

UPTIME INSTITUTE

Data Center Site Infrastructure Tier Standard: Topology

Prepared by Uptime Institute Professional Services

Copyright ©2009 by the Uptime Institute

Building 100
2904 Rodeo Park Drive East
Santa Fe, NM 87505

All rights reserved.

The Uptime Institute's Publications are protected by international copyright law. The Uptime Institute (Institute) requires written requests at each and every occasion that the Institute's intellectual property or portions of the Institute's intellectual property 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 marketing materials, and handouts for seminars and courses. For more information, please visit www.uptimeinstitute.org/resources to download a Copyright Reprint Permission Request Form.

Uptime
Institute™

Abstract: The Institute Tier Standard: Topology is an objective basis for comparing the functionality, capacity, and expected availability (or performance) of a particular site infrastructure design topology against other sites, or for comparing a group of sites. This standard describes criteria to differentiate four classifications of site infrastructure topology based on increasing levels of redundant capacity components and distribution paths. This standard focuses on the definitions of the four Tiers and the performance confirmation tests for determining compliance to the definitions. The Commentary, in a separate section, provides practical examples of site infrastructure system designs and configurations that fulfill the Tier definitions as a means to clarify the Tier classification criteria.

Keywords: data center, infrastructure, Tier, classification, Tiers, Tier level, topology, availability, reliability, redundant, Concurrent Maintenance, Concurrently Maintainable, Fault Tolerance, Fault Tolerant, Operational Sustainability, dual power, functionality, performance, metrics

Copyright ©2009 by the Uptime Institute
Building 100
2904 Rodeo Park Drive East
Santa Fe, NM 87505

All rights reserved.

The Uptime Institute's Publications are protected by international copyright law. The Uptime Institute (Institute) requires written requests at each and every occasion that the Institute's intellectual property or portions of the Institute's intellectual property 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 marketing materials, and handouts for seminars and courses. For more information, please visit www.uptimeinstitute.org/resources to download a Copyright Reprint Permission Request Form.

Introduction

This introduction is not part of the Uptime Institute Data Center Site Infrastructure Tier Standard: Topology. It provides the reader with context for the application of the standard.

This Uptime Institute Data Center Site Infrastructure Tier Standard: Topology is a restatement of the content previously published as the Institute white paper Tier Classifications Define Site Infrastructure Performance. Selected content of this white paper has been reedited into an ANSI Standards Model format. Future updates or changes to the Institute Tier Standard: Topology shall be accomplished through a review and recommendation process consistent with other recognized Standards bodies.

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 Information Technology (IT) environment in the computer room.

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 yields a Tier II site rating.

This very stringent definition is driven by senior executives who have approved multi-million dollar investments for an objective report of actual site capabilities. Any exceptions and exclusions footnoted in the approval documents will be quickly lost and forgotten. If a site has been advertised within an organization as being Fault Tolerant (Tier IV), it will be inconsistent 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 mean time between failures (MTBF) component statistical reliability to generate a predictive availability and then using that number to match the empirical availability results with those of sites representing the different Tier classifications. 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.

Finally, the Tier Standard focuses on the topology and performance of an individual site. High levels of end-user availability may be attained through the integration of complex IT architectures and network configurations that take advantage of synchronous applications running on multiple sites. However, this Standard is independent of the IT systems operating within the site.

Copyrights

This document is copyrighted by the Uptime Institute, Inc. The Institute—in making this document available as a reference to governmental agencies, public institutions, and private users—does not waive any rights in copyright to this document.

Participants

The original contents of the Data Center Site Infrastructure Tier Standard: Topology were developed by the following individuals:

W. Pitt Turner, IV

John H. Seader

Vincent E. Renaud

With editorial contribution by:

Julian S. Kudritzki

Kenneth G. Brill

Contents

- 1. Overview 1
 - 1.1. Scope 1
 - 1.2. Purpose 1
 - 1.3. References 1
- 2. Tier Classification Definitions 1
 - 2.1. Tier I – Basic Data Center Site Infrastructure 1
 - 2.2. Tier II – Redundant Site Infrastructure Capacity Components 2
 - 2.3. Tier III – Concurrently Maintainable Site Infrastructure 2
 - 2.4. Tier IV – Fault Tolerant Site Infrastructure 3
 - 2.5. Engine-Generator Systems 3
 - 2.6. Ambient Temperature Design Points 4
 - 2.7. Tier Requirements Summary 4
- 3. Commentary for Application of the Tier Standard: Topology 5
 - 3.1. Outcome-Based Tier Standard 5
 - 3.2. Impact of Ambient Design Conditions 5
 - 3.3. Tier Functionality Progression 6
 - 3.4. Fractional or Incremental Tier Classification 7
 - 3.5. Non-Compliance Trends 7

1. Overview

1.1 Scope

This standard establishes four distinctive definitions of data center site infrastructure Tier classifications (Tier I, Tier II, Tier III, Tier IV), and the performance confirmation tests for determining compliance to the definitions. The Tier classifications describe the site-level infrastructure topology required to sustain data center operations, not the characteristics of individual systems or subsystems. This standard is predicated on the fact that data centers are dependent upon the successful and integrated operation of several separate site infrastructure subsystems, the number of which is dependent upon the individual technologies (e.g., power generation, refrigeration, uninterruptible power sources, etc.) selected to sustain the operation.

Every subsystem and system integrated into the data center site infrastructure must be consistently deployed with the same site uptime objective to satisfy the distinctive Tier requirements.

Compliance with the requirements of each Tier 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.

Commentary on the Tier standard is in a separate section that focuses on examples and conceptual-level illustrations of the various ways to design and configure each Tier topology within each Tier level. The commentary section also offers guidance in the application and implementation of the Tier definitions. In addition, the commentary section includes discussion and examples to aid in understanding Tier concepts as well as information on common design topology shortfalls.

1.2 Purpose

The purpose of this standard is to equip design professionals, data center operators, and non-technical managers with an objective and effective means for identifying the anticipated performance of different data center site infrastructure design topologies.

1.3 References

American Society of Heating, Refrigerating, and Air-Conditioning Engineers, 2005 ASHRAE Handbook – Fundamentals. Publication.

Uptime Institute, Inc. Fault Tolerant Power Compliance Specification, Version 2.0. Specification.

Uptime Institute, Inc. Continuous Cooling Is Required for Continuous Availability. White paper.

2. Site Infrastructure Tier Standards

2.1 Tier I: Basic Site Infrastructure

2.1.1 The fundamental requirement:

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

2.1.2 The performance confirmation tests:

- a) There is sufficient capacity to meet the needs of the site.
- b) Planned work will require most or all of the site infrastructure systems to be shut down affecting computer equipment, systems, and end users.

2.1.3 The operational impacts:

- a) The site is susceptible to disruption from both planned and unplanned activities. Operation (Human) errors of site infrastructure components will cause a data center disruption.
- b) An unplanned outage or failure of any capacity system, capacity component, or distribution element will impact the computer equipment.
- c) 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.

2.2 Tier II: Redundant Site Infrastructure Capacity Components

2.2.1 The fundamental requirement:

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

2.2.2 The performance confirmation tests:

- a) Redundant capacity components can be removed from service on a planned basis without causing any of the computer equipment to be shut down.
- b) Removing distribution paths from service for maintenance or other activity requires shutdown of computer equipment.

2.2.3 The operational impacts:

- a) The site is susceptible to disruption from both planned activities and unplanned events. Operation (Human) errors of site infrastructure components may cause a data center disruption.
- b) An unplanned capacity component failure may impact the computer equipment. An unplanned outage or failure of any capacity system or distribution element will impact the computer equipment.
- c) 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.

2.3 Tier III: Concurrently Maintainable Site Infrastructure

2.3.1 The fundamental requirements:

- a) A Concurrently Maintainable data center has redundant capacity components and multiple independent distribution paths serving the computer equipment. Only one distribution path is required to serve the computer equipment at any time.
- b) All IT equipment is dual powered as defined by the Institute's Fault Tolerant Power Compliance Specification, Version 2.0 and installed properly to be compatible with the topology of the site's architecture. Transfer devices, such as point-of-use switches, must be incorporated for computer equipment that does not meet this specification.

2.3.2 The performance confirmation tests:

- a) 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.
- b) There is sufficient permanently installed capacity to meet the needs of the site when redundant components are removed from service for any reason.

2.3.3 The operational impacts:

- a) The site is susceptible to disruption from unplanned activities. Operation errors of site infrastructure components may cause a computer disruption.
- b) An unplanned outage or failure of any capacity system will impact the computer equipment.
- c) An unplanned outage or failure of a capacity component or distribution element may impact the computer equipment.
- d) Planned site infrastructure maintenance can be performed by using the redundant capacity components and distribution paths to safely work on the remaining equipment.
- e) During maintenance activities, the risk of disruption may be elevated. (This maintenance condition does not defeat the Tier rating achieved in normal operations.)

2.4 Tier IV: Fault Tolerant Site Infrastructure

2.4.1 The fundamental requirements:

- a) A Fault Tolerant data center has multiple, independent, physically isolated systems that provide redundant capacity components and multiple, independent, diverse, active distribution paths simultaneously serving the computer equipment. The redundant capacity components and diverse distribution paths shall be configured such that "N" capacity is providing power and cooling to the computer equipment after any infrastructure failure.
- b) All IT equipment is dual powered as defined by the Institute's Fault Tolerant Power Compliance Specification, Version 2.0 and installed properly to be compatible with the topology of the site's architecture. Transfer devices, such as point-of-use switches, must be incorporated for computer equipment that does not meet this specification.
- c) Complementary systems and distribution paths must be physically isolated from one another (compartmentalized) to prevent any single event from simultaneously impacting both systems or distribution paths.
- d) Continuous Cooling is required. For more information see the Institute white paper Continuous Cooling Is Required for Continuous Availability.

2.4.2 The performance confirmation tests:

- a) A single failure of any capacity system, capacity component, or distribution element will not impact the computer equipment.
- b) The system itself automatically responds ('self heals') to a failure to prevent further impact to the site.
- c) 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.
- d) There is sufficient capacity to meet the needs of the site when redundant components or distribution paths are removed from service for any reason.

2.4.3 The operational impacts:

- a) The site is not susceptible to disruption from a single unplanned event.
- b) The site is not susceptible to disruption from any planned work activities.
- c) The site infrastructure maintenance can be performed by using the redundant capacity components and distribution paths to safely work on the remaining equipment.
- d) During maintenance activity where redundant capacity components or a distribution path shut down, the computer equipment is exposed to an increased risk of disruption in the event a failure occurs on the remaining path. This maintenance configuration does not defeat the Tier rating achieved in normal operations.
- e) Operation of the fire alarm, fire suppression, or the emergency power off (EPO) feature may cause a data center disruption.

2.5 Engine-Generator Systems

Tier III and IV engine-generator systems 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.

2.5.1 Site on Engine-Generator Power

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 while they are carrying the site on engine-generator power.

2.5.2 Manufactures' Run Time Limitation

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.

2.5.3 Regulatory Run Time Limitation

Engine-generator systems often have an annual regulatory limit on operating hours driven by emissions. These environmental limits do not impact the consecutive hours of operation constraint established in this section.

2.6 Ambient Temperature Design Points

The effective capacity for data center facilities infrastructure equipment shall be determined at the peak demand condition based on the climatological region and steady state operating set points for the data center. All manufactures' equipment capacities shall be adjusted to reflect the extreme observed temperatures and altitude at which the equipment will operate to support the data center.

2.6.1 Extreme Annual Design Conditions

The capacity of all equipment that rejects heat to the atmosphere shall be determined at the Extreme Annual Design Conditions that best represents the data center location in the most recent edition of the ASHRAE Handbook – Fundamentals. (Each ASHRAE Handbook is revised and published every 4 years.) The design Wet Bulb (WB) temperature shall be the listed Extreme Max WB value and the design Dry Bulb (DB) temperature for design shall be the "n=20 years" value.

2.6.2 Computer Room Set points

The capacity for computer room cooling equipment shall be determined at the return air temperature, and relative humidity established by the owner for steady state data center operations.

2.7 Tier Requirements Summary

A summary of the preceding requirements defining the four distinct Tier classification levels is in Table 1.

Table 1: Tier Requirements Summary

	Tier I	Tier II	Tier III	Tier IV
Active Capacity Components to Support the IT Load	N	N+1	N+1	N After any Failure
Distribution Paths	1	1	1 Active and 1 Alternate	2 Simultaneously Active
Concurrently Maintainable	No	No	Yes	Yes
Fault Tolerance	No	No	No	Yes
Compartmentalization	No	No	No	Yes
Continuous Cooling	Load Density Dependent	Load Density Dependent	Load Density Dependent	Class A

3. Commentary for Application of the Tier Standard - Topology

This Commentary is not part of the Data Center Site Infrastructure Tier Standard: Topology. It provides the reader with context for the application of the standard.

3.1 Outcome-Based Tier Standard

The definitions used in the Institute's Tier Standard are necessarily and intentionally very broad to allow innovation and client manufacture and 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 operational performance outcomes that define the four Tiers of site infrastructure are very straightforward. Many designs that pass a checklist approach will fail an operational performance requirements approach. This means that, in addition to the rigorous application of engineering principles, there is still considerable judgment and flexibility in the design for uptime and how subsystems are integrated to allow for multiple operating modes.

Tier I sites typically experience 2 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 3 maintenance windows over a 2-year period and have 1 unplanned outage each year. The redundant components of Tier II topology provide some maintenance opportunity leading to just 1 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 (1) 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:

- a) The impact of the failure on overall availability is not represented by the likelihood of a system failure and
- b) The Institute does not recognize fractional Tier ratings. See section 3.4.

3.2 Impact of Ambient Design Conditions

The sustainable effective capacity of most cooling and power generating equipment is impacted by the actual ambient conditions in which it operates. These components typically require more energy to operate and provide less useable capacity as altitude and ambient air temperatures rise.

A common practice for conventional facilities is to select design values applicable to most but not all anticipated hours of operation of that facility. This results in an economical choice of equipment that meets requirements most of the time. This is not appropriate for data centers that are expected to operate on a 24 X Forever basis.

Using a dry bulb temperature for design that is exceeded 2% of the time results in selection of a component that is undersized 175 hours of the year. Although this may seem to imply that the owner runs an operational risk for a little over one week each year, these hours actually occur incrementally spread over several days. The 2% design value could result in actual conditions exceeding the design parameters of the equipment several hours every afternoon for a 1- to 2-month period. A 0.4% value, considered conservative by many design professionals, still results in equipment performing below requirements approximately 35 hours each year.

Another example concerning ambient conditions arises when selecting heat rejection systems for split system direct expansion cooling system. Many manufactures provide product selection tables based on 95°F ambient outside conditions. These components will only produce the nominal capacity listed when operating in up to 95°F outside air. These component capacities must be adjusted downward to provide the required capacity when temperatures exceed 95°F.

3.3 Tier Functionality Progression

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.

3.3.1 Tier I

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.

3.3.2 Tier II

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.

3.3.3 Tier III

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.

3.3.4 Tier IV

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.

3.4 Fractional or Incremental Tier Classification

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).

Consistent, across-the-board application of Tier topology concepts for electrical, mechanical, automation, and other subsystems is required for any site to satisfy the Tier standards defining any classification level. Selecting the appropriate topology solution based on the IT availability requirements to sustain well-defined business processes, and the 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.

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. 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.

- a) 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.
- b) 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.
- c) Including a UPS system patterned after a Tier IV system within a site having a Tier II power distribution backbone yields a Tier II Certification.

3.5 Non-Compliance Trends

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 unmaintainable 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. 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.



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