data center provider. The goal of any organization isn't to just build an infrastructure for today — rather — the environment must be capable of handling the needs of tomorrow as well. This is why when selecting a colocation data center, the mentality must change from "this is just another data center hosting our equipment" to "this is now our partner who will help our business grow and stay more resilient."

Creating a partnership with the colocation provider will create a cohesive extension to your company's already existing IT team. These folks will take daily charge to ensure that your systems stay healthy, viable and running. Remember, never build for today — always build for tomorrow.

Bill Kleyman is a veteran, enthusiastic, technologist with experience in datacenter design, management, and deployment. His architecture work includes virtualization and cloud deployments as well as business network design and implementation. Bill enjoys writing, blogging, and educating colleagues about technology. Currently, he has published over 100 articles, several podcasts and webcasts and has participated in many other forms of interactive media. He has also published whitepapers on cloud computing and managing IT environments. Bill is a featured writer, contributor and author and his work has been published on Business Week, Tech Target, Data Center Knowledge (INET Interactive) and The Champion Community (UBM DeusM). Bill stays busy as a virtualization and cloud solutions architect at MTM Technologies where he designs solutions for some of the world's largest organizations.



Data Center Knowledge (DCK) is a leading online source of daily news and analysis about the data center industry. We cover the latest developments and trends driving the powerful growth in demand for mission-critical facilities, the challenges and opportunities presented by high-density computing and its impact on power and cooling, and the evolution of the industry to include cloud computing and modular data centers.