NTT America

# Which Data Center?

## Key factors in choosing who to trust with your IT Infrastructure

***Tarif Abboushi***
*Vice-President of Data Center Infrastructure*
*NTT America*

May, 2009

NTT Communications

# Which Data Center?

Key factors in choosing who to trust with your IT Infrastructure

## Introduction

A decade ago, the term "mission critical facilities" related to such things as hospitals, power stations and military installations. It is testimony to the role Information and Communications Technology now plays in our lives that today the term is used almost exclusively for data centers.

Data centers have become the information processing nerve centers and storage cores of corporate, institutional and governmental entities worldwide. The process of selecting a data center can be complex and time consuming because of the number of options available and the downside of getting it wrong. This paper is intended as a high-level primer for selecting the right data center. It outlines the key parameters to be considered by decision makers tasked with evaluation and procurement of data center facilities to meet their organizations' needs.

## Decision Criteria

Decision criteria can be grouped onto three broad categories:

**Reliability:** How reliable the data center is in terms of physical security, power, processes etc. and how these relate to uptime or availability

**Efficiency:** The power usage of the facilities and equipment, and how well "green" technologies and designs are exploited to minimize operational cost and carbon emissions.

**Vendor:** Whether a hosting provider or a facilities management company, it is important to examine such things as contractual terms, financial strength of the vendor and Service Level Agreements.

## Data Center Reliability

Data center reliability is measured in terms of availability (or uptime), which in turn is a function of infrastructure design topology and operational sustainability, including human factors. The Telecommunications Industry Association (TIA) introduced the concept of tiered reliability for data centers in April 2005 with the publication of TIA-942 Telecommunications Infrastructure Standards for Data Centers. Similarly, the Uptime Institute's Tier Classification and Performance Standard (http://uptimeinstitute.org/content/view/411/) defines a four-Tier system that has become a de facto industry benchmark for data center availability. As applicable to design topology, the Tiers can be summarized thus:

| Tier | Description | Representative Site Availability |
|---|---|---|
| I | **Basic Infrastructure**: Non-redundant infrastructure components and distribution paths. Annual site shutdown necessary for some maintenances operations. | 99.67% |
| II | **Redundant Components Infrastructure**: Redundant capacity components, non-redundant distribution paths. Susceptible to downtime from planned and unplanned events. | 99.75% |
| III | **Concurrently Maintainable Infrastructure**: Redundant capacity components and distribution paths. Susceptible to downtime from unplanned events. | 99.98% |
| IV | **Fault-tolerant Infrastructure**: Multiple, independent infrastructure systems each with redundant capacity components and distribution paths. | 99.99% |

To translate the design topology-based tier definitions into practical terms that will better enable IT executives to make reasoned data center selection decisions, the Uptime Institute has published more meaningful site outage information. The chart is based on measurement of actual availability of 16 data centers with infrastructure topologies meeting the four tier definitions, including human factors (which account for up to 70% of total site failures), over a period of up to ten years, with uptime measured from the perspective of the IT client's operations in the white space.

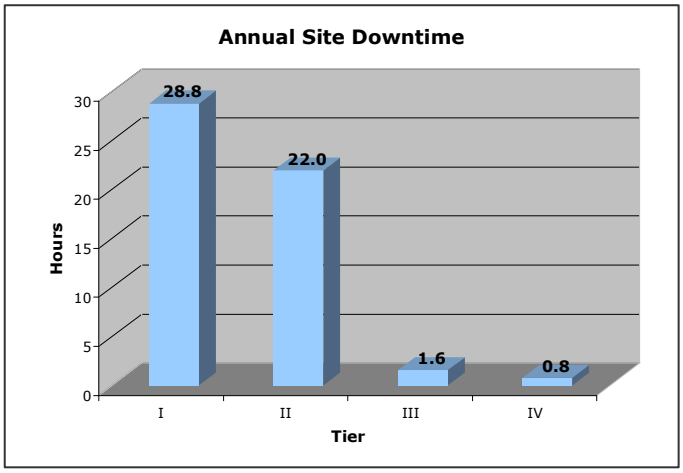


The obvious first question is, what would cause a Tier IV site to suffer an annual downtime of 0.8 hours? The answer is:

*0.8 hours per year = 4 hours every five years*

This allows for the fact that it would take IT clients 4 hours to recover from even a momentary power outage, such as one caused by a fire alarm or inadvertent operation of an EPO (Emergency Power Off) system.

This information allows IT executives charged with securing data center premises to perform cost-benefit analyses upon which sound decisions can be based. The quantum leap in reliability is clearly between Tier II and Tier III sites. The incremental reliability advantage of Tier IV sites over Tier III sites entails a disproportionate increase in cost. This explains why design topologies that are Tier II-compliant or Tier III-compliant are deployed at most third-party data centers built today on a speculative basis by data center space providers with a national portfolio.

Which Tier an organization needs depends on business requirements. Companies that can justify the cost of extremely high availability by the higher cost of data center outage-related disruption should opt for Tier IV space. For companies with high availability requirements that can tolerate the cost of disruption due to unplanned events, Tier III facilities are appropriate. Tier II may suffice for companies that will not suffer from a disruption due to data center outages because their revenue stream is not dependent on real-time products or services delivery, or start-up companies that must restrict IT spending in order to allocate the lion's share of venture capital funds to developing their core competencies. Tier I facilities are the domain of companies with low availability requirements and little or no IT-dependent revenue stream.

## Energy Efficiency

Data centers are notoriously power-dense. Whereas a typical office building consumes on average under 10Watts/sf, data centers routinely consume 300+Watts/sf, and in some cases much higher. Data center energy efficiency affects both a facility's operational cost and carbon footprint ("greenness").

The primary metric used to measure data center energy efficiency is Power Usage Effectiveness, PUE, which is the ratio of total power consumed by a data center to the power consumed by the IT equipment deployed in it. A legacy data center might have a PUE as high as 2.5 or more. A newer-vintage design utilizing an air economizer-based cooling system and other data center best practices (e.g. blanking panels and cold aisle containment) can record a PUE in the 1.3 range and lower, depending on geographical location and year-round weather patterns.

The Green Grid, an industry group concerned with data center energy efficiency, promotes the use of Data Center Infrastructure Efficiency (DCIE) as the energy efficiency metric. DCIE is simply the reciprocal of PUE.

As the industry has matured, a wide range of best practices have been identified and adopted that improve the efficiency of data center infrastructure. Nowhere has the positive impact been more pronounced than in cooling systems. Examples include:

**Examples of "best practice" in lowering PUE**
- Air economizer-based cooling system
- Blanking panels
- Cold aisle containment

- the configuration of equipment racks in layouts that provide distinct cold (supply air) aisles and hot (exhaust air) aisles
- the use of ceiling plenums as a contained return air path with ducts connecting the plenums to the cooling units' return air intake grilles
- the deployment of air curtains for cold aisle containment, thus ensuring the cold supply air is delivered only where it is needed, at the front of equipment chassis, and preventing the cold air from bypassing the equipment and flowing directly into the exhaust air stream

## The Vendor

Whether considering data center services directly with a mission-critical facilities development and management company or through a hosting provider, there are many factors a customer must consider when choosing who they will trust with their IT infrastructure over the years to come.

### Service Level Agreements

Service Level Agreements (SLA) are a critical component of data center collocation contracts. They articulate the recourses (e.g. billing rebates) available to the customer in the event of non-performance of the data center mission-critical infrastructure or systems. They cover the elements that are crucial to availability (e.g. power, cooling, network performance as measured by latency and packet loss) as well as parameters like response time for trouble tickets or remote hands requests.

Customers should read and understand the SLA language in their contracts. A weak (customer-unfriendly) SLA suggests the provider has a low level of confidence in the deployed mission-critical infrastructure and systems. The higher the tier level, the stronger should be the SLA.

### Certification

Third party certification, such as SAS70-II, is a must for some data center end-users. A certification report presents the results of reputable third party audits of the systems and controls deployed in running and securing the data center. It spells out what controls are in place and how effective they are in maintaining the desired operational standards.

SAS 70

Type II

### Global versus Regional

For companies with data center requirements that extend overseas to different countries or continents, picking a provider with a global footprint may result in levels of standardization of facilities, services and products (e.g. SLA terms) that would be rare when dealing with different vendors in different countries. Standardization should result in a more streamlined and consistent collocation experience, regardless of location.

### Financial Stability

As with all mission-critical systems vendors, careful consideration should be given to a provider's financial health. Moving IT infrastructure from one facility to another can be an expensive, disruptive and stressful exercise. Before signing a 2, 3 or 5-year contract, be sure your vendor will be around for at least the contractual term.

## Summary

Picking an enterprise hosting data center in which to deploy IT and communications infrastructure systems can be a challenging endeavor. Simplify your due diligence by focusing on the key parameters:

| | |
|---|---|
| **Data center reliability/availability** | Is the tier level compatible with your organization's outage-tolerance profile? |
| **Data center energy efficiency** | How green is the building? What is the facility's PUE/DCIE? What are the monthly recurring power and cooling system costs? |
| **Service Level Agreements** | Is a strong SLA offered, one that appropriately reflects the facility's tier level compatibility? |
| **Data center certification** | Read the certification reports to learn how effectively the site is run and managed. If the site is not certified, step up your due diligence. |
| **Geographic footprint of the vendor** | Are your vendor's offerings as geographically diverse as your needs? |
| **Vendor's financial stability** | Will the vendor be around for the long run? |


*Tarif Abboushi, P.E. is Vice-President of Data Center Infrastructure at NTT America, Inc. He has been involved in the evaluation, acquisition, design, construction, testing and commissioning, maintenance and operation of mission-critical facilities for almost twenty years, either as a consulting engineer or with a data center owner-operator. He holds a BSEE from Manchester University (UK) and a Master's in Engineering-Economic Systems from Stanford University.*

*NTT America is North America's natural gateway to the Asia-Pacific region, with strong capabilities in the U.S. market. NTT America is the U.S. subsidiary of NTT Communications Corporation, the global data and IP services arm of a Fortune Global 500 telecom leader: Nippon Telegraph & Telephone Corporation (NTT). NTT America provides world-class Enterprise Hosting, managed network, and IP networking services for enterprise customers and service providers worldwide. For additional information on NTT America, visit us on the Web at www.nttamerica.com.*



NTT Communications

NTT America
www.nttamerica.com