Explain: Tier 1 / Tier 2 / Tier 3 / Tier 4 Data Center

by VIVEK GITE *on JUNE 7, 2008* · 43 COMMENTS· *last updated at* JANUARY 29, 2011

Q. What is data center tiers? What is tier 1 data center? Which tier / level is the best for maximum uptime?

A. Tier 1 to 4 data center is nothing but a standardized methodology used to define uptime of data center. This is useful for measuring:

a) Data center performance
b) Investment
c) ROI (return on investment)
b)

Tier 4 data center considered as **most robust and less prone** to failures. Tier 4 is designed to host mission critical servers and computer systems, with fully redundant subsystems (cooling, power, network links, storage etc) and compartmentalized security zones controlled by biometric access controls methods. Naturally, the simplest is a Tier 1 data center used by small business or shops.

- **Tier 1** = Non-redundant capacity components (single uplink and servers).
- **Tier 2** = Tier 1 + Redundant capacity components.
- **Tier 3** = Tier 1 + Tier 2 + Dual-powered equipments and multiple uplinks.
- **Tier 4** = Tier 1 + Tier 2 + Tier 3 + all components are fully fault-tolerant including uplinks, storage, chillers, HVAC systems, servers etc. Everything is dual-powered.

Data Center Availability According To Tiers

The levels also describes the availability of data from the hardware at a location as follows:

- Tier 1: Guaranteeing 99.671% availability.
- Tier 2: Guaranteeing 99.741% availability.
- Tier 3: Guaranteeing 99.982% availability.
- Tier 4: Guaranteeing 99.995% availability.

Further recommended readings:

- Uptime Institute
- TIA-942 Data Center Standards Overview - 102264AE
- Telecommunications Infrastructure Standard for Data Centers

Data Center Tier Standards | Tier 1-4 Overview | Colocation America

Data center tier standards exist to evaluate the quality and reliability of a data center's server hosting ability. The Uptime Institute uses a somewhat mysterious four-tier ranking system as a benchmark for determining the reliability of a data center. This proprietary rating system begins with Tier I data centers, which are basically warehouses with power and ends with Tier IV data centers, which offer 2N redundant power and cooling in addition to a **99.99% uptime guarantee**.

A Tier III **data center** is concurrently maintainable, allowing for any planned maintenance activity of power and cooling systems to take place without disrupting the operation of computer hardware located in the data center. In terms of redundancy, Tier III offers "N+1" availability. Any unplanned activity such as operational errors or spontaneous failures of infrastructure components can still cause an outage. In other words, Tier III isn't completely fault tolerant. A Tier IV data center is fault-tolerant, allowing for the occurrence of any unplanned activity while still maintaining operations. Tier IV facilities have no single points of failure. The basic concept is that a Tier IV design requires double the infrastructure of a Tier III design. Note that both Tier III and Tier IV data center specifications require IT equipment to have dual power inputs to permit maintenance of power distribution components between the UPS and IT equipment.

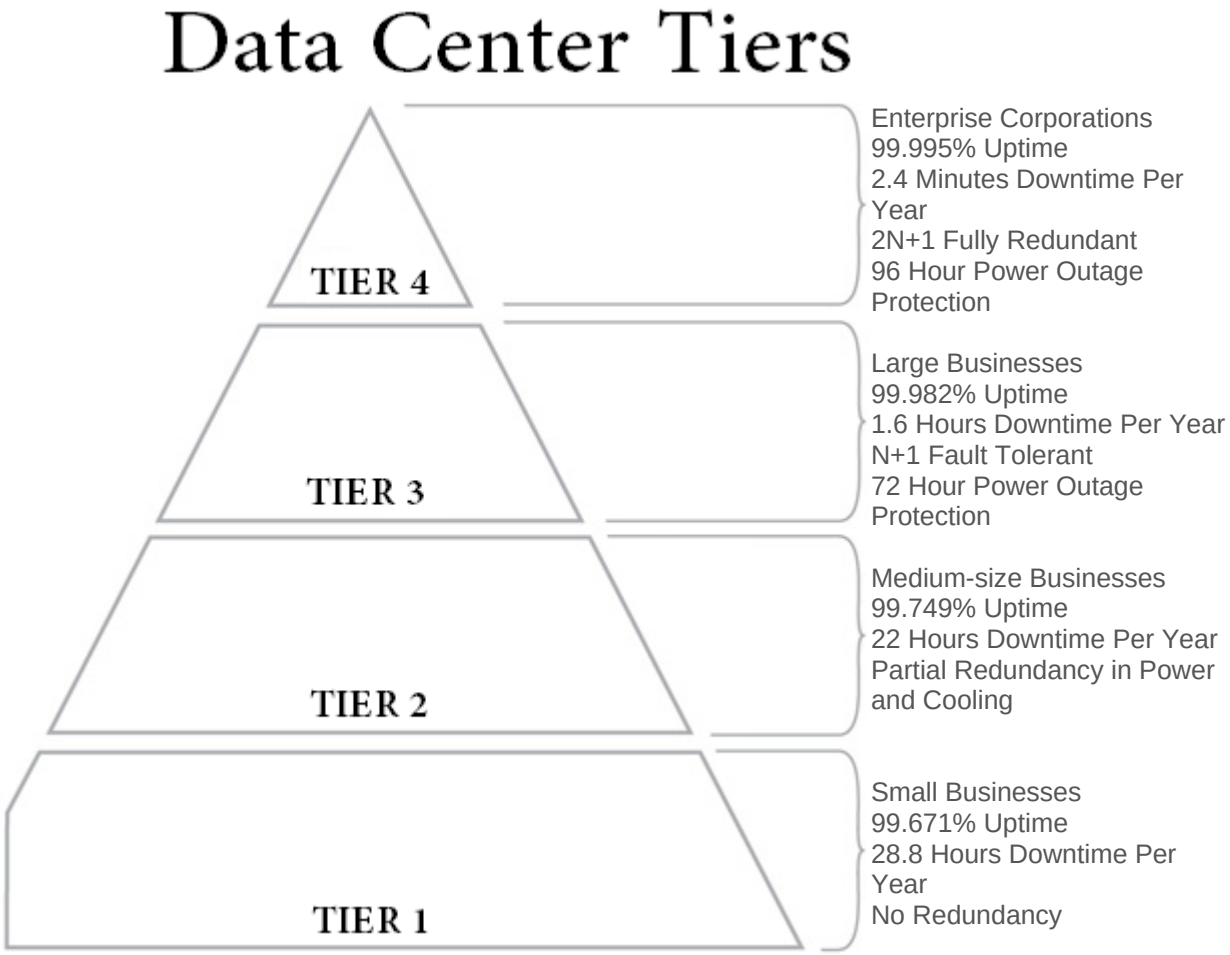

Unfortunately, the Uptime Institute has chosen not to fully publish the evaluation criteria for these different tier levels. Few data centers have tier certifications from the Uptime Institute. Only 38 facilities or design documents for facilities have official tier certifications at this point; these are primarily enterprise data centers. The result is that the Uptime Institute's definitions have been misused by the industry, ignorantly in many cases. Facility builders, designers and

owners have tried to tweak the terminology slightly to give it their own unique flavor. Enterprises should question any Tier IV claims by data center providers because it is difficult to get customers to pay the rates necessary to monetize the Tier IV investment of roughly double that of a Tier III facility.

Tier I Research has established a bi-level data center quality rating system, with the following criteria:

Premium Data Center:

Facility must meet N+1 redundancy standards in each of the following areas: uninterruptible power supply (UPS), backup generators and HVAC systems. Preferably, the facility should be on multiple power grids. The facility must have 24/7 staffing and continuous video surveillance. Biometric security is preffered, but not required. VESDA (or better) fire detection and dry-pipe pre-action or gas-based fire suppression are required. Multiple fiber providers must serve the facility through two or more entrance facilities. Data center resiliency is based not only on physical power and cooling infrastructure, but also on network communications redundancy and fire detection and suppression. These are key attributes missing from the Uptime Institute's data center tier ratings, which focus on the power and cooling infrastructure.

Standard Data Center:

Facilities with sufficient power and cooling capacity and basic fire suppression are in this category, as are data centers with non-redundant HVAC and UPS units and only a single backup generator. The bottom line is that data center customers are looking for high availability and security for IT equipment, which implies redundancy for power, cooling and network connections; **fire detection**and suppression; the ability for concurrent maintainability and the 24x7 staffing and security found in a premium data center.

Data Center Tiers

What exactly is a data center?

A commercial **data center (or datecenter or data centre)** is a facility designed to house computer servers at a climate-controlled environment. It is built to withstand calamities and environmental disasters such as storms or fires. Data centers are equipped with redundancy systems that can constantly produce clean power such as uninterrupted power systems (UPS). The best data centers ensure continuous operation in virtually any circumstance.

Enterprise class data centers are ideal for large organizations that need to share data between various divisions that do not share the same physical presence. By hosting data in a climate-controlled environment, you can minimize the cost of sharing data between various divisions of the same organizations. A **wholesale data center** provider rents / leases space to businesses and individuals who wish to host data in a controlled environment. This kind of arrangement can disperse the cost of building storm-resistant and climate-controlled structures, thus providing all tenants with maximum reliability.

Since data centers are designed to house sensitive computing equipment, they usually have several elements in common. Each **data center** comes with physical security, fire suppression, environment control, stand-by power, and plenty of computing equipment. As the center of computer operations, datacenters need to be properly designed to minimize the risk of losses due to disasters and facilitate efficient disaster recovery efforts as well.

Standards for data center tiers

The **Uptime Institute**, an organization based in Santa Fe, New Mexico, defines and holds the copyright to the four Tier levels, the Tier Classification (Tier I, Tier II, Tier III, Tier IV) were created to describe site-level infrastructure required to sustain data center operations.

Tier I Data Center

Now that medium-sized businesses continue to experience growth, they have become perfect candidates for the huge storage systems that a **Tier I Data Center** offers. It claims better capacity, reliability, performance, and manageability over a simple office setting.

Tier classifications also focus on the availability of facilities based on their connectivity, cooling components, and power-generating equipment. They describe the degree to which a certain facility remains resilient to failures of electrical, cooling, and technical systems. Resilience to such failures is

also determined by the architecture and redundancy of the overall design of facilities. In general, a **Tier I Data Center** is considered as the least.

Tier 1 Basic Data Center

Tier-1 facilities do not have redundant capacity components. They only provide basic cooling and power with no excess capacity for backup. Redundancy in the MEP distribution paths is not present either. Unplanned outages or failure of distribution elements will have a impact on computer systems. The maintenance required for **data center infrastructure** when replacing components or performing utility work will affect the facility as if there were an unplanned outage.

Tier I data center typically experiences 2 different 12-hour,site-wide shutdowns and 1.2 equipment or distribution failures on average each year. The annual impact of maintenance and outages is 28.8 hours per year, or 99.67 percent availability

If your business involves various computer systems, peripherals, and other components used in telecommunications or for file/data storage, you need to have a datacenter facility to house all of these. You'll commonly find backup (technically called redundant) power supplies, connections for effective data communication, a controlled environment, and various security measures to make sure that everything you have in the data center is kept safe. But not all datacenters are alike, because some are built more complex than others. Still, they follow an effective identifying system to determine how these datacenters are built and designed, and that's by using four different tiers, with the first as the simplest, and the fourth, usually the most advanced. This article will discuss the basics of **Tier 2 datacenter**.

What is **Tier II data center**?

The most basic datacenter classification available is the **Tier 1** datacenter. It is defined to just contain one path for the distribution,with no backup components in case of failure. Tier I datacenter typically experience 28.8 hours of outages per year, or 99.71% availability.

On the other hand, the **Tier 2 datacenter** could be considered as an upgraded form of the Tier 1 datacenter. It also contains a single path for distribution but has some redundant components, which can be removed for planned maintenance without impacting IT equipment. On average tier II sites have 1 unplanned outage each year that affect the IT operations' environment. The annual impact of maintenance and unplanned outages is 22 hours per year or 99.75% availability.

Tier III Data Center

While browsing through web hosting providers, you will eventually come across a company that talks about how they run a state-of-the-art **Tier 3 data center** without any explanation of what it means. Most hosting providers tend to use the term loosely so it can be hard to understand how each of them operates. It is important to understand what tier classifications stand for because they can help you choose a reputable hosting company that can satisfy your business needs.

The Uptime Institute, an organization that specializes in enhancing the reliability of data centers, developed the tier classifications. They are determined through a score that demonstrates how redundant data centers are and what their uptime guarantees are. In this system, **Tier 1** is considered as the least reliable and **Tier 4** as the most consistent and trustworthy. The former does not offer redundancy for components and power so it can only provide a 99.67% uptime guarantee. The latter offers dual-powered systems that were designed to be fault-tolerant,thus giving them a 99.99% rate of reliability.

Tier 3 seems to offer the best of the aforementioned tiers because it provides multiple cooling and power sources,which gives data center an uptime of 99.98%. It is also cheaper to maintain than **Tier IV**, hence making it a great alternative for corporations and large organizations.
Why choose Tier 3 data centers?

Through tier classifications, designers can understand the different methods used for qualifying various aspects of data centers. It also provides clients with the chance to compare one **data center** to another in the most objective manner. Since data centers are the most critical part of any business, every organization needs to ensure the reliability of its **datacenter**. This is the main reason why data centers that were built according to **Tier III** specifications are beneficial to those who want to take advantage of a 99.98% rate of availability. Here are some of the best qualities of **Tier 3** data centers.

Tier III data centers were specifically designed to satisfy two important requirements, concurrent maintainability and redundancy. Tier 3 possesses concurrent maintainability for all cooling components and power distribution systems. When a component becomes unavailable due to maintenance issues or power outages, it should never affect the normal functioning of the entire infrastructure. This **data center** tier can also enable planned maintenance activities without disrupting the operation of computing systems. It also includes an N+1 redundancy component and multiple distribution paths for power and cooling. Operations experience shows that tier III facilities have 1.6 hours of unplanned failures, providing 99.98% percent availability.

Tier IV Data Center

To provide an objective basis for determining specific data center needs, the **Uptime Institute** established four distinctive Tier classifications. These Tier classifications describe an **infrastructure** topology required to sustain data center operation, tier IV classification builds on **tier III** by adding fault tolerance to the site-level infrastructure topology. In Tier IV fault tolerance extends to each and every system or component (eg, power generation, cooling, uninterruptible power sources etc.).

Tier 4 provides fault tolerant **data center infrastructure** topology, so that events affecting the data center computer room are reduced to 0.8 hours on an annual basis for 99.995+% availability.
Tier IV: Fault Tolerant

- Fulfills **Tier 1** and **Tier 2** and **Tier 3** requirements
- Multiple, independent, diverse, active distribution paths.
- 'N' capacity for power and cooling to the IT equipment after any infrastructure failure
- **Not susceptible to disruption from a single unplanned event.**
- 99.995+% availability

REQUIREMENTS	**Source: The Uptime Institute**			
	Tier I	Tier II	Tier III	Tier IV
Number of delivery paths (power and cooling)	Only 1	Only 1	1 active 1 passive	2 active
Redundancy	N	N+1	N+1	2(N+1) or S+S
Compartmentalization	No	No	No	Yes
Concurrently Maintainable	No	No	Yes	Yes
Fault tolerant to worse event	None	None	None	Yes

TIER REQUIREMENTS	**Source: The Uptime Institute**			
	Tier I	Tier II	Tier III	Tier IV
Building type	Tenant	Tenant	Standalone	Standalone
Staffing	None	1 Shift	1+ Shifts	"24 by forever"
Useable for critical load	100% N	100% N	90% N	90% N
Initial watts per square foot	20-30	40-50	40-60	50-80

Ultimate watts per square foot	20-30	40-50	100-150	150+
Uninterruptable cooling	None	None	Maybe	Yes
Support space:raised floor	20%	30%	80-90%	100+%
Raised floor height	12″	18″	30-36″	30-36″
Floor loading pounds per SF	85	100	150	150
Utility voltage	208,480	208,480	12-15kV	12-15kV
Single points of failure	many + human error	many + human error	many + human error	many + human error
Annual site-caused IT downtime (actual)	28.8 hrs	22.0 hrs	1.6 hrs	0.4 hrs
Site availability	99.671%	99.749%	99.982%	99.995%
Months to Implement	3	3 to 6	15 to 20	15 to 20
Year first deployed	1965	1970	1985	1995

COLOCATION | COLLOCATION | COLOCATION

Service provided by Data Center companies that allows third parties a location to house their network, server, and storage systems. The companies often provide the hardware and also offer inter-connectivity to other telecommunication providers and service providers.

WHOLESALE DATA CENTER POD

This is a fully built data center solution often within a multi-tenant data center property. Often these are fully commissioned and they are often ready for a company to move in its servers and racks.

POWERED SHELL | POWER BASED BUILDING

These are properties, typically warehouses that are in shell condition that possess a number of elements that would make for an ideal data center retrofit. This includes sufficient power to the site, multiple fiber providers and permits necessary to build a data center. Often these properties will have ample ceiling heights (over 15'), good column spacing, and heavy floor loading.

MODULAR DATA CENTERS

These systems are typically purpose built container modules that can outfit racks, servers, and cooling systems.

CLOUD COMPUTING

Cloud computing is a general term for anything that involves delivering hosted services over a network, typically the Internet. There are both public and private clouds. There are many types of public clouds but generally these services are broadly divided into three categories: Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS)and Software-as-a-Service (SaaS).

A cloud service has three distinct characteristics that differentiate it from traditional hosting. It is sold on demand, typically by the minute or the hour; a user can have as much or as little of a service as they want at any given time; and the service is fully managed by the provider (the user needs nothing but a computer and Internet access).

It is often misunderstood that the clouds will eliminate the need for data centers; as cloud computing needs to be hosted somewhere.

TIER 1, TIER II, TIER III, TIER IV DATA CENTER

This is a designation that is awarded by Uptime Institute that quantifies the data center infrastructure. The higher the Tier the greater the availability, reliability, security and other components of the data center.

Understanding “Uptime” and Data Center Tier Levels

- By: Julius Neudorfer
  March 21st, 2012

The data center industry is laden with inter-related terms such as Uptime, Tier Levels, Availability, Redundancy and Reliability. In order to make an informed decision it is important to understand what they mean and what is the actual significance for any proposals you are considering in your “Build vs Buy” analysis.

The concept of “Uptime” was pioneered by the Uptime Institute which was founded in 1993 and introduced its well defined Tier Classification system: I, II, III and IV, of which Tier IV represents the highest level of projected availability. Today, its Tier Certification system is globally recognized and its members are mostly Fortune 100-sized companies having multiple data centers averaging 50,000 square feet. And while not everyone subscribes to the Uptime Institute officially, marketing references to *Tier 2,3 or 4 are common among those seeking to ascribe a certain level of design or construction to a data center’s overall availability or system redundancy, which may, or may not be totally accurate.

The terms “N, N+1 and 2N”, typically refer to the number of power and cooling components that comprise the entire data center infrastructure systems. Wherein “N” is the minimum rating of any component (such as a UPS, generator or cooling unit) required to support the critical load. An “N” system is not redundant at all, and the failure of any component will cause an outage, effectively describing a tier 1 type facility. N+1 and 2N, represent increasing levels of component redundancies and power paths, roughly mapping to the tiers 2-4, however it is important to understand that redundant components in themselves do not guarantee continuous availability, nor insure compliance with an Uptime Institute certified data center tier level.

Moreover, besides redundancy, the ability to do planned maintenance or emergency repairs on systems may involve the necessity to take them offline. This involves the key concept of “concurrent maintainability” which permits systems to be bypassed, without impacting the availability of the computing equipment. The Uptime Institute is well recognized in clearly defining concurrent maintainability in their Tier Level system. This is one of the key criteria in the design or certification of Tier III and Tier IV data centers.

Besides the level of infrastructure redundancy at the facility level, consistency of procedures for operations, maintenance and support of the critical infrastructure systems are key to ensuring continuous availability. Moreover, the Uptime Institute has now also established a related sub-category called Operational Sustainability to define and evaluate data center’s operational procedures, as an addition to their well recognized Tier Classification rating system.

Whether you chose to build or buy, you should examine all of these factors very closely to understand what is being promised if it is a brand new building and/or what the operating history of a proposed data center facility provider has been. You can download the complete Data Center Knowledge – Data Center Build vs Buy report, which includes all the articles in the series plus additional sidebars and analysis, compliments of Digital Realty Trust.

Julius Neudorfer is the CTO and founder of North American Access Technologies, Inc. and writes for Data Center Knowledge on issues and strategies relevant to seniors business executives.

A Refresher on Data Center Tiers

Posted by Mary Beth Hamilton on Tuesday, August 3rd, 2010

At Eze Castle Integration we often reference data center tiers (i.e. Tier II and Tier III) in our written materials and assume readers will automatically understand the value of these distinctions. In some cases this might be a safe assumption, but you know what they say about assuming so we'll do a refresher in this blog post.

Data center tiers – Tier I to IV – represent a standardized method to define the uptime of a data center. The tiers are useful in measuring:

- Data center performance
- Investment
- Return on investment (ROI)

The four tiers, as classified by The Uptime Institute, include the following:

- **Tier 1:** composed of a single path for power and cooling distribution, without redundant components, providing 99.671% availability. This is the simplest and typically used by small businesses.
- **Tier II:** composed of a single path for power and cooling distribution, with redundant components, providing 99.741% availability.
- **Tier III:** composed of multiple active power and cooling distribution paths, but only one path active, has redundant components, and is concurrently maintainable, providing 99.982% availability.
- **Tier IV:** composed of multiple active power and cooling distribution paths, has redundant components, and is fault tolerant, providing 99.995% availability.

As you can see, Tier IV is considered the most robust and least prone to failures. It is designed to host mission critical servers and computer systems. The cost of colocating in a data center increases as the tier increases. Tier II and Tier III facilities generally meet the uptime requirements of most hedge funds and investment management firms.

Here is a graphical representation of the tiers

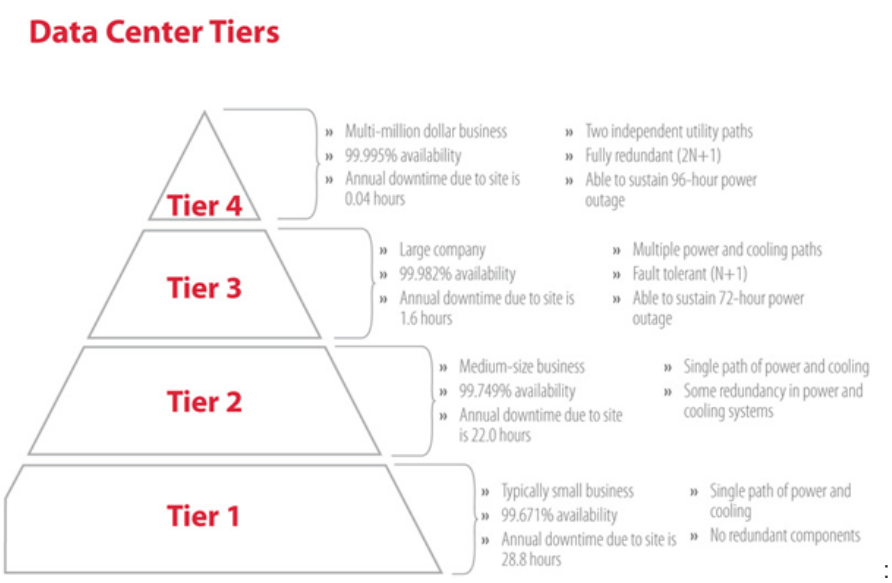

:

Data center tiers are becoming increasingly important as hedge fund and investment firms look to cloud computing infrastructures to increase agility, reduce operating costs, and simplify IT infrastructure and application management. Our Eze Cloud infrastructure is built across multiple Tier II and Tier III data centers to deliver the availability and performance investment firms require.

As a hedge fund evaluates its data center options for colocation or explores why cloud computing may be right for the hedge fund it is essential to understand the differences between each tiers. We hope this article was helpful in defining the distinctions.

Also, be sure to checkout our article on hosted business applications and what hedge funds must consider.

Ten Unique Data Center Designs

Posted By Colleen Miller On July 29, 2010 @ 1:35 pm In Data Center Design,EvoSwitch | 6 Comments

Across the globe, many data center designs use an amazing amount of creativity in locating and deploying unique data centers. These are not your run-of-the-mill server rooms, but truly one-of-a-kind facilities that have a personality all their own.

There is more to data center design than the straightforward mechanical approach, according to Rob Snevely, author of *Enterprise Data Center Design and Methodology* [1]. "The detailed process of data center design appears on the outset to be a purely mechanical process involving the layout of the area, computations to determine equipment capacities, and innumerable other engineering details," Snevely writes. "They are, of course, essential to the design and creation of a data center, however, the mechanics alone do not a data center make. The use of pure mechanics rarely creates anything that is useful, except perhaps by chance." He goes on to discuss design philosophies that go back to classical times.

The data centers we present below have gone beyond the ordinary, incorporating innovative methods of using the layout, building or environment to create unique data center designs that work for the 21st century.

“The Spy Who Loved Me” Bunker

The Bahnhof “James Bond Villain” Data Bunker in Stockholm, Sweden [2]

[2]Nestled in the granite rock of the Vita Berg Park 100 feet beneath Stockholm, a unique facility operated by Bahnhof, one of Sweden’s largest ISPs, has multiple unusual features: greenhouses, waterfalls, German submarine engines, simulated daylight and the ability to withstand a hit from a hydrogen bomb. Albert France-Landrd, the architects who designed the project in a former military anti-atomic-bomb shelter, stated that the starting point of the project was to consider the rock as a living organism. From the space-themed conference room to the NOC with a jungle-theme featuring many lush green plants, this data center is truly tricked-out.

Super-Sized Data Center

Microsoft Container Data Center in Chicago [3]

[4]Encompassing more than 700,000 square-feet, Microsoft’s $500 million Chicago data center is one of the largest data centers ever built. And it’s highly unique as well. The lower level houses 40-foot former shipping containers packed with web servers in a garage-like setting, while a second story houses traditional raised-floor data center space. Parking spaces in the ”container canyon” are occupied by containers packed with servers and, in some cases, equipment to power and cool the servers. A central spine provides hookups to the containers on the lower level and power distribution equipment is located on the mezzanine level. With double-stacked containers, portable stairs are used to provide access to the second-level container.

The Greenest of the Green

IBM Green Data Center, campus of Syracuse University [5]

[5]IBM calls this data center at Syracuse University its “greenest” facility to date. Although it is small in size at 12,000 square foot, the facility uses a combination of approaches to reduce power and cooling impacts. The facility includes: micro-turbine engines fueled by natural gas, absorption chillers, the reuse of waste heat in nearby buildings, DC power distribution, and rear-door liquid cooling for each cabinet. Through running this data center, IBM and Syracuse are developing data on the most effective way to use these technologies in tandem in a data center for the best collective power savings.

Saving Power and Water

Microsoft Data Center in Dublin, Ireland [6]

[6]Microsoft’s new data center in Dublin uses far less energy and water than typically consumed in other data centers of this scale. The facility will power much of Microsoft’s global cloud computing operation. The Dublin center operates at a Power Usage Effectiveness (PUE) of 1.25, Microsoft says, compared to averages of about 2.0 for the industry and

1.6 for other Microsoft facilities. The design innovations driving this project's efficiency include a "free cooling" system that uses outside air to cool the data center and a server pod design that employs hot-aisle containment to support warmer operating temperatures inside the server space.

Circular Data Center in a Silo

The CLUMEQ Supercomputer in Quebec [7]

[7]The CLUMEQ supercomputing center in Quebec partnered with Sun Microsystems to convert a huge silo into a data center. Previously housing a Van de Graaf particle accelerator, the cylindrical building is 65 feet high and 36 feet wide with two-foot thick concrete walls. The CLUMEQ Colossus cylinder features an interior "hot core" (as opposed to a hot aisle) in the center of the building and uses the outside ring of the facility as the cold air plenum. The cabinets are arranged in a ring on each floor, facing the outside of the silo. The floors supporting each ring of cabinets are comprised of grates rather than solid flooring to facilitate airflow through the facility.

Free Cooling With Fresh Sea Air

HP Data Center in Wynyard, England [8]

[8]HP's innovative new data center in Wynyard, England leverages cool fresh air from the North Sea making extensive use of free cooling – the use of outside air rather than refrigeration to keep servers cool. Although many large data centers are now using free cooling, the HP Wynyard facility features an innovative airflow scheme that uses the entire lower floor of the facility as a cooling plenum. The Wynyard data center, which was begun by EDS prior to its acquisition by HP, is notable for its implementation of multiple energy efficiency measures.

Rock On! Underground Mine in Pennsylvania

Iron Mountain's Energy Efficient Bunker [9]

[10]Iron Mountain, a traditional leader in document storage, uses underground "nuke proof" data bunkers to store its customers' digital data in its own data center. The huge underground facility in western Pennsylvania is part of the company's push into the data center sector. Iron Mountain bought the former limestone mine, which houses 1.7 million square feet of space, in 1998 and created an energy-efficient data center known as Room 48 in a section of the former mine.

Reduce, Reuse, Recycle

i/o Data Centers Phoenix ONE Data Center [11]

[11]The Phoenix ONE data center is an example of reusing an existing property. Its building was finished in 2005 to house the former bottling operations of LeNature's Beverages. The 538,000-square-foot facility includes more than 80,000 square feet of premium office space,

allowing i/o Data Centers to use the site as both a data center and its corporate headquarters. The company is using both solar power and thermal storage to save on energy costs for its huge operation.

Which Came First: Chickens or Computers?
Yahoo Computing Coop in Lockport, NY [12]

[12]Located not far from Niagara Falls in upstate New York, the Yahoo Computing Coop is the result of a five-year effort to completely rethink the company's data center design. The Yahoo team refined its approach and is building faster, cheaper data centers that will be dramatically more efficient than previous facilities.And yes, each building does look like a huge chicken coop. "The building itself is an air handler," says Scott Noteboom, the Director of Data Center Operations for Yahoo. "The entire building is meant to breathe, and there's a lot of louvers and dampers to control the airflow."

Free to Be: Carbon Free!
EvoSwitch Climate Neutral Data Center in Amsterdam [13]

[13]Being both carrier neutral and carbon neutral is the strategy behind EvoSwitch, located in the Netherlands. EvoSwitch is doubling the size of its Amsterdam data center to meet customer demand for space in its facility. The 100,000 square foot (9,285 square meter) data center houses the European infrastructure for Wikipedia. EvoSwitch is supported by 20 megawatts of power capacity that is generated entirely from sustainable sources by Delta power, which uses a mix of solar, wind and biomass in its power generation.

---

[7] The CLUMEQ Supercomputer in Quebec: **http://www.datacenterknowledge.com/archives/2009/12/10/wild-new-design-data-center-in-a-silo/**

[8] HP Data Center in Wynyard, England: **http://www.datacenterknowledge.com/inside-hps-green-north-sea-data-center/**

[9] Iron Mountain's Energy Efficient Bunker: **http://www.datacenterknowledge.com http://www.datacenterknowledge.com/iron-mountains-energy-efficient-bunker/**

[10] Image: **http://www.datacenterknowledge.com/iron-mountains-energy-efficient-bunker/**

[11] i/o Data Centers Phoenix ONE Data Center: **http://www.datacenterknowledge.com/an-inside-look-at-phoenix-one-data-center/**

[12] Yahoo Computing Coop in Lockport, NY: **http://www.datacenterknowledge.com/archives/2010/04/26/yahoo-computing-coop-the-shape-of-things-to-come/**

[13] EvoSwitch Climate Neutral Data Center in Amsterdam: **http://www.datacenterknowledge.com/inside-evoswitch-climate-neutral-data-center/**