

WHITE PAPER

Tier Classifications Define Site Infrastructure Performance

By W. Pitt Turner IV, PE, John H. Seader, PE, Vince Renaud, PE, and Kenneth G. Brill

Table of Contents

Executive Summary.....	3	Table 1: Tier Requirements Summary.....	13
Tier Level Information Is Divided into “Tier Performance Standards” and “Commentary” Sections	3	Table 2: Common Attributes Found in Data Centers That Are Unrelated to Tier Requirements	14
Purpose of this White Paper	3	Table 3: Comparing IT Solutions for Reliability, Availability, and Serviceability to Site Infrastructure	15
Tier Performance Standards.....	4	Illustration 1: Illustrative Electrical System Topology - Tier I.....	16
Tier I: Basic Site Infrastructure	4	Illustration 2: Illustrative Electrical System Topology - Tier II	16
Tier II: Redundant Capacity Components Site Infrastructure.....	4	Illustration 3: Illustrative Electrical System Topology - Tier III	17
Tier III: Concurrently Maintainable Site Infrastructure	4	Illustration 4: Illustrative Electrical System Topology - Tier IV	18
Tier IV: Fault Tolerant Site Infrastructure.....	5	Frequently Asked Questions	19
Tiers and Engine-Generator Plants.....	5		
Determining a Site’s Tier Rating for Design Topology	5		
Tier Commentary	6		
The Institute’s STANDARDS Are Based on a Specific Performance Outcome.....	6		
Tier Functionality Progression.....	6		
Align Business Requirements with Tiers	7		
IT Availability Success Is Dependent upon Successful, Fully Integrated Operation of All Site Infrastructure Systems	8		
Operational Sustainability: More Than Tiers	9		
Uptime Institute Tier Certification.....	10		
Conclusion.....	11		
Definition of Terms.....	11		
Referenced White Papers.....	12		
About the Authors.....	12		
About the Uptime Institute	13		
Tables and Illustrations.....	13		

Executive Summary

Internationally accepted within the uninterruptible uptime community, the Uptime Institute's (Institute) Tier Classification and Performance Standard is an objective basis for comparing the functionality, capacities, and relative cost of a particular site infrastructure design topology against others, or to compare groups of sites. This paper includes a summary of the performance-based outcome requirements of a four-Tier system and provides discussion and illustrations of each classification. Actual site availability performance, combining both design topology and Operational Sustainability (site location, topology decisions about robustness and operability, construction implementation, and site management and staffing), are presented by Tier Classification. Significant cautions about Tier misapplication are provided. Within each Classification, a number of different design topology solutions will satisfy the requirements, but some will be more robust and operator-friendly than others. This white paper focuses primarily on site infrastructure design topology, yet Operational Sustainability plays a much more significant role in the site availability actually achieved.

Tier Level Information Is Divided into “Tier Performance Standards” and “Commentary” Sections

Responding to user questions and suggestions, this white paper has been updated where appropriate and reorganized into two separate sections:

- **Tier Performance Standards** are now described in their own section, similar to the format of engineering documents. The standards focus on the definitions of the Tiers and the performance confirmation tests for determining compliance to the definitions. These are absolute criteria. Performance is measured by outcome-based confirmation tests and operational impacts. This method of measurement differs from a prescriptive design approach or a checklist of required equipment, neither of which addresses nor guarantees any particular performance outcome.
- **Tier Commentary** is in a separate section that focuses on examples and conceptual-level illustrations of the various ways to design and configure each Tier. In addition, the commentary section includes discussion and examples to aid in Tier understanding as well as information on common design topology shortfalls. Comparison tables of common data center attributes and expected availability are provided. The commentary section also offers guidance in the comprehension, design, implementation, and use of the Tier definitions.

Purpose of this White Paper

This white paper:

- Equips non-technical managers with a simple and effective means for identifying the anticipated performance of different data center site infrastructure design topologies.
- Warns decision makers that site availability is ultimately an art form combining both design topology and Operational Sustainability to optimize life-cycle cost.
- Provides IT-based definitions and the summary of the performance-based outcome requirements for each Tier level.
- Reports actual site availability performance by Tier Classification as experienced by end-users operating major sites. These calculations include human-caused site failures (70 percent of the typical total), site infrastructure equipment failures, and downtime for site maintenance. Also included is end-user downtime while computer systems are being restored as a result of a site event.
- Warns readers that component/system counts or Mean Time Between Failures (MTBF) engineering calculations are meaningless for determining Tier compliance unless these approaches also include human-factor-related failures and end-user information recovery time after a site failure. End-user impact is an essential component for calculating true site performance and is usually missed in other metrics.
- Cautions senior executives that self-proclaimed or self-certified Tier claims often turn out to be misleading, incomplete, or just plain wrong.
- Outlines the need for third-party validation of site resiliency choices (site selection, design topology, and Operational Sustainability) **before** decision makers commit to projects involving tens or hundreds of millions of dollars.
- Refers readers to the Institute's white paper *Cost Model: Dollars per kW plus Dollars per ft² of Computer Floor* which contains information on using kW as the primary cost driver in estimating relative data center construction costs.
- Warns decision makers that spending significantly less than the amounts contained in the aforementioned white paper will most likely result in failure to meet Tier requirements. Conversely, spending the budgeted amount or more does not guarantee meeting Tier requirements.
- Focuses on performance of an individual site. It is understood that higher levels of end-user availability may be attained through use of multiple sites. (Many cautions apply!)
- Provides a commentary on the common attributes found in data centers that are not related to Tier.

¹Please see Uptime Institute white paper *Operational Sustainability and Its Impact on Data Center Uptime Performance, Investment Value, Energy Efficiency, and Resiliency*.

Tier Performance Standards

Tier I: Basic Site Infrastructure

The fundamental requirement(s):

- A Tier I basic data center has non-redundant capacity components and a single, non-redundant distribution path serving the computer equipment.

The performance confirmation test(s):

- Planned work will require most or all of the site infrastructure systems to be shut down affecting computer equipment, systems, and end users.
- An unplanned outage or failure of any capacity system, capacity component, or distribution element will impact the computer equipment.
- There is sufficient capacity to meet the needs of the site.

The operational impact(s):

- The site is susceptible to disruption from both planned and unplanned activities. Operation errors or spontaneous failures of site infrastructure components will cause a data center disruption.
- The site infrastructure must be completely shut down on an annual basis to safely perform necessary preventive maintenance and repair work. Urgent situations may require more frequent shutdowns. Failure to regularly perform maintenance significantly increases the risk of unplanned disruption as well as the severity of the consequential failure.

Tier II: Redundant Capacity Components Site Infrastructure

The fundamental requirement(s):

- A Tier II data center has redundant capacity components and a single, non-redundant distribution path serving the computer equipment.

The performance confirmation test(s):

- Redundant capacity components can be removed from service on a planned basis without causing any of the computer equipment to be shut down.
- Removing distribution paths from service for maintenance or other activity requires shut down of computer equipment.
- An unplanned outage or failure of any capacity system or distribution element will impact the computer equipment. An unplanned capacity component failure may impact the computer equipment.

The operational impact(s):

- The site is susceptible to disruption from both planned activities and unplanned events. Operation errors or spontaneous failures of site infrastructure components may cause a data center disruption.

- The site infrastructure must be completely shut down on an annual basis to safely perform preventive maintenance and repair work. Urgent situations may require more frequent shutdowns. Failure to regularly perform maintenance significantly increases the risk of unplanned disruption as well as the severity of the consequential failure.

Tier III: Concurrently Maintainable Site Infrastructure

The fundamental requirements(s):

- A Concurrently Maintainable data center has redundant capacity components and multiple independent distribution paths serving the computer equipment. Typically, only one distribution path serves the computer equipment at any time.
- All IT equipment is dual powered and installed properly to be compatible with the topology of the site's architecture. See elaboration in "The operational impact(s)" section below.

The performance confirmation test(s):

- **Each and every** capacity component and element in the distribution paths can be removed from service on a planned basis without impacting any of the computer equipment.
- An unplanned outage or failure of any capacity system will impact the computer equipment.
- An unplanned outage or failure of a capacity component or distribution element may impact the computer equipment.
- There is sufficient permanently installed capacity to meet the needs of the site when redundant components have been removed from service.

The operational impact(s):

- The site is susceptible to disruption from unplanned activities. Operation errors or spontaneous failures of site infrastructure components may cause a computer disruption.
- Planned site infrastructure maintenance can be performed by using the redundant capacity components and distribution paths to safely work on the remaining equipment.
- In order to establish Concurrent Maintainability of the critical power distribution system between the UPS and the computer equipment, Tier III sites require all computer hardware have dual power inputs as defined by the Institute's Fault Tolerant Power Compliance Specification, Version 2.0. Transfer devices, such as point-of-use switches, must be incorporated for computer equipment that does not meet this specification.
- During maintenance activities, the risk of disruption may be elevated. (This maintenance condition does not defeat the Tier rating achieved in normal operations.)

Tier IV: Fault Tolerant Site Infrastructure

The fundamental requirement(s):

- A Fault Tolerant data center has multiple, independent, physically isolated systems that each have redundant capacity components and multiple, independent, diverse, active distribution paths simultaneously serving the computer equipment.
- All IT equipment is dual powered and installed properly to be compatible with the topology of the site's architecture. See elaboration in "The operational impact(s)" section below.

The performance confirmation test(s):

- A single failure of any capacity system, capacity component, or distribution element will not impact the computer equipment.
- The system itself automatically responds (self-heals) to a failure to prevent further impact to the site.
- **Each and every** capacity component and element in the distribution paths can be removed from service on a planned basis without impacting any of the computer equipment.
- Complementary systems and distribution paths must be physically isolated from one another (compartmentalized) to prevent any single event from simultaneously impacting both systems or paths.
- Continuous Cooling is required. See the Institute white paper *Continuous Cooling Is Required for Continuous Availability*.
- There is sufficient capacity to meet the needs of the site when redundant components or distribution paths have been removed from service.

The operational impact(s):

- The site is not susceptible to disruption from a single unplanned event.
- The site is not susceptible to disruption from any planned work activities.
- In order to establish Concurrent Maintainability of the critical power distribution system between the UPS and the computer equipment, Tier IV sites require all computer hardware to have dual power inputs as defined by the Institute's Fault Tolerant Power Compliance Specification, Version 2.0. Transfer devices such as point-of-use switches must be incorporated for computer equipment that does not meet this specification.
- The site infrastructure maintenance can be performed by using the redundant capacity components and distribution paths to safely work on the remaining equipment.
 - During maintenance activities, the risk of disruption may be elevated.
 - During maintenance activity where one path is shut down, the computer equipment is exposed to an increased risk of disruption should a failure occur on the

remaining path. This maintenance condition does not defeat the Tier rating achieved in normal operations.

- Operation of the fire alarm, fire suppression, or the emergency power off (EPO) feature may cause a data center disruption.

Tiers and Engine-Generator Plants

- Tier III and IV engine-generator plants are considered the primary power source for the data center. The local power utility is an economic alternative. Disruptions to the utility power are not considered a failure, but rather an expected operational condition for which the site must be prepared.
- A Tier III or IV engine-generator system, along with its power paths and other supporting elements, shall meet the Concurrently Maintainable and/or Fault Tolerant performance confirmation tests when they are carrying the site on engine-generator power.
- Engine generators for Tier III and IV sites shall not have a limitation on consecutive hours of operation when loaded to "N" demand. Engine generators that have a limit on consecutive hours of operation at "N" demand are appropriate for Tier I or II.
- Engine-generator systems often have an annual limit on operating hours driven by emissions. These environmental limits do not impact the consecutive hours of operation constraint established in this section.

Determining a Site's Tier Rating for Design Topology

Determining a site's actual Tier rating for design topology is not a complicated process, although it is one that is rarely done correctly. *Table 1* includes the requirements for the Tier Performance Standards and how they apply by Tier.

Simply put, **the Tier topology rating for an entire site is constrained by the rating of the weakest subsystem that will impact site operation.** For example, a site with a robust Tier IV UPS configuration combined with a Tier II chilled water system will yield a Tier II site rating.

This very stringent definition is driven by the expectations of senior executives who have signed off on multi-million dollar investments, as well as to factually report actual site capabilities. Any exceptions and exclusions footnoted in the approval documents will soon be lost and forgotten. If a site has been advertised within an organization as being Fault Tolerant (Tier IV), it will be intolerable (and perhaps career ending) to have to plan a site shutdown at any time in the future—regardless of any "fine print" exclusions that diligently identified the risk. For this reason, **there are no partial or fractional Tier ratings.** A site's Tier rating is not the average of the ratings for the critical site infrastructure subsystems. The site's Tier rating is the lowest of the individual subsystem ratings.

Similarly, the Tier rating cannot be claimed by using calculated MTBF component statistical reliability to generate a predictive availability and then using that number to match the actual

measured availability results shown later in *Table 2*. Statistically valid component values are not available, partly because product life cycles are getting shorter and no independent, industry-wide database exists to collect failure data.

Even if accurate failure values existed, this approach often fails to include the information recovery time for the end users, which in many cases is measured in hours. It also fails to consider the 70 percent of failures caused by human interaction. If these real factors are not considered in component failure calculations, this is truly a case of “garbage in, garbage out” and any conclusions are going to be off by a factor of at least five to eight. A calculated reliability of 0.9999, which considers only component and system equipment failures, ignores human interaction as a source of failure, and does not consider that information recovery time does NOT define a site as being Tier IV! The only way to determine Tier level is to objectively determine a site topology’s ability to respond to planned and unplanned events.

Please refer to the section “Uptime Institute Tier Certification” below for further information.

Tier Commentary

The Institute’s STANDARDS Are Based on a Specific Performance Outcome

The requirements used in the Institute’s Tier Performance Standard are necessarily and intentionally very broad to allow innovation and client manufacturer/equipment preferences in achieving the desired level of site infrastructure performance or uptime. The individual Tiers represent categories of site infrastructure topology that address increasingly sophisticated operating concepts, leading to increased site infrastructure availability. The performance outcomes that define the four Tiers of site infrastructure are very straightforward. Recent initiatives by several groups to replace the Institute’s Tier concepts with component counts and checklists have lost the focus that what ultimately counts is uptime performance. Most designs that pass a checklist approach will fail a performance requirements approach. What this means is that, in addition to the rigorous application of engineering principles, there is still considerable art to the design of uptime and how subsystems are integrated to allow for multiple operating modes.

Tier Functionality Progression

Tier I solutions acknowledge the owner’s desire for dedicated site infrastructure to support IT systems. Tier I infrastructure provides an improved environment over that of an ordinary office setting and includes: a dedicated space for IT systems; a UPS to filter power spikes, sags, and momentary outages; dedicated cooling equipment not shut down at the end of normal office hours; and an engine generator to protect IT functions from extended power outages.

Tier II solutions include redundant critical power and cooling capacity components to provide an increased margin of safety against IT process disruptions due to site infrastructure equipment failures. The redundant components are typically extra UPS

modules, chillers, heat rejection equipment, pumps, cooling units, and engine generators. A malfunction or normal maintenance will result in loss of a capacity component.

Owners who select Tier I and Tier II solutions to support current IT technology are typically seeking a solution to short-term requirements. Both Tier I and Tier II are usually tactical solutions, i.e., driven by first-cost and time-to-market more than life-cycle cost and uptime (or availability) requirements. Rigorous uptime requirements and long-term viability usually lead to the strategic solutions found more often in Tier III and Tier IV site infrastructure. Tier III and Tier IV site infrastructure solutions have an effective life beyond the current IT requirement. Strategic site infrastructure solutions enable the owner to make strategic business decisions concerning growth and technology, unconstrained by current site infrastructure topology.

Tier III site infrastructure adds the concept of Concurrent Maintenance beyond what is available in Tier I and Tier II solutions. Concurrent Maintenance means that **each and every** capacity or distribution component necessary to support the IT processing environment can be maintained on a planned basis without impact to the IT environment. The effect on the site infrastructure topology is that a redundant delivery path for power and cooling is added to the redundant critical components of Tier II. Maintenance allows the equipment and distribution paths to be returned to like new condition on a frequent and regular basis. Thus, the system will reliably and predictably perform as originally intended. Moreover, the ability to concurrently allow site infrastructure maintenance and IT operation requires that **each and every** system or component that supports IT operations must be able to be taken offline for scheduled maintenance without impact to the IT environment. This concept extends to important subsystems such as control systems for the mechanical plant, start systems for engine generators, EPO controls, power sources for cooling equipment and pumps, isolation valves, and others.

Tier IV site infrastructure builds on Tier III, adding the concept of Fault Tolerance to the site infrastructure topology. Similar to the application of Concurrent Maintenance concepts, Fault Tolerance extends to **each and every** system or component that supports IT operations. Tier IV considers that any one of these systems or components may fail or experience an unscheduled outage at any time. The Tier IV definition of Fault Tolerance is based on a single component or path failure. However, the site must be designed and operated to tolerate the cumulative impact of every site infrastructure component, system, and distribution path disrupted by the failure. For example, the failure of a single switchboard will affect every subpanel and equipment component deriving power from the switchboard. A Tier IV facility will tolerate these cumulative impacts without affecting the operation of the computer room.

The progressive nature of functionality from Tier I to Tier II and Tier III to Tier IV is demonstrated in the schematic illustrations found at the end of this white paper. The examples show the addition of components, systems, and distribution paths, as described

above. **Although these illustrations are not recommended design solutions for any particular set of requirements**, the four electrical topologies are illustrative of the Tier Classification concepts. Substantial levels of detail (e.g., segmented electrical buses, isolation breakers, etc.) were omitted from these sketches for purposes of clarity. Mechanical system functionality progresses similarly through increasing Tier levels.

Consistent, across-the-board application of Tier concepts for electrical, mechanical, automation, and other subsystems is absolutely required for any site to satisfy the Tier standards. However, site infrastructure has been occasionally described by others in the industry in terms of fractional Tiers (e.g., Tier 2.5), or incremental Tiers (Tier III +, Enhanced Tier III, or Tier IV lite). Fractional or incremental descriptions for site infrastructure are not appropriate and are misleading.

- A site that has an extra (redundant) UPS module but needs all the installed cooling units running to keep the computer room temperature within limits does not meet the redundancy requirements for Tier II.
- A switchboard that cannot be shut down without affecting more than the redundant number of secondary chilled water pumps (reducing the available capacity to less than N) is not Concurrently Maintainable and will not be Certified as Tier III.

The four Tier Standard Classifications address topology, or configuration, of site infrastructure, rather than a prescriptive list of components to achieve a desired operational outcome. For example, the same number of chillers and UPS modules can be arranged on single power and cooling distribution paths resulting in a Tier II solution (Redundant Components) or on two distribution paths that may result in a Tier III solution (Concurrently Maintainable). Compare the Tier II and Tier III schematic illustrations at the end of this paper. Both topologies contain the same N+1 capacity redundancy for engine generators and UPS modules, but the alternate distribution paths characterize Tier III.

Align Business Requirements with Tiers

Each industry has a unique uptime need driving the site-infrastructure Tier level required. After careful alignment of IT availability objectives with site infrastructure performance-expectations, an informed company may select a site representing any of the Tier Classifications.

Tier I is appropriate for firms such as:

- Small businesses where information technology primarily enhances internal business process.
- Companies whose principal use of a web-presence is as a passive marketing tool.
- Internet-based start-up companies without financially enforceable customer quality-of-service commitments.

These companies typically do not have an established revenue stream or an identifiable financial impact of disruption due to data

center failure. Sometimes companies with an established revenue stream will select Tier I topology because their applications have a low availability requirement, such as only during a 5.5-day business week. Other companies may select Tier I topology if they plan to abandon the site when the business requirements exceed the Tier I functionality.

Tier II is appropriate for firms such as:

- Call centers where multiple sites are available.
- Internet-based companies without serious financial penalties for quality-of-service commitments.
- Small businesses whose information technology requirements are mostly limited to traditional normal business hours, allowing system shutdown during off hours.
- Scientific research, e.g., chip design, oil exploration, seismic processing, or long-term weather modeling, which typically does not have online or real-time service delivery obligations.

These companies typically do not depend on real-time delivery of products or services for a significant part of their revenue stream or are contractually protected from damages due to lack of system availability. Occasionally, companies will select Tier II infrastructure if they have become burdened with impacts due to nuisance equipment outages associated with Tier I sites. A large number of institutional and educational organizations select Tier II infrastructure because there is not a significant impact of disruption due to a data center failure. Some companies have successfully used Tier II infrastructure to provide off-site electronic vaulting for offline data.

Typical applications for Tier III facilities are:

- Companies that support internal and external clients 24x7, such as service centers and help desks, but can accept short periods with limited service due to a site failure.
- Businesses whose IT resources support automated business processes, so the impact on clients of system shutdowns is manageable or acceptable.
- Companies spanning multiple time zones with clients and employees spanning regional areas.
- Internet-based companies or co-location providers that have quality-of-service commitments with serious financial ramifications.

Companies selecting Tier III infrastructure usually have high-availability requirements for ongoing business or have identified a significant cost of disruption due to a planned data center shutdown. These companies are willing to accept the impact-of-disruption risk of an unplanned event. However, Tier III is appropriate for companies who expect the functionality requirements to increase over time and do not want to abandon the data center. Sometimes, these companies design a Tier III site to be transparently upgraded to Tier IV. This requires Tier IV functionality to be provided for in the original planning concept.

Tier IV is justified most often for:

- Companies with an international market presence delivering “24 by forever” services in a highly competitive client-facing market space or where processes are continuous (international in- and outbound wire transfers, etc.).
- Businesses based on E-commerce, market transactions, or financial settlement processes.
- Large, global companies where client access to applications and employee exploitation of information technology is a competitive advantage.
- Internet-based companies or co-location providers that have quality-of-service commitments with serious financial ramifications.

Companies who have extremely high availability requirements for ongoing business or for whom there is a profound and clearly identified cost of disruption due to any data center shutdown select Tier IV site infrastructure. These companies will know the cost of a disruption, usually in terms of both actual dollar costs and impact to market share. The cost of disruption makes the case for investment in high-availability infrastructure a clear business advantage.

Selecting the appropriate topology solution based on the IT availability requirements to sustain well-defined business processes without substantial financial consequences for downtime provides the best foundation for investment in data center facilities. It is preferable for the owner’s focus during the data center design and delivery process to be on the consistent application of the Tier Performance Standard rather than on the details that make up the data center site infrastructure.

Including a criteria or an attribute of a higher Tier Classification in the design does not increase the overall Tier Classification. However, deviation from a Tier objective in any subsystem will prevent a site from being Certified at that Tier. For example, a UPS system patterned after a Tier IV system within a site having a Tier II power distribution backbone will yield a Tier II site. The most significant deviations from the Tier Standard found in most sites can be summarized as inconsistent solutions.

Frequently, a site will have a robust, Fault Tolerant electrical system patterned after a Tier IV solution but will utilize a Tier II mechanical system that cannot be maintained without interrupting computer room operations. **This results in an overall Tier II site rating.** Most often, the mechanical system fails Concurrent Maintenance criteria because of inadequate coordination between the number and location of isolation valves in the chilled water distribution path. Another common oversight is branch circuiting of mechanical components, which results in having to shut down the entire mechanical system to perform electrical maintenance. If more than the redundant number of chillers, towers, or pumps is de-energized for electrical maintenance, computer-room cooling is impacted.

Electrical systems often fail to achieve Tier III or Tier IV criteria due to design choices made in the UPS and the critical power distribution path. UPS configurations that utilize common input and output switchgear are almost always un-maintainable without

computer room outages and will fail the Tier III requirements even after spending many hundreds of thousands of dollars. Topologies that include static transfer switches in the critical power path for single-corded IT devices will likely fail both the Fault Tolerance criteria and the Concurrent Maintenance criteria.

Consistent application of standards is necessary to have an integrated solution for a specific data center. It is clear that the IT organization invests heavily in the features offered by newer computer equipment technology. Often, as the electrical and mechanical infrastructures are defined and the facility operations are established, there is a growing degree of inconsistency in the solutions incorporated in a site. As shown in *Table 3*, each segment must be integrated to deliver the overall data center solution. An investment in one segment must be met with a similar investment in each of the other segments if any of the elements in the combined solution are to have the desired effect on IT availability. A well-executed data center master plan or strategy should consistently resolve the entire spectrum of IT and facility requirements.

IT Availability Success Is Dependent upon Successful, Fully Integrated Operation of All Site Infrastructure Systems

The Tier Classifications were created to consistently describe the site-level infrastructure required to sustain data center operations, not the characteristics of individual systems or subsystems. Data centers are dependent upon the successful and integrated operation of at least 16 separate site infrastructure subsystems. Every subsystem and system must be consistently deployed with the same site uptime objective to satisfy the distinctive Tier requirements. The most critical decision-making perspective owners and designers must consider when making inevitable tradeoffs is what effect does the decision have on the life cycle integrated operation of the IT environment in the computer room.

The Institute defines site availability from the perspective of an IT user. Any site incident or event that affects information availability as experienced by end users detracts from site infrastructure availability. It is understood that there are IT or network device failures that will also detract from end-user availability. The site downtime clock starts running from the moment IT operations are first affected until the time they are fully restored. Thus, site downtime is not the 15 seconds of a utility power failure but the total time users are down until IT availability is restored. For Tier I and Tier II topologies, downtime for site infrastructure maintenance, which includes the time to bring IT systems down, perform the site maintenance, and restore IT availability, typically has a bigger availability impact than a UPS system failure. Based on operating experience of monitored sites, the typical maintenance outage at Tier I and Tier II sites is 12 hours. The time for IT to recover from a typical outage such as momentary power loss is 4 hours at sites of any Tier.

The Institute has measured actual performance of 16 data centers having site infrastructure topologies meeting the four Tier definitions. These figures have been converted to end-user availability values representative of each Tier Classification. These

empirically determined values include equipment failures as well as Operational Sustainability and human factors over a period of up to 10 years with uptime measured from the perspective of the IT client's end users. This "real world" site availability is strikingly different from the probability of system failure that is often calculated (with human interaction and end-user recovery omitted) using values from the Institute of Electrical and Electronics Engineers (IEEE) Gold Book for recommended practices for reliable power systems or guidelines from the IEEE Orange Book for emergency and standby power. A representative site infrastructure availability of 99.95 percent (about 4.4 hours of downtime per year) is not equivalent to a statistical reliability of 0.9995 (1 in 2,000 chance of a failure). Similarly, as outlined earlier, a calculated statistical reliability of 0.9995 does not indicate a site is "better than Tier III."

Tier I sites typically experience two separate 12-hour, site-wide shutdowns per year for maintenance or repair work. In addition, across multiple sites and over a number of years, Tier I sites experience 1.2 equipment or distribution failures on average each year. The annual impact of maintenance and unplanned outages is 28.8 hours per year, or 99.67 percent availability.

Operations experience shows that, on average, Tier II sites schedule three maintenance windows over a 2-year period and have one unplanned outage each year. The redundant components of Tier II topology provide some maintenance opportunity leading to just one site-wide shutdown each year and reduce the number of equipment failures that affect the IT operations environment. The annual impact of maintenance and unplanned outages is 22 hours per year, or 99.75 percent availability.

Tier III topology is Concurrently Maintainable, so annual maintenance shutdowns are not required, which allows an aggressive maintenance program improving overall equipment performance. Experience in actual data centers shows that operating better maintained systems reduces unplanned failures to a 4-hour event every 2.5 years, or 1.6 hours on an annual basis. Tier III sites demonstrate 99.98 percent availability.

Tier IV provides robust, Fault Tolerant site infrastructure, so that facility events affecting the computer room are empirically reduced to one 4-hour event in a 5-year operating period, or 0.8 hours on an annual basis. Individual equipment failures or distribution path interruptions may still occur, but the effects of the events are stopped short of the IT operations environment. Tier IV sites consistently demonstrate 99.99 percent availability.

The representative availability percentages reported are a characteristic of the operating experience of multiple sites within each Tier Classification. A site with a measured infrastructure availability of 99.90 percent—midway between Tier II (99.75 percent) and Tier III (99.98 percent)—has an operating experience consistent with sites having Tier II topology but does not achieve the availability of Tier III sites. Availability calculations do not determine the Tier Classification. Even more importantly, infrastructure with a statistical probability of failure of 0.9990

cannot be represented as a "Tier 2.5" site, since the impact of the failure on overall availability is not represented by the likelihood of a system failure and there are no fractional Tier ratings.

Independent of site infrastructure experience, IT organizations often describe data center availability objectives as "Five Nines," or 99.999 percent of uptime. This is a very aggressive goal, especially if compared to the observed consequences of a single site outage. While a site outage is assumed to be promptly restored (which requires "24 by forever" staffing), it can still require up to 4 hours for IT to recover information availability and restore end-user functionality, even if the likelihood of database corruption or a server power supply failure are set aside. In reality, facility failures often reveal previously unknown IT architecture, hardware, or software issues.

If a momentary site infrastructure power outage results in a 4-hour end-user disruption, how relevant is an objective of 99.999 percent availability? Based on a single site outage of 4 hours, it will take 45.6 years of 100 percent uptime to restore cumulative site availability back to the 99.999 percent objective (4 hours x 60 minutes an hour ÷ 5.26 minutes per year = 45.6 years).

Even a Fault Tolerant and Concurrently Maintainable Tier IV site will not satisfy an IT requirement of Five Nines (99.999 percent) uptime. The very best a Tier IV site can hope for is 100 percent uptime for multiple years, but at some point, a site failure is almost certain. *Table 2* uses 99.995 percent for representative Tier IV site availability, but this assumes that a site outage occurs not more than once every 5 years. Even with a properly designed Tier IV configuration, a false fire alarm or the unintended operation of the EPO feature almost certainly will cause a site failure. Unless human activity issues are continually and rigorously addressed, at least one failure is likely over a 5-year period.

Clients with very high availability requirements for their end users often consider multiple sites working in concert. This requires a more complex IT architecture and network configuration but will provide the ability to have the unfortunate EPO or fire event at one site be transparent to the end users.

Operational Sustainability: More Than Tiers

The Tier Performance Standard provides objective criteria to consistently evaluate the implementation of the selected operational concepts in a design or existing site infrastructure. The standard does not direct the specific design solution or technology the owner or design team must use to reach the site performance objective. Owners are free to choose any number of UPS configurations, products, or manufacturers—as long as the result can meet the target Tier confirmation tests. Moreover, the use of static or rotary UPS systems, fuel-cell technologies, direct-expansion cooling, or air- or water-cooled chillers are left to the owner. The Tier Standards have attained wide acceptance because they allow the owner to include such concerns as first cost, operations complexity, proven technologies, and product availability as appropriate, while

still focusing on the desired operational outcome of the completed facility.

In addition to availability, other owner requirements must be addressed in infrastructure design. Protection of data or physical assets is independent of the site infrastructure Tier Classification. Increasing IT equipment power densities requires additional design considerations beyond redundancy in the power and cooling systems. Project requirements such as video surveillance and gaseous fire suppression are frequently necessary to meet an owner’s regulatory compliance or insurance requirements and are completely separate from IT availability objectives. The key understanding required for a successful data center operation is to differentiate between Tier Performance Standard criteria, owner risk and cost tolerance, and industry best practices.

Considerations of cost, risk tolerance, and best practices clearly point to a wider number of site infrastructure characteristics than shown in *Table 2*. The Institute refers to these other factors as Operational Sustainability. Experience by the authors with the Tier Standard since its inception indicates that over time, Operational Sustainability characteristics become a much more important contributor to ultimate investment value and long-term availability.

Investments in improved Operational Sustainability characteristics often lead to increased life-cycle availability and account for much of the cost variance within individual Tier solutions. Improved Operational Sustainability characteristics decrease the cost or risk of completing maintenance or speed the recovery from site infrastructure incidents or failures. Less costly and less risky maintenance means the work is more likely to be completed, keeping the equipment in better condition and calibration. More operations-centric designs make operations easier, so fewer mistakes are made.

Illustrative Examples:

The following exemplify site infrastructure characteristics that impact Operational Sustainability, while not affecting the overall Tier Classification of the solution.

- A topology that can switch the power source for all mechanical components so they continue to run when any electrical panel is shut down eliminates an operations constraint to maintenance. Procedures that require critical cooling equipment to be shut down during recurring electrical system maintenance may not be allowed if another chiller is out of service for repairs. Missed maintenance leads to decreased reliability.
- A design that mounts critical components in difficult-to-reach areas or limits access space in the central plant may increase the time required to maintain important systems. The increased time window may eliminate the ability to schedule the maintenance activity or extend restoration of the redundancy after a component failure.

- Installing engine generators and switchgear inside the facility (with adequate access space) eliminates the effects of weather and time-of-day on access, safe maintenance, and repair activities.
- In order to improve stability, the combined load on a critical system is often limited to 90 percent nameplate or nominal capacity, excluding redundant components, over a sustained period of time. This is referred to as a redline rating.
- Compartmentalization, a Tier IV requirement, provides benefits for Tier III sites. The effects of evacuation requirements for areas affected by refrigerant leaks can be limited to the number of redundant chillers by careful Compartmentalization. Chillers that are necessary to keep the computer room cool can continue to operate while those in a separate compartment are shut down to purge the refrigerant.
- Compartmentalization of the primary and maintenance electrical distribution paths also provides a major advantage to a site. If an arc flash or electrical fire (an unplanned event) occurs in a Tier III site, the site could be disrupted. However, if the maintenance path is physically separated from the normal path, Compartmentalization will permit the site to rapidly recover on a power path through a completely different space from where the fire occurred.

Uptime Institute Tier Certification

The Uptime Institute has retained the exclusive legal right to review, assess, and Certify data centers and other critical facilities to the Institute’s Tier Classification System. Furthermore, the Institute retains the exclusive right to utilize its comprehensive database of emerging industry problems and best design practices to inform Operational Sustainability enhancement recommendations.

The Uptime Institute’s Tier Certification involves additional criteria beyond the high-level information provided herein. The three steps in this process are:

- Design Certification: The deliverables include a review of the project’s 100 percent design document package, a report of Tier deficiency findings as well as enhancement opportunities for Operational Sustainability, an electronic badge, and a formal letter stating the Tier Certification level represented in the design documents. This document-review step allows project owners to proceed to construction with confidence that a critical aspect of the site infrastructure’s configuration was not missed. Institute staff is also available to review the design at major milestones to ensure that neither budget nor schedule is exhausted on non-compliant solutions.
- Constructed Facility Certification: The deliverables include detailed findings from a physical, on-site review of the constructed environment, a review of commissioning documentation,² observation of demonstrations of the site’s Concurrent Maintenance and/or Fault Tolerance functionality, a wall plaque suitable for hanging, and a

formal Institute letter stating the Tier level attained.

- **Operational Sustainability Rating:** This rating addresses the site availability factors beyond Tier of design topology and includes an assignation of the re-Certification period (1, 3, or 5 years, exclusive of major capacity or functionality upgrades or modifications). Deliverables include findings and recommendations on the staffing and management program (metrics, training, processes, and procedures) used to operate the site; an electronic badge; a plaque suitable for wall mounting; and a formal Institute letter stating the level of Tier and Operational Sustainability achieved.

Only sites listed at uptimeinstitute.org/certifiedsites are Uptime Institute Design, Facility, or Operational Sustainability Certified. Any site not listed by the Institute is self-certified.

The Institute has reviewed the claims of dozens of self-certified sites. To date, four sites have attained, on initial review, the Tier level promised by the design engineer or desired by the client. The others required expensive re-work. For these projects, the gap between client expectations and the actual Tier level achieved was consistently at least one level, and often two levels. The gap is not always attributable to a funding shortfall as the owner often spent more than enough to have achieved the desired outcome. In addition to disappointment and career risk, the result is wasted investment often measured in millions of dollars. In the majority of cases, the Institute's Certification process will achieve the desired result AND save money because the Institute's specialized professionals can often suggest optimal design solutions that are elegant in their simplicity and operability.

Conclusion

Data center owners have the responsibility for determining what Tier of functionality is appropriate or required for their sites. As such, it is a business decision to determine the Tier necessary to support site availability objectives. Part of this decision is to balance the IT operational practices with the facility practices that support the IT world. Once selected, however, the desired Tier must be uniformly implemented across all systems.

Definition of Terms

- **Computer Hardware Equipment**
This terminology encompasses all information technology equipment required at a data center to perform the information processing work. It includes servers, storage, network, communication, and all other information technology components.
- **Concurrent Maintenance**
This is originally an IT term, which entails a site functionality such that any work can be performed on a planned basis without impacting the end user. In the site infrastructure

world, this means that **each and every** capacity component or distribution element can be repaired, replaced, serviced, tested, etc., on a planned basis without affecting the computer equipment or processing.

- **Fault Tolerant**

This means that a system is capable of sustaining an unplanned event without disrupting the computer's end user. The concept of Fault Tolerance originated in the IT environment. In the site infrastructure world, it means that the computer equipment will not be impacted by a facility failure (fire alarms and EPO exposures are notable exceptions). This benefit requires multiple, independent sources and multiple, independent, diversely-routed distribution paths, so a failure on one source or path does not impact the other. This also requires use of computer equipment that meets the Institute's *Fault Tolerant Compliant Power Specification, Version 2.0*. Computer equipment that does not meet that specification will require additional components, such as a point-of-use switch, to stay online during site infrastructure maintenance activity or to provide Fault Tolerance at the IT device.

- **Point-of-Use Switch**

A fast-acting, rack-mounted switch installed to allow single-corded loads to exploit the dual-power-path benefits of Tier III and IV topology.

- **Redundant Capacity Components**

The number of active capacity components in a system beyond the minimum number of units required to support the IT load is referred to as redundant. If one unit of capacity is required to support the computer equipment, more than one unit of capacity is installed. Terms such as N+1 (need plus one) or N+2 are commonly applied to capacity component counts but not distribution paths. Ideally, with redundant capacity components, failure of a capacity unit is transparent and system performance is unaffected. In the real world, capacity unit failures can cause system failures.

- **Site Infrastructure**

A typical data center site is composed of at least 16 major mechanical, electrical, fire protection, security, and other subsystems. Each has additional site subsystems and components.

- **Source**

Computer equipment consumes electricity and gives off heat in the process. Continuous, uninterruptible operation requires multiple sources of power and cooling independent of each other. These power and cooling sources are usually composed of many different components to form a reliable system.

²Full commissioning of the site infrastructure, including Level 5 Integrated Systems Operations Testing, is required for Tier Certification.

- **Useable Capacity**

This is the maximum amount of actual load that can be supported by the N level of capacity. Typically, the maximum amount of useable load is less than the nameplate capacity to allow for component aging, installation errors, and to provide for contingencies.

Referenced White Papers

The following white papers are referenced herein and can be found at uptimeinstitute.org/whitepapers:

- *Operational Sustainability and Its Impact on Data Center Uptime Performance, Investment Value, Energy Efficiency, and Resiliency*
- *Cost Model: Dollars per kW plus Dollars per ft² of Computer Floor*
- *Fault Tolerant Compliant Power Specification, Version 2.0*
- *Continuous Cooling Is Required for Continuous Availability*
- *How to Meet “24 by Forever” Cooling Demands of Your Data Center*
- *Site Uptime Procedures for Safely Performing Work in a Live Data Center*

About the Authors

This white paper is the collective result of the experience of the authors who have personally designed and operated major data centers. This paper is further grounded in a database of over 4,000 AIRs and 400 data center failures. This work has been accomplished independent of any Engineer-of-Record or manufacturer affiliation. Since 1980, this expanding team has guided, justified, and overseen billions of dollars of data center investments for major, global companies that are absolutely dependent upon truly continuous availability for business survival.

Vince Renaud, PE, is a Distinguished Fellow and Certification Authority for the Uptime Institute and a Principal of ComputerSite Engineering. Mr. Renaud has provided leadership and strategic direction to maintain the highest level of infrastructure availability. In varying roles from data center owner and operator to consultant, Mr. Renaud has provided planning, design, construction, operation, and maintenance of mission critical facilities for the Department of Defense and Fortune 100 companies on a world-wide basis.

John H. Seader, PE, is a Distinguished Fellow and Certification Authority for the Uptime Institute and a Principal of ComputerSite Engineering. Mr. Seader’s career in critical facilities spans 14 years and includes responsibilities ranging from planning, engineering, design, and construction to start-up and operation for clients such as the Department of Defense, Sabre, and Williams Communication. Prior to joining ComputerSite Engineering, Mr. Seader was a Senior Technology Manager for Deloitte Consulting Outsourcing, LLC.

W. Pitt Turner IV, PE, is a Distinguished Fellow and Senior Certification Authority for the Uptime Institute and a faculty member for the Institute’s Site Uptime Network®. As a Principal of ComputerSite Engineering, Mr. Turner has personally guided billions of dollars in client site infrastructure investments. Prior to joining ComputerSite Engineering in 1993, Mr. Turner was a Senior Project Manager for Pacific Bell’s Fairfield Data Center; he was responsible for concept development, design, construction, and start-up for a 200,000-ft² facility. His work included the benchmarking of other data centers to help establish business-process improvements.

Additional editorial guidance provided by Terry Altom, PE, and Ken Brill.

Mr. Altom is a Consultant for ComputerSite Engineering who specializes in site infrastructure audits and custom strategic-level consulting engagements.

Mr. Brill is the founder of the Uptime Institute, its Executive Director, and a Principal of ComputerSite Engineering, which he also founded.

Previously, ComputerSite Engineering, Inc., a sister company of the Institute, was exclusively authorized to perform site and resiliency certifications. This authorization has been withdrawn, but projects currently identified or in process will be allowed to finish. Whether under ComputerSite Engineering or the Uptime Institute, the same team of Certification Authorities (Vince Renaud, Hank Seder, and Pitt Turner) will continue to oversee Institute Certifications of Tier and Operational Sustainability.

About the Uptime Institute

The Uptime Institute, Inc. is a pioneer in creating and operating knowledge communities for improving uptime effectiveness in data center Facilities and Information Technology organizations. The 100 members of the Institute’s Site Uptime Network® are committed to achieving the highest levels of availability and many are Fortune 500 companies. Members learn interactively from each other as well as from Institute sponsored 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. For the industry as a whole, the Institute publishes white papers, offers a Site Uptime Seminar Series, Site Uptime Symposium, and Data Center Design Charrette Series on critical uptime-related topics. The Institute also conducts sponsored research and product certifications for industry manufacturers. For data center owners, the Institute Certifies data center Tier level and Operational Sustainability.

Tables and Illustrations

Table 1: Tier Requirements Summary

	Tier I	Tier II	Tier III	Tier IV
Active Capacity Components to Support IT Load	N	N+1	N+1	N after any failure
Distribution Paths	1	1	1 active and 1 alternate	2 simultaneously active
Concurrently Maintainable	No	No	Yes	Yes
Fault Tolerance (single event)	No	No	No	Yes
Compartmentalization	No	No	No	Yes
Continuous Cooling*	load density dependent	load density dependent	load density dependent	Yes (Class A)

*For additional information on Continuous Cooling, refer to the Uptime Institute white paper *Continuous Cooling Is Required for Continuous Availability*.

Table 2: Common Attributes Found in Data Centers that are Unrelated to Tier Requirements

Tier I sites have their roots in the mainframe environments of the 1970s. Tier IV became possible with the advent of dual-powered computers in the mid 1990s. Tier II facilities were widespread in the 1980s; Tier III started in the mid 1980s and is the most common site infrastructure currently being implemented, although many are designed for future transparent upgrades to Tier IV. A responsible approach to site infrastructure investment is to understand clearly the availability objectives necessary to support the owner's current and future business requirements and then to consistently design, build, and operate the site to conform to those needs.

The following table depicts various attributes commonly associated with data centers, but these attributes are not requirements of the Tier definitions. For example, the presence of a raised floor or the height of the floor are not criteria at any Tier. (The recommended height of raised floors, when used, is most directly correlated to power density and operability.)

	Tier I	Tier II	Tier III	Tier IV
Building Type	Tenant	Tenant	Stand-alone	Stand-alone
Staffing shifts Staff/shift	None None	1 Shift 1/Shift	1+Shifts 1-2/Shift	"24 by Forever" 2+/Shift
Useable for Critical Load	100% N	100% N	90% N	90% N
Initial Build-out kW per Cabinet (typical)	<1kW	1-2 kW	1-2 kW	1-3 kW
Ultimate kW per Cabinet (typical)	<1 kW	1-2 kW	>3 kW ^{†,*}	>4 kW ^{†,*}
Support Space to Raised-Floor Ratio	20%	30%	80-90+%	100+%
Raised-Floor Height (typical)	12 inches	18 inches	30-36 inches	30-42 inches
Floor Loading lbs/ft (typical)	85	100	150	150+
Utility Voltage (typical)	208, 480	208, 480	12-15 kV	12-15 kV
Single Points-of-Failure	Many + Human Error	Many + Human Error	Some + Human Error	Fire, EPO + Some Human Error
Representative Planned Maintenance Shut Downs	2 Annual Events at 12 Hours Each	3 Events Over 2 Years at 12 Hours Each	None Required	None Required
Representative Site Failures	6 Failures Over 5 Years	1 Failure Every Year	1 Failure Every 2.5 Years	1 Failure Every 5 Years
Annual Site-Caused, End-User Downtime (based on field data)	28.8 hours	22.0 hours	1.6 hours	0.8 hours
Resulting End-User Availability Based on Site-Caused Downtime	99.67%	99.75%	99.98%	99.99%
Typical Months to Plan and Construct	3	3-6	15-20	15-30
First Deployed	1965	1970	1985	1995

^{*} 3.5 kW per cabinet over large areas is acceptable for traditional air-cooling designs.

[†] Higher kW/cabinet densities require a greater ratio of support space to computer floor (at least 1:1 at 3 kW/cabinet, 2:1 at 6 kW/cabinet, 3:1 at 9 kW/cabinet, etc.) Generally, deeper raised floors are required for higher densities.

^{*} Most sites have a difficult time maintaining stable and predictable cooling for racks in the 1-2 kW range. Major process improvements are required before entertaining rack densities above 2 kW. See Institute white paper *How to Meet "24 by Forever" Cooling Demands of your Data Center*.

Table 3: Comparing IT Solutions for Reliability, Availability, and Serviceability to 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 Generator Dual Power S + S
Mechanical Infrastructure	Redundant Components Fan 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

Illustration 1: Illustrative Electrical System Topology - Tier I

Note: This diagram illustrates basic Tier I electrical distribution concepts. This diagram should not be interpreted to represent a standard nor compliant electrical system topology or a solution fulfilling any particular set of requirements.

Tier Certification requires consistent application of Tier concepts to all relevant critical subsystems that comprise data center site infrastructure.

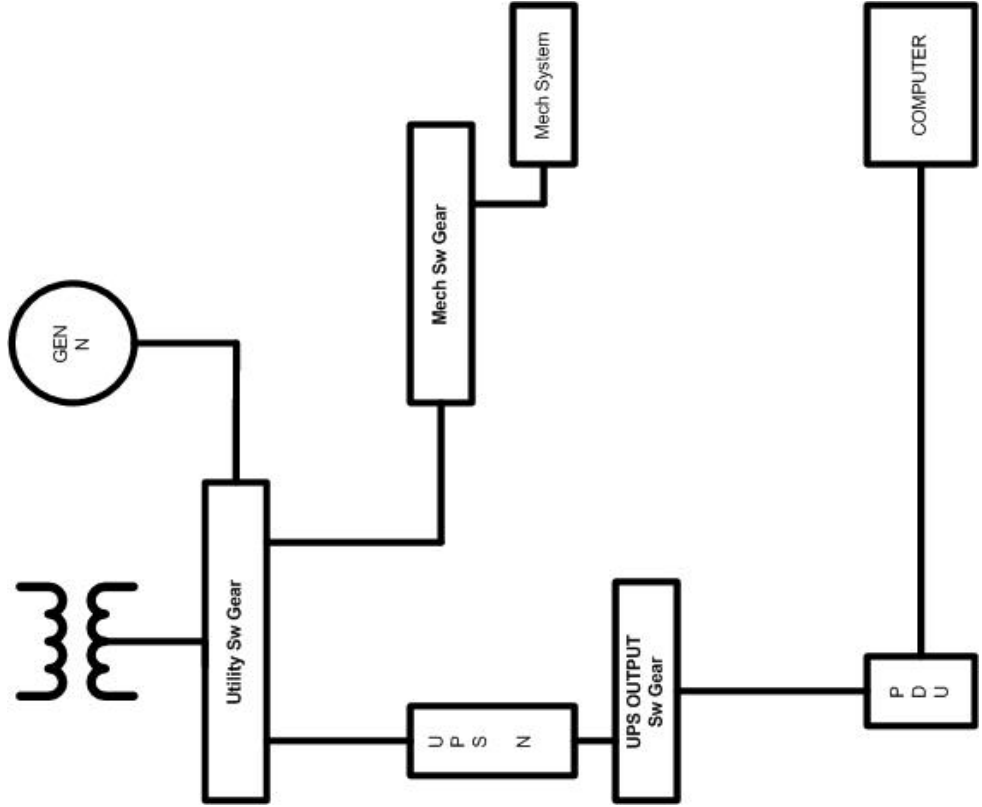


Illustration 2: Illustrative Electrical System Topology - Tier II

Note: This diagram illustrates a basic Tier II electrical distribution concept. Note the additional (redundant) equipment required to move from Tier I to Tier II. This diagram should not be interpreted to represent a standard nor compliant electrical system topology or a solution fulfilling any particular set of requirements.

Tier Certification requires consistent application of Tier concepts to all relevant critical subsystems that comprise data center site infrastructure.




Illustration 3: Illustrative Electrical System Topology - Tier III

Note: This diagram illustrates a Tier III electrical distribution concept. Note the additional (redundant) equipment required to move from Tier II to Tier III. This diagram should not be interpreted to represent a standard nor compliant electrical system topology or a solution fulfilling any particular set of requirements.

Tier Certification requires consistent application of Tier concepts to all relevant critical subsystems that comprise data center infrastructure.

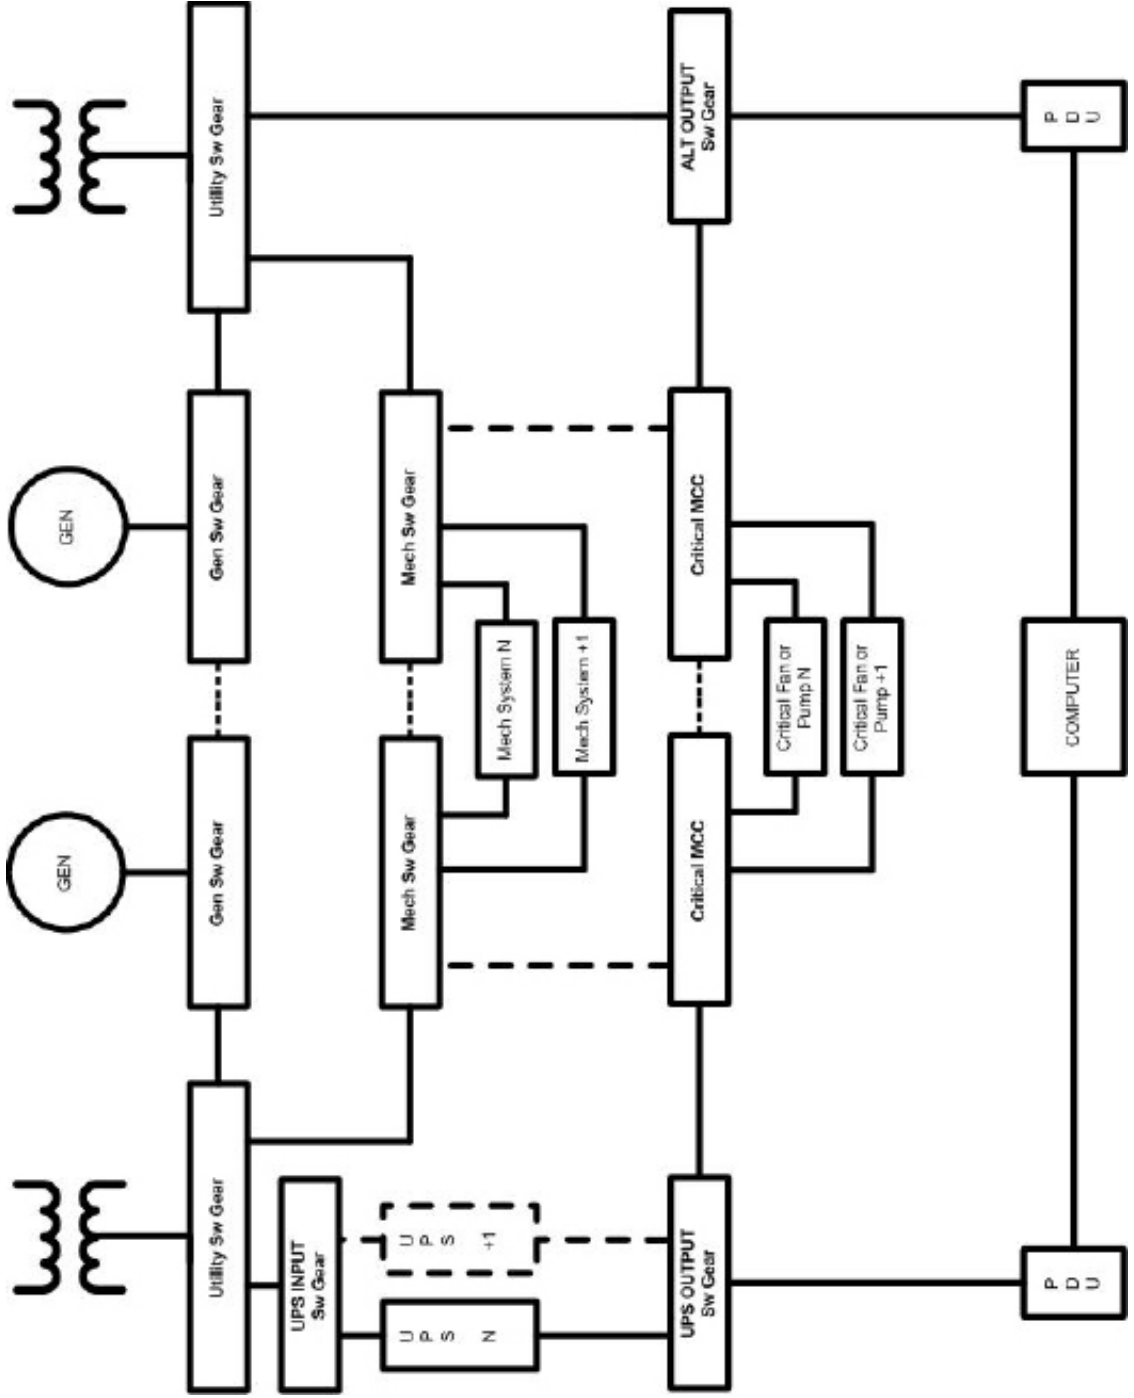
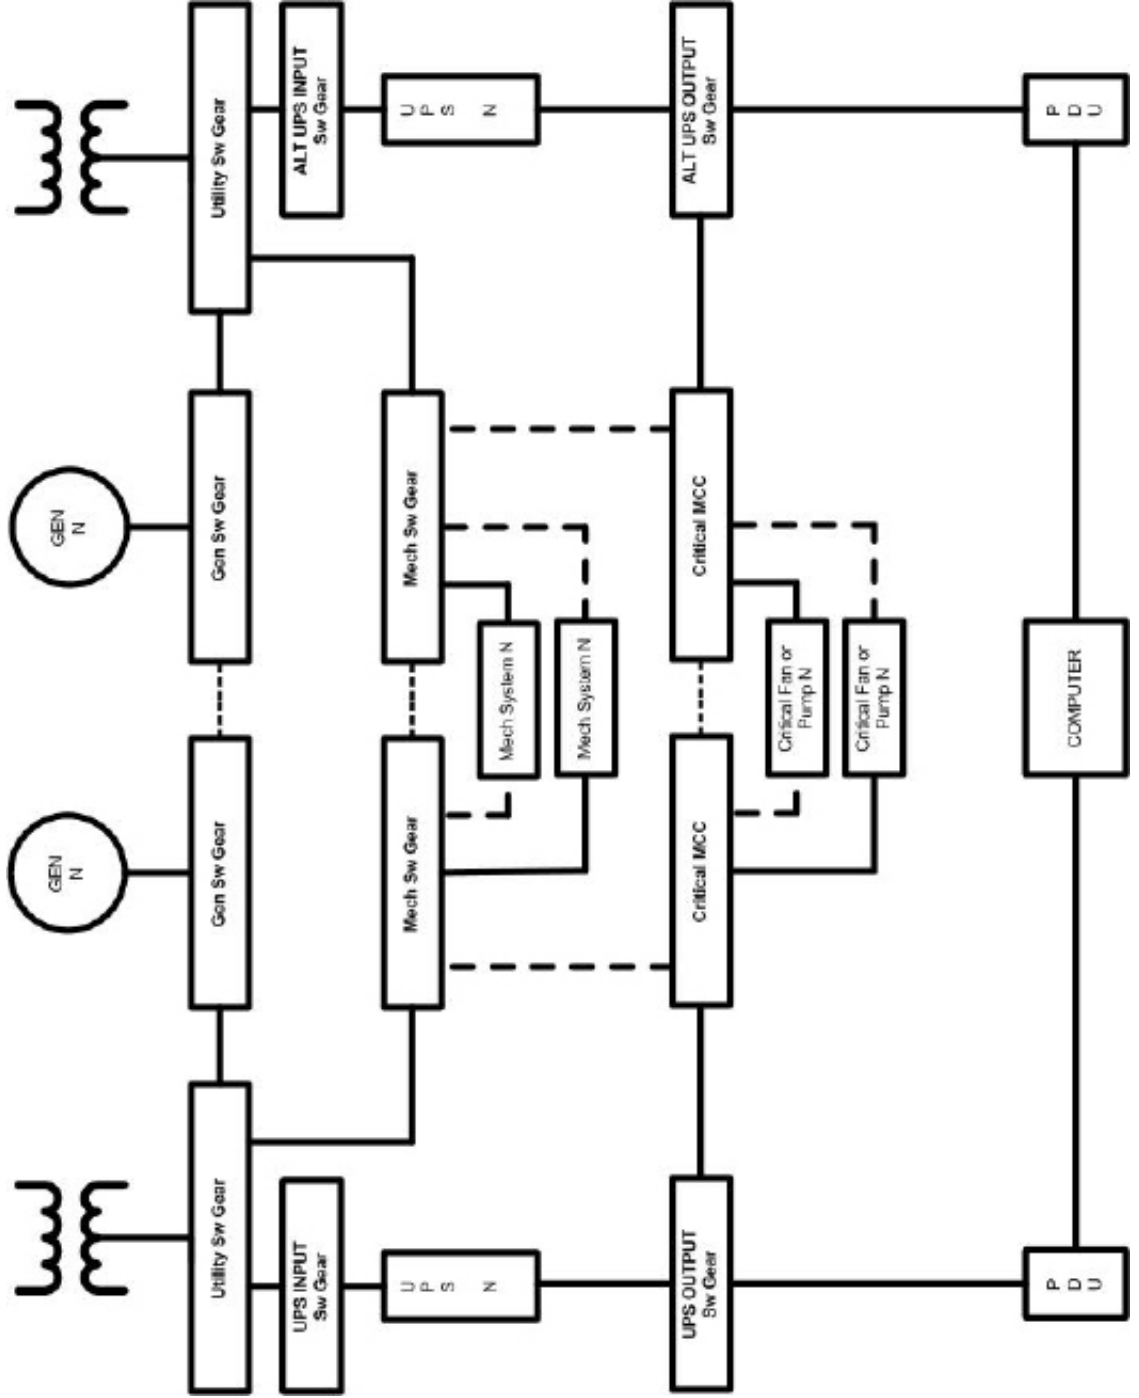


Illustration 4: Illustrative Electrical System Topology - Tier IV

Note: This diagram illustrates a Tier IV electrical distribution concept. Note the additional (redundant) equipment required to move from Tier III to Tier IV. This diagram should not be interpreted to represent a standard nor compliant electrical system topology or a solution fulfilling any particular set of requirements. Omitting the redundant UPS modules would not change the Tier IV rating.

Tier Certification requires consistent application of Tier concepts to all relevant critical subsystems that support the computer room.



Frequently Asked Questions

1. *Is there a specification or checklist of components for each particular Tier?*
A specification or checklist is not provided because having the “right number of components” provides no assurance of Concurrent Maintainability or Fault Tolerance. This position has been repeatedly validated by design reviews and site due-diligence activities of projects driven by checklists.
2. *Why doesn't the Tier white paper 'just tell us what to do'?*
A prescriptive approach (telling a client what the one-line should be, for example) limits innovation and the introduction of new concepts or solutions. This is not in the best interest of a client or the industry in general. Innovation is essential to continued development of the industry. Furthermore, Tier of topology determinations are driven by owner's business requirements.
3. *I have an N+3 data center. What Tier is it?*
There is no effective way to ascertain the Tier based on the number of redundant components. The “N+” phrase is almost exclusively applied to capacity components and rarely to distribution paths. The Tier standard addresses both.
4. *What does it mean to have outcome-based Tiers?*
Outcome-based means the availability discussed in the Tier white paper is based on the availability outcome one would expect given a certain Tier. Inevitably, the site infrastructure failure will happen. Outcome-based approaches go beyond the time required to restore the site infrastructure and include the time to restart hardware, software, application, etc. Supporting the end user is why a data center exists, and outcome-based analysis includes the impact to them.
5. *Are the Tiers based on facility availability alone or on the resulting impact to the end users from a facility failure?*
Availability levels associated with the Tiers are based on the time required to restore service to the end users. This includes the time required to restore the site infrastructure plus the time to restart hardware, software, application, and etc.
6. *Do the Tiers require me to include equipment that I may not need for my business requirements?*
The business requirements are the basis for any discussions on the topology of the site infrastructure. If a company only needs a Tier II facility (for a call center, for example), that is all they should build. The close connection between a business need and what is included in the facility is another reason the checklist approach is not used in this document. System capabilities should be built on business requirements, not someone else's checklist.
7. *Why must the engine-generator system be Concurrently Maintainable or Fault Tolerant when it is operating? Why not use the utility as a power source when doing maintenance on the engine-generator system?*
For sites that are serious about availability (Tier III and IV sites), the engine-generator system is the primary reliable source of power for the site. The power utility is an economical alternative provider but does not offer the stable, disruption-free power required for a data center. Using the power utility as a source of power during maintenance of the engine-generator system exposes the site to raw utility power in the event the UPS system goes to bypass. Utility power can be disrupted even in good weather by vehicle accidents and other events.
8. *My company spent a lot of money to get the Tier we needed. Why do we need to bother with Operational Sustainability?*
Extensive research indicates that the primary differentiators in the number and frequency of site failures are the factors discussed in the Operational Sustainability white paper. Sites that have incorporated these factors into their operation have substantially better results than their peer sites with the same Tier that have not.
9. *My engineer did a MTBF calculation. Does this tell me what Tier the site is?*
MTBF is a statistical analysis that attempts to describe the probability of a failure. These studies do not include the primary contributor to site failure: human actions and decisions. Some do not incorporate expected or normal events such as the power utility failing and the ensuing system changes necessary to transfer the site to the engine-generator system. MTBF analysis is not outcome-based as described above. The Institute assumes the probability of a failure is one: it is inevitable. The challenge is to design and manage a site in a way that reduces the number of failures and increases the interval between them.

The Uptime Institute's publications are protected by international copyright law. The Institute requires written requests whenever the Institute's literature or portions of the Institute's literature are reproduced or used. The Institute copyright extends to all media—paper, electronic, and video content—and includes use in other publications, internal company distribution, company websites, and handouts for seminars and courses. For more information, please visit uptimeinstitute.org/resources to download a Copyright Reprint Permission Request form.

Uptime
Institute

2904 Rodeo Park Drive East, Building 100, Santa Fe, NM 87505
505.986.3900 · Fax 505.982.8484 · uptimeinstitute.org

© 2008 Uptime Institute, Inc.

TUI3026F