



Industry Standard Tier Classifications Define Site Infrastructure Performance

By W. Pitt Turner IV, P.E., John H. (Hank) Seader, P.E. and Kenneth G. Brill

One of the most common sources of confusion in the field of uninterruptible uptime is what constitutes a reliable data center. All too often, reliability is in the eye of the beholder: what is acceptable to one person or company is inadequate to the next. Competing companies with data centers of radically different infrastructure capabilities are all claiming to deliver high availability.

With the continuously increasing pressure on high availability and the explosive growth of the Internet comes an increased demand for computer hardware reliability. Information technology customers expect availability of “Five Nines” or 99.999%. Unfortunately, the substantial investment a business frequently makes to achieve Five Nines, in its computer hardware and software platforms, is likely to be insufficient unless matched with a complementary site infrastructure (power, cooling, and other environmental support systems) that can support their availability goals.

The Uptime Institute, Inc.® (*The Institute*) developed a tiered classification approach to site infrastructure functionality that addresses the need for a common benchmarking standard. *The Institute*’s system has been in practice for 10 years. It includes actual measured availability figures for site availability ranging from 99.67% to more than 99.99%. It is important to note that this range of availability is substantially less than the current Information Technology (IT) expectations for Five Nines, which leads to the conclusion that site availability gates overall IT availability.

Over the last 40 years, data center infrastructure designs have evolved through at least four distinct stages, which are captured in *The Institute*’s classification system. Tier I first appeared in the early 1960s, Tier II in the 1970s, Tier III in the late 1980s and early ’90s, and Tier IV in 1994 with the United Parcel Service Windward project, which was the first site to assume the availability of dual-powered computer equipment. *The Institute* participated in the development of Tier III concepts and pioneered the creation of Tier IV.

Invention of Tier IV was made possible by Ken Brill, Executive Director of *The Institute*, who, in 1991, envisioned a future when all computer hardware would come with dual power inputs (US Patent 6,150,736). During construction of the \$50 million Windward project, United Parcel Service worked with IBM and other computer hardware manufacturers to provide dual-powered computer hardware. The significance of Mr. Brill’s insight has subsequently been confirmed by billions of dollars in site infrastructure investment.

Dual-power technology requires at least two completely independent electrical systems. These dual systems supply power via diverse power paths to the computer load, by effectively moving the last point of electrical redundancy from the Uninterruptible Power Supply (UPS) system downstream to a point inside the computer hardware itself. Mr. Brill’s intuitive conclusion

has since been confirmed by *The Institute*’s research that has determined that 98% of all site infrastructure failures occur between the UPS and the computer load. Since completion of the Windward project in 1994, System plus SystemSM (S+S) Tier IV electrical designs have become common and the number of computer hardware projects with dual inputs has grown.

The advent of dual-powered computer hardware in tandem with Tier IV electrical infrastructure is an example of site infrastructure design and computer hardware design simultaneously achieving higher availability. Even with the significant improvements in computer hardware design made over the past 10 years, many data centers constructed in the last five years—and even today claim Tier IV functionality, but actually deliver only Tier I, II, or III—are falling behind in their capability to match the availability required by the information technology they support. The purpose of this paper is to outline what functionality and attributes are required for the different tier levels.⁷

Defining the Tiers

The tier classification system involves several definitions. A site that can sustain at least one unplanned, worst-case infrastructure failure with no critical load impact is considered fault tolerant. A site that is able to perform planned site infrastructure activity without



shutting down critical load is considered concurrently maintainable (fault tolerance level may be reduced during concurrent maintenance). It is important to remember that a typical data center site is composed of at least 20 major mechanical, electrical, fire protection, security and other systems, each of which has additional subsystems and components. All of these must be concurrently maintainable and/or fault tolerant for the site to meet the requirement of fault tolerant and/or concurrently maintainable.

Some sites built with fault-tolerant S+S electrical concepts failed to incorporate the mechanical analogy, which involves dual mechanical systems. Such sites are classified Tier IV electrically, but only achieve a Tier III mechanically. Another common mistake is only looking at first level failures and not the subsequent failures that will sometimes be triggered by the first failure.

The following list summarizes the high level characteristics of each tier. The availability numbers shown are actual numbers for many sites which combine both tier requirements as well as the associated tier attributes.

■ **Tier I**

Tier I is composed of a single path for power and cooling distribution, without redundant components, providing 99.671% availability.

■ **Tier II**

Tier II is composed of a single path for power and cooling distribution, with redundant components, providing 99.741% availability.

■ **Tier III**

Tier III is 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**

Tier IV is composed of multiple active power and cooling distribution paths, has redundant components, and is fault tolerant, providing 99.995% availability.

Tier I Data Center Infrastructure
Basic Data Center

A Tier I data center is susceptible to disruption from both planned and unplanned activity. It has computer power distribution and cooling, but it may or may not have a raised floor, a UPS, or an engine generator. The critical load on these systems is up to 100% of N. If it does have UPS or generators, they are single-module systems and have many single points-of-failure. The infrastructure should be completely shut down on an annual basis to perform preventive maintenance and repair work. Urgent situations may require more frequent shutdowns. Operation errors or spontaneous failures of site infrastructure components will cause a data center disruption.

Tier II Data Center Infrastructure
Redundant Components

Tier II facilities with redundant components are slightly less susceptible to disruptions from both planned and unplanned activity than a basic data center. They have a raised floor, UPS, and engine generators, but their capacity design is N+1, which has a single-wired distribution path throughout. Critical load is up to 100% of N. Maintenance of the critical power path and other parts of the site infrastructure will require a processing shutdown.

Tier III Data Center Infrastructure
Concurrently Maintainable

Tier III level capability allows for any planned site infrastructure activity without disrupting the computer hardware operation. Planned activities include preventive and programmable maintenance, repair and replacement of components, addition or removal of capacity components, testing of components and systems, and more. For large sites using chilled water, this means two independent sets of pipes. Sufficient capacity and distribution must be available to simultaneously carry the load on one path while performing maintenance or testing on the other path. Unplanned activities such as errors in operation or spontaneous failures of facility

This chart illustrates Tier requirements.

	Tier I	Tier II	Tier III	Tier IV
Number of Delivery Paths	Only 1	Only 1	1 Active 1 Passive	2 Active
Redundancy	N	N + 1	N + 1	S + S or 2 (N + 1)
Compartmentalization	No	No	No	Yes
Concurrently Maintainable	No	No	Yes	Yes
Fault tolerance to Worst Event	None	None	None	Yes



This chart illustrates the tier attributes of the sites from which the actual availability numbers were drawn.

	Tier I	Tier II	Tier III	Tier IV
Building Type	Tenant	Tenant	Stand-alone	Stand-alone
Staffing	None	1 Shift	1 + Shifts	“24 by Forever”
Useable for Critical Load	100% N	100% N	90% N	90% N
Initial Gross Watts per Square Foot (W/ft²) (typical)	20-30	40-50	40-60	50-80
Ultimate Gross W/ft² (typical)	20-30	40-50	100-150 ^{1,2,3}	150+ ^{1,2}
Uninterruptible Cooling	None	None	Maybe	Yes
Support Space to Raised-Floor Ratio	20%	30%	80-90% ²	100+%
Raised-Floor Height (typical)	12”	18”	30-36” ²	30-36” ²
Floor Loading lbs/ft² (typical)	85	100	150	150
Utility Voltage (typical)	208, 480	208, 480	12-15 kV ²	12-15 kV ²
Single Points-of Failure	Many + human error	Many + human error	Some + human error	None + human error
Annual Site-Caused IT Downtime (actuals)	28.8 hours	22.0 hours	1.6 hours	0.4 hours
Site Availability	99.671%	99.749%	99.982%	99.995%
Months to Implement	3	3-6	15-20	15-20
Year First Deployed	1965	1970	1985	1995
Construction Cost (± 30%)^{1,2,3}				
Raised Floor	\$220/ft ²	\$220/ft ²	\$220/ft ²	\$220/ft ²
Useable UPS Output	\$10,000/kW	\$11,000/kW	\$20,000/kW	\$22,000/kW

¹ 100 W/ft² maximum for air-cooling over large areas, water or alternate cooling methods greater than 100 W/ft² (added cost excluded).
² Greater W/ft² densities require greater support space (100% at 100 W/ft² and up to 2 or more times at greater densities), higher raised floor, and, if required over large areas, medium voltage service entrance.
³ Excludes land; architectural, engineering, and commissioning fees; permits and other fees; interest; and abnormal civil costs. These can be several million dollars. Assumes minimum of 15,000 ft² of raised floor, architecturally plain, one-story building, with power backbone sized to achieve ultimate capacity with installation of additional components or systems. Make adjustments for NYC, Chicago, and other high cost areas.

infrastructure components will still cause a data center disruption. The critical load on a system does not exceed 90% of N. Many Tier III sites are designed with planned upgrades to Tier IV when the client’s business case justifies the cost of additional protection. The acid test for a concurrently maintainable data center is the ability to accommodate any planned work activity without disruption to computer room processing.

Tier IV Data Center Infrastructure Fault Tolerant

Tier IV provides site infrastructure capacity and capability to permit any planned activity without disruption to the critical load. Fault-tolerant functionality also provides the ability of the site infrastructure to sustain at least one worst-case, unplanned failure or event with no critical load impact. This requires simultaneously active distribution paths, typically in S+S configuration. Electrically, this means two separate UPS systems in which each system has N+1 redundancy. The combined critical load on a system does not exceed 90% of N. Because of fire and electrical safety codes, there will still be downtime exposure due to fire alarms or persons initiating an Emergency Power Off

(EPO). Tier IV requires all computer hardware have dual power inputs as defined by *The Institute’s* Fault Tolerant Power Compliance Specifications Version 2.0, which can be found at www.uptimeinstitute.org. The acid test for a fault tolerant data center is the ability to sustain an unplanned failure or operations error without disrupting computer room processing. In consideration of this acid test, compartmentalization requirements must be addressed.

Solving Incompatible Five Nines Expectations

Even a fault tolerant and concurrently maintainable Tier IV site will not satisfy an IT requirement of Five Nines (99.999%) uptime. The best a Tier IV site can deliver over time is 99.995%. This assumes a site outage occurs only as a result of a fire alarm or EPO and that such an event occurs not more than once every five years. Only the top 10 percent of Tier IV sites will achieve this level of performance. Unless human activity issues are continually and rigorously addressed, at least one additional failure is likely over five years. While the site outage is assumed to be



instantaneously restored (which requires “24 by forever” staffing), it can still require up to four hours for IT to recover information availability.

Tier IV’s 99.995% uptime is an average calculated over five years. An alternative calculation using the same underlying data is 100% uptime for four years and 99.954% for the year in which the downtime event occurs.

Higher levels of site uptime can be achieved by protecting against accidental activation of the real need for fire protection and EPOs. Preventatives include high sensitivity smoke detection, limiting fire load, signage, extensive training, staff certification, limited number of non-staff in critical spaces, and treating employees and

contracted staff well to increase pride in their work. All of these measures, if taken, can reduce the risk of failures.

Other solutions include placing the redundant parts of the IT computing infrastructure in different site infrastructure compartments so that a site infrastructure event cannot simultaneously affect all IT systems. Another alternative is focusing special effort on business-critical and mission-critical applications so they do not require four hours to restore. These operational issues can improve the availability offered by any data center, and are particularly important in a Four Nines Tier IV data center housing IT equipment that requires Five Nines availability.

This chart illustrates how these ideas are mapped over the architecture of site infrastructure.

	RELIABILITY	AVAILABILITY	SERVICEABILITY
Information Technology	Clustering RAID and DASD Token Ring Console Automation Change Management	Logical Partitions Clustering Mirrored Data Hot Backup Business Continuity	Hot Pluggable Hot Microcode Updates “Call Home” Remote Service
Electrical Infrastructure	UPS Dual Power S + S	Engine Generator Dual Power S + S	Engine Generators Dual Power S + S
Mechanical Infrastructure	Redundant Components Fans and Pumps on UPS	Thermal Storage	Dual Pipe Thermal Storage
Facility Operations	Passive Automation Change Management MAPS/Certification Simulation	24 by “Forever” Staffing Compartmentalization Failure Bypass Options On-Site Spares	Work Performed during Regular Hours In-House Knowledge In-House Supervision

Authorship

Pitt Turner is a professional engineer, a distinguished fellow of *The Institute*, and a Principal in ComputerSite Engineering, Inc[®]. He has guided more than \$1.6 billion in site infrastructure investment for primarily Fortune 50 clients.

Hank Seader developed the original idea for the Tier concept. At the time, he was a facility manager for a major data center and wanted a simple way to convey complex reliability concepts to his senior management. Currently, Hank is a member for the ComputerSite Engineering team.

Ken Brill is Executive Director of *The Institute*, and a Principal in ComputerSite Engineering. He is the founder of the Site Uptime Network[®] and the inventor of dual power distribution technology for high availability data centers.

This white paper may be quoted, with recognition of *The Institute*, or reproduced and distributed in its entirety at no charge. *The Institute* exclusively reserves the right to certify and determine the tier ranking of site infrastructures as summarily defined in this white paper. *The Institute*’s comprehensive site certification process involves significant additional criteria beyond the information provided in this document. Sites reviewed and certified by *The Institute* can be seen at www.uptimeinstitute.org. This white paper is posted and maintained on The Uptime Institute’s website at www.upsite.com/TUIpages/tuiwhite.html.

The Uptime Institute is a pioneer in creating and operating knowledge communities for improving uptime effectiveness in data center Facilities and Information Technology organizations. The 68 members of *The Institute*’s *Site Uptime Network* are committed to achieving the highest levels of availability with many being Fortune 100 companies. They interactively learn from each other as well as from *The Institute* facilitated meetings, site tours, benchmarking, best practices, uptime effectiveness metrics and abnormal incident collection and trend analysis. From this interaction and from client consulting work, *The Institute* prepares white papers documenting best practices for use by Network members and for the broader uninterruptible uptime industry. *The Institute* also conducts sponsored research and offers insightful seminars and training in site infrastructure management.

© 2001-2005 The Uptime Institute, Inc.



Building 100, 2904 Rodeo Park Drive East, Santa Fe, NM 87505
 Fax (505) 982-8484 Phone (505) 986-3900 E-mail tui@upsite.com