

AD _____

Award Number:

W81XWH-08-1-0585

TITLE:

Advanced patient Data Replication and Recovery

PRINCIPAL INVESTIGATOR:

Perez, David
Hendrian, Andrew

CONTRACTING ORGANIZATION:

Eisenhower Medical Center
Rancho Mirage, CA 92270

REPORT DATE:

October 2009

TYPE OF REPORT:

Annual

PREPARED FOR: U.S. Army Medical Research and Materiel Command
Fort Detrick, Maryland 21702-5012

DISTRIBUTION STATEMENT:

Approved for public release; distribution unlimited

The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision unless so designated by other documentation.

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 01-10-2009		2. REPORT TYPE Annual Report		3. DATES COVERED (From - To) 8 SEP 2008 - 7 SEP 2009	
4. TITLE AND SUBTITLE Advanced Patient Data Replication and Recovery at Eisenhower				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER W81XWH-08-1-0585	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Perez, David Hendrian, Andrew E.				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Eisenhower Medical Center Rancho Mirage, CA 92270				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Medical Research and Materiel Command Fort Detrick, Maryland 21702-5012				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT 1. The move to electronic medical records (EMR) necessitates that clinical data protection and recovery best practices support extremely low RPO and RTO. Only near-real time and synchronized backups to offsite, disk based, storage technologies support sub-four-hour RPO and RTO while still protecting the data from local loss. 2. Objectives include: A. Lower the risk of clinical patient data loss to clinical staff B. Support clinicians dependence on EMR data by making it less prone to loss C. An "IT healthcare best practice" will be redefined for off-site real-time data replication of electronic medical records which will lower RPO and RTO to less than the current levels of 24 and 48 hours. D. Target objectives for RPO and RTO will be 1 to 4 hours.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 36	19a. NAME OF RESPONSIBLE PERSON USAMRMC
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (include area code)

Table of Contents

	<u>Page</u>
Introduction.....	3
Body.....	5
Key Milestones.....	5
Reportable Outcomes.....	13
Conclusion.....	14
Appendices.....	15

Introduction

Advanced Patient Data Protection (APDAPT)

Eisenhower Medical Center (EMC) is attempting to lower the risk of losing patient data, as well as the risk incurred by lengthy recovery processes in the case of a data loss, by making available, in near real-time, a duplicate electronic medical record which includes radiological images. EMC has made a multi-million dollar investment in the digitization of patient data; the Electronic Medical Record (EMR). Moving from a process firmly entrenched in the use of paper forms, verbal authorizations, and hand written notes, EMC has digitized the creation, storage and retrieval of the patient chart or EMR. The EMR is comprised of patient vital signs, nurse notes, medications administered, doctors' orders, dietary and radiology orders, radiological studies and results, lab orders and results as well as transcriptions, etc. Undertaking the change to an EMR has required a massive departure from decades-old processes that were intrinsically tied to paper records and manual procedures.

This change necessitates that clinical data protection and recovery best practices support extremely low recovery point objectives (RPO) and recovery time objectives (RTO.) Recovery Point Objective describes the acceptable amount of data loss measured in time.

The Recovery Point Objective (RPO) is the point in time to which you must recover lost data as defined by your organization. This is generally a definition of what an organization determines is an "acceptable loss" in a disaster situation. Traditional backup strategies for many hospitals have focused on business or financial data protection. These strategies considered a 24 hour backup window an acceptable risk. Backups would be created every 24 hours. If a data loss occurred 23 hours and 59 minutes later, there would be no back up for that previous time period back to the previous back up. This represents "lost" data, and the lost data would be recovered from other sources, such as printed reports, statements, etc. This recovery of a "days" worth of data was deemed acceptable. When considering other data types, such as clinical data, that risk is no longer acceptable. For EMC the RPO has been 24 hours. Based on this RPO the data must be restored to within 24 hours of the disaster. All data from the point of the disaster to 24 hours later will have to be manually recovered through other means. In healthcare, with records no longer paper-based, this could prove to be impossibility with some clinical data sets, like verbal pharmacy and lab orders.

The Recovery Time Objective (RTO) is the duration of time and a service level within which a business process must be restored after a disruption in order to avoid unacceptable consequences associated with a break in business continuity. The RTO includes the time for trying to fix the problem without a recovery, the recovery itself, tests, as well as communication to those who use the systems affected. This time frame is usually an objective or goal for an organization, not a mandate. Strategy is often selected that will not meet the RTO. EMC's strategy will be to find a solution that will meet the objective.

While tape based backup have been traditionally the medium of choice, the time it takes to recall off-site tapes and restore from tape can add dozens of hours to a recovery. As such, EMC has decided to investigate a disk-based storage technology that can support sub four-hour RPO and RTO. Moving a fault-tolerant copy of the data offsite using such a technology will protect the data from local loss.

EMC endeavors to achieve the following in regards to clinical data:

1. Lower the risk that EMC will be unable to access patient data from the EMR due to data loss.
2. Reduce the risk of loss of PACS data elements in the event of disaster to EMC's local databases.
3. Increase the availability of the EMR data by lowering the figures for RPO and RTO from 24 and 48+ hours respectively. Target objectives for RPO and RTO will be 1 to 4 hours.
4. Re-define an "IT healthcare best practice" for the protection and recoverability of electronic patient data through the utilization of an off-site, real-time, replication of electronic medical records which will lower RPO and RTO to less than the current levels of 24 and 48 hours (respectively) in the event of data loss within the EMR.

Body and Key Milestones:

Eisenhower Medical Center's approach to reducing risk to the electronic medical record by reducing the time needed to recover from a data loss is to replicate our clinical data to an asynchronous disk-based technology outside of the seismic disaster zone of southern California. The progress achieved so far has been limited due mainly to resource constraints. However, as described in the following pages, we have made significant progress in this project. In January of 2009 we finalized on a hardware/software mix to protect the EMR data. We also finalized the data to be protected as well. A remote hosting site selection criterion was also established and approved by Eisenhower VP/CIO. The remote host search has begun and will be concluded by late 2009. Data replication from our existing archive platform was also begun and completed. The survey to query clinicians as to the prioritization of importance of clinical computer applications was also created and will be sent for protocol approval.

- Select a platform of software and hardware for replication and storage of our medical records and PACS images.

Jan 2009

Eisenhower Medical Center (www.emc.org) selected EMC² (www.emc.com) Centera line of data storage technologies to be the host to clinical data as a secondary pool. Centera is a mechanism for storing information that can be retrieved based on its content, not its storage location. It is used for high-speed storage and retrieval of fixed content, such as medical records and clinical modalities. One of the best aspects of the Centera is that archived information becomes immutable—it cannot be changed. Centera protects the authenticity and security of archived information by applying a hash algorithm to every file. If someone were to change even one character, the system recognizes that change and saves it as a new object. When dealing with a legal medical record that contains PACS images or EMR data, record immutability is of supreme importance. Additionally the US Army currently utilizes the Centera line for image storage. The Eisenhower Medical Center Centera storage pool is comprised of 192 terabytes (TB) of raw capacity. A services and hardware contract was negotiated and executed between Eisenhower Medical Center and EMC2 Corp. in December of 2008. Delivery of hardware and software was completed in late January of 2009 to our campus. Additionally, data center infrastructure needed alterations and additions to accommodate the storage. New network cables were run to existing network infrastructure and eight additional power circuits (208V, 30 Amp) were installed to accommodate the power requirements of the two new storage cabinets.

Our McKesson Horizon Medical Imaging (HMI) PACS system has native application programming interfaces (API) to the Centera storage archive pool. Because of this, the HMI can do duplicate writes to multiple storage targets. A major challenge we have is that our McKesson Horizon Patient Folder application (HPF) cannot do native API calls to multiple storage targets. As such, a work around was needed to enable the multiple data streams from HPF. This difficulty

was discovered rather late in the project. The resolution was to have the HPF application write to a third-party application (Legato Disk Extender) which then wrote to the primary Centera. The primary Centera (PC) would then run native EMC² Replication protocol to the Secondary Centera (SC), essentially duplicating the archive. It was this method that was followed to originally populate the SC with data from the PC. Once the SC ships offsite, it will then utilize the same protocol to “catch-up” and synchronize with the PC.

- Identify the type of clinical data to be replicated.
Jan 2009

The data to be protected consists of archived medical records from 2005 as well as radiological exams performed since 2004. While EMC has several source applications that create electronic medical records, the records being archived are from our Horizon Patient Folder (HPF) application (a McKesson product), which is an EMR aggregator application. HPF takes feeds from other EMR systems and compiles them by patient and date of service as well as location of service. Radiological exams originate from our McKesson Horizon Medical Imaging application (HMI) and include in all images from the following modalities: x-ray, CT scan (16 and 64 slice), mammography, magnetic resonance imaging (MRI.) Eisenhower was collecting about .1 TB of medical record data per month and a little over .1 TB of radiology data each month in 2007. That number has risen to .2 TB for EMR and .15 TB for PACS data each month. See Appendix B: PACS Study FY2010 for details on growth of studies and storage requirements for PACS images. The numbers actually represent a stalled growth trend due to the economic downturn, which has negatively affected EMC’s patient volumes. Had the economy remained at the level prior to the housing crash of 2008, the growth rates would have been much higher. We expect this growth to pick up over the next 12 months.

- Begin to replicate current clinical data set to the selected technology platform.
May 2009

Power and data-networking resources were installed in the existing data center in Rancho Mirage. The Centera cabinets (2) were delivered, installed, and configured by EMC² and Eisenhower staff. Peer-to-peer archiving was configured between the source and target data locations. This was to make the initial copy of our onsite archive to the new archive Centera. The devices were setup side-by-side and a “trickle” synchronization was initiated. After 47 days of “trickle” data synchronization, the target and source pools were identical on both the primary Centera and the secondary Centera. Trickle synchronization was employed to ensure that the “live” data pool was not disturbed and that clinical workflow was not interrupted by system latency due to more aggressive synchronization rates. Aggressive synchronization would have sped up the process of copying the data pool of 40+ TB of data by three weeks time. However, archive PACS image retrieval would have been significantly slowed in

the interim. EMC decided that clinician impact should be minimized and the slower trickle synchronization would suffice for the first copy of data. Eisenhower radiologists and medical records recorders were not impacted by the trickle copy.

- Identify and select appropriate location criteria for the off site replicated data.
August 2009

The site criteria have been agreed upon and defined as meeting the following items:

1. The remote-site must be outside the regional disaster area of southern California as defined by the United States Geologic Survey National Seismic Hazard Survey. “The 2008 National Seismic Hazard Maps represent the ‘best available science’ (regarding ground movement vectors and probabilities) based on input from scientists and engineers that participated in the update process.” (U.S. Geological Survey Open-File Report 2008–1128: 2008 National Seismic Hazard Maps - Peterson, Mark D., et. al. pg 40.) See Seismic Hazard map below for relative distances from Eisenhower Medical Center’s location in Rancho Mirage, CA. Eisenhower Medical Center and possible remote sites have been superimposed over the National Seismic Hazards map. Note that EMC sits in one of the most at risk areas in the western United States as far as the rate of peak ground acceleration probability. Our goal is to move our secondary archive outside of these high risk areas into areas with much lower peak ground acceleration rates.

Documentation for the 2008 Update of the United States National Seismic Hazard Maps




Figure 39. Map of peak ground acceleration (PGA) for 10-percent probability of exceedance in 50 years in the Western United States in standard gravity (g).

The next map shows Eisenhower Medical Center's position relative to major known faults in southern California.



Vertical faults such as the San Andreas (red band from top left to bottom right) are shown as a thin strip. Faults that are at an angle to the surface are shown as wider ribbons as they lie beneath broad areas (the nearest fault to you might be a few miles beneath your home). Areas that seem to have few faults can still experience strong shaking from earthquakes on unmapped faults or from large earthquakes on distant faults. (Putting Down Roots in Earthquake Country, Lucile M. Jones, USGS and Mark Benthien, SCEC 2008) These USGS and SCEC resources have been instrumental in determining seismic safe-zones nearby to Eisenhower Medical Center.

2. The remote-site must not be in a coastal area.
3. The remote site city must be in a zone of seismic activity that is less than the seismic activity of the southern California seismic zone. This zone runs the entire length of the state of California and comprises the width (from west coast of Los Angeles basin to the eastern state line due west of L.A.) of southern California.
4. The remote site city must be drivable within 10 hours of EMC via Interstate highway 10, Interstate highway 15, or Interstate highway 8. The naming of these thoroughfares is not a requirement to use them in the event of an emergency. Other roads may be used.

5. The remote site data center must be within a 1-hour drive of a commercial airport in the remote site city that receives traffic from the major commercial freight and passenger airlines.
6. The remote site city should be within 2 hours flight time from either Ontario International Airport in Ontario, Ca or Palm Springs International Airport in Palm Springs, CA.
7. The remote site city airport should be accessible via a non-stop flight from Palm Springs airport or Ontario, California airports with Palm Springs airport being the most desired starting point. This is a preference, not a hard requirement.
8. The remote site data center facility should be designed and built to Tier III or greater standards as defined by the Uptime institute. However the site does not need to be certified by the Uptime Institute as Tier III. The major requirements in Tier III standards are: 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. Additionally, all IT equipment is dual powered and installed properly to be compatible with the topology of the site's architecture. The site would have Tier II or Tier III power utility redundancy and Tier III physical and technical security. A copyrighted Uptime Institute whitepaper titled "Tier Classifications Define Site Infrastructure Performance" has been included in Appendix C. More information regarding the Tier standards can be found there. These principles were heavily relied upon by Eisenhower Medical Center to create these criteria. 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 institute prepares white papers documenting best practices for use by the industry.
9. The site would need to be able to host Eisenhower staff with appropriate workspaces, including telephone and internet service, for several weeks if needed.
10. Hotels must be available within 15 miles or 30 minutes of the selected sites.
11. The remote-site should have multiple data utility Internet Service Providers (ISP) providing service to the hosts.
12. TIA -942 Telecommunications Infrastructure Standards for Data Centers would be required to be followed for electrical grounding and data pathing, fire suppression, networking and cooling for the selected remote site.

13. The ongoing cost of the hosted remote-site must be sustainable by Eisenhower Medical Centers annual operating budget limitations and be approved by the Vice President / CIO of Eisenhower Medical Center.

This safety zone for the remote-site has been determined to be at least 200 miles to the north of the Coachella valley or 300 miles east of the valley. Sites west and south were deemed as inappropriate. Westward locations offered no respite from the seismic risks associated with our current location. Southward locations lacked necessary infrastructure to accommodate this projects goals. The most immediate sites are Las Vegas, NV and Phoenix, AZ. Other sites were also reviewed in the cities of Denver, Boise, and Salt Lake City.

Milestones not accomplished in the originally expected timeframe:

- Protocol (survey) approval and administration.
The research survey has taken longer than expected to create. Competing projects also hindered the researchers' efforts to complete the survey. Departmental priorities included in-sourcing an IT Help Desk, Anesthesiologist Documentation System implementation, major upgrade to our ED and nursing Electronic Medical Record, various major construction projects, the implementation of Computerized Physician Order Entry (CPOE) systems, as well as staffing limitations due to the economic slowdown of 2008-2009. Additionally, the subjects of the test (the clinical staff of Eisenhower Medical Center i.e. nurses, phlebotomists, radiologists, hospitalists, anesthesiologists, case managers, rehabilitation therapists, pharmacists, lab technicians, affiliated physicians, Emergency Department doctors, etc.) were also deeply engaged in the conversion of the paper chart to an electronic medical record. Protocol approval is still required for Eisenhower Medical Center as well as for the US Army. While the survey has been created, it has received neither EMC or US Army protocol approval. We expect this approval to be completed shortly.
- Move hardware to the off site alternative location.
There has been a delay in selecting a remote-site. That delay translated into a delay in getting the copied hardware to the remote-site. This milestone will soon be completed in late 2009 or early 2010 once the site is selected.
- Setup feed from the primary data source to keep the remote database up to date.
The above noted delays and staff constraints contributed to not achieving this milestone. Completion of this task is estimated to occur in early 2010 in conjunction to the selection of a suitable site and delivery of the hardware platform to that site.
- Create a fail-over mode that will move the application data connection

The above noted delays and staff constraints contributed to not achieving this milestone. The new time line for completion of this milestone is February 2010.

Next Steps:

Once the above milestones have been accomplished EMC will administer a survey (attached at Appendix A) to a wide range of clinicians (nurses, doctors, pharmacists, clinical technicians, etc.) The desired objective is to help set the framework for a clinical IT best practice for disaster recovery priorities for clinical application sets. The survey will be submitted to thousands of clinicians across the country to cast as wide a net as possible in order to gather as much opinion as possible regarding the relative importance of clinical information systems. The outcome will be to have a standardized disaster recovery priority for clinical information and documentation systems. Additionally, EMC will begin testing the failover to the replicated EMR archive once it has been setup in the remote data center site. The failover tests will then be timed and retimed to determine the expected archive failover durations. Failover procedures will also be prepared and tested and refined.

Reportable Outcomes:

While the overall objective of the research is still in the early phases, we have garnered some valuable lessons. Almost right off the bat, as project manager, I battled resource constraints as competing project demands strained and then drained away resources from this research project. While the competing projects were prioritized as high visibility by the administration of the hospital, skilled resources were, none the less, needed for this research project. The role of technical project manager proved essentially impossible to backfill. As a result, the research portion of the project was stalled for about 6 months. In order to further resist the drain of resources by future projects we have put the Advanced Data Protection project on the organizations official project portfolio so there is visibility by the administration into this project. Since that time, we have had no resource constraints due to competing projects.

One of the most surprising lessons centered on our research objective of categorizing our clinical applications in regards to disaster recovery. Categorization or prioritization of clinical systems will hopefully give us insight into the recovery order of systems in the event of a major outage. Pulling together a detailed survey of clinicians to identify what, to them, is the most critical clinical applications is central to identifying a healthcare IT best practice for protecting electronic medical records and securing patient safety. We expect to gain much valuable data from this survey. However, in a non-affiliated project of moving our data center from one location to a newly constructed data center, as we prepared to take down system after system, our clinician community identified systems critical to their business/clinical success. This listing proved to be surprising. While EMR systems were near the top of the list, the systems involved in the medication administration cycle were identified as *most crucial*. These systems included the pharmacy system, the medication administration bar-coding system, the pharmacy automation robot (which fills pharmacy orders) and the drug dispensing kiosks located near the nurse stations in the hospital. These systems have become so central in patient safety initiatives, and nurses have become so accustomed to the automation achieved, that reverting to manual downtime procedures proved to be unsettling to many nurses. Many of these same nurses, only 4 years previous, were dubious of the need for such systems. Manual processes centering on the medication administration and reconciliation process had to be resurrected and refined in order to achieve the safety level EMC was accustomed to before the data center move. Additionally nurses and pharmacists had to be re-familiarized and retrained in these manual and paper processes and procedures.

Conclusions:

Results are as yet, inconclusive since the study survey has not been administered and the hardware solution is still not at a selected remote data center. Once the secondary data repository is secure and operational, recovery time objective (RTO: the duration within which the EMC EMR business process must be restored after a disruption in order to avoid unacceptable consequences associated with a break in EMR continuity) can then be quantified along with the recovery point objective (RPO.) There does not appear to be any technical reasons at this point why our original objectives of less than 4 hour RPO/RTO cannot be achieved. However, the study is not far enough in actual testing scenarios to begin analysis.

Additionally, the IT Healthcare best practices will be further refined after administration and analysis of the protocol results. These best practices sounding prioritization of clinical system protections and restoration priorities among the many clinical systems (EMR's, lab, pharmacy, radiology, transcription, and other systems) will enhance the way in which IT in healthcare can further support the mission critical nature of healthcare.

Appendices:

- A. Clinician Survey for Application Restoration Priority
- B. EMC PACS Study FY2010
- C. Tier Classifications Define Site Infrastructure Performance; The Uptime Institute, Inc, W. Pitt Turner IV, P.E., et.al.

Available at <http://www.iswest.com/colocation/TierClassification.pdf>

Clinical Information Systems and Data Availability Survey:

1. If clinical information systems were not available, which ONE system and its data set would be most urgent to providing care to your patient? i.e. Which system would you want restored to working order first?
 - a. Laboratory system/data
 - b. Radiology system/data
 - c. Electronic documentation system/data (Documentation, Meds Admin, etc)
 - d. Pharmacy system/data
 - e. Patient Folder system
 - f. STAR Patient registration system/data
 - g. Other systems (please note which system)
2. Which would be the next system or data set that you would require most urgently?
 - a. Laboratory system/data
 - b. Radiology system/data
 - c. Electronic documentation system/data (Documentation, Meds Admin, etc)
 - d. Pharmacy system/data
 - e. Patient Folder system
 - f. STAR Patient registration system/data
 - g. Other systems (please note which system)
3. Please list the remaining systems in order of descending urgency. You may add to the list any system not recorded above.
 - a.
 - b.
 - c.
 - d.
 - e.
 - f.
 - g.
4. For the time frames listed below rate the ease or difficulty with which you are able to continue to provide care without the use of these information systems by circling the appropriate number? 1 represents the lowest impact to providing care and a 5 represents the greatest difficulty in providing care for the designated time frame.
 - a. 1 hour 1 2 3 4 5
 - b. 2 hours 1 2 3 4 5
 - c. 4 hours 1 2 3 4 5
 - d. 8 hours 1 2 3 4 5
 - e. 24 hours 1 2 3 4 5
 - f. 2 days 1 2 3 4 5
 - g. 5 days 1 2 3 4 5
 - h. 7 days 1 2 3 4 5
5. Are you aware of Eisenhower Medical Center's downtime procedures for extended downtimes in the event the following systems are not available?
 - a. Laboratory system/data
 - b. Radiology system/data
 - c. Electronic documentation system/data (Documentation, Meds Admin, etc)
 - d. Pharmacy system/data
 - e. Patient Folder system
 - f. STAR Patient registration system/data

- g. Other systems (please record which system)
- 6. How long can you provide appropriate patient care utilizing downtime procedures?
 - a. 4 hours
 - b. 8 hours
 - c. 24 hours
 - d. 2 days
 - e. 5 days
 - f. 7 days
 - g. 30 days
 - h. other



EISENHOWER MEDICAL CENTER
Information Systems
PACS Review
March 2010

PACS (Horizon Medical Imaging version 11.5.1)

- 1) Originally installed November 2003 for EIC. These servers are 6.5 years old.
- 2) Additional servers installed January 2005 for the hospital implementation. These servers are 5 years old.
- 3) All workstations from every phase are still in production. Oldest radiologist workstations are 6.5 years old.
- 4) IRL network is 100 Mb/s. Need upgrade to 1000 Mb/s or Gigabit.
- 5) Need to convert EIC, Breast Center and Hospital Basement to new 20 Gb/s backbone. Dolores Hope is already on this network.
- 6) PACS handling nearly twice the image size that it was sized for in 2005 due to increase in hospital imaging and number of remote sites.

PACS Data Analysis

	2005	2009	Variance '05-'09	2010 (est)	Variance (est) '05-'10
Number of Studies (all sites)	157,821	188,906	20%	220,000	39%
Total Image Size (Terabytes)	5.8	10.3	78%	12	107%
Number of Imaging Modalities	85	105	24%	123	45%

Notable Facts

Hospital CT	3.5x more images (size) in 2009 than in 2005
	Added Monterey
End of 2009	EIC
	Added EDOC, Argyros EIC/BC/ICC/Executive
2010	Health
End of 2010	Palm Springs EIC/BC/ICC

EIC=Eisenhower Imaging Center
BC=Breast Center
ICC=Immediate Care/Express Clinic

Author: Scott McCabe, Senior RIS/PACS Administrator
smccabe@emc.org 760-340-3911 x3620
Revised: March 19, 2010
Version: 1.0



Tier Classifications Define Site Infrastructure Performance

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

Widely accepted within the uninterruptible industry, The Uptime Institute's Tier Performance Standards are an objective basis for comparing the capabilities of a particular design topology against others or to compare groups of sites. This paper defines a four Tier system providing discussion and illustrations of each classification. Significant cautions about Tier misapplication are provided. While the paper focuses primarily on design topology, sustainability (how the site is operated once constructed) plays a more significant role in what site availability is actually achieved. Actual site performance figures combining both design topology and sustainability are presented by Tier classification.

This white paper:

- Equips non-technical managers with a simple and effective means for identifying different data center site infrastructure design topologies.
- Provides IT based definitions and performance requirements for each Tier Level.
- Provides actual 5-year availability for 16 major sites by Tier classification.
- Warns that site availability is a combination of both design topology and "sustainability" with considerable optimization "art" involved.
- Warns that component/system counts or MTBF analysis plays no role in determining Tier compliance partially because each fails to include sustainability factors which account for 70% of all infrastructure failures.
- Cautions "self proclaimed" Tier claims all too often turn out to be misleading, incomplete, or wrong.
- Outlines need for third-party validation of site selection, design, and sustainability decisions before committing to multi-million dollar projects.
- Provides a commentary on typical Tier attributes.

Background

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

With the continuously increasing pressure on high availability comes an increased demand for computer

hardware reliability. Information technology customers expect availability of "Five Nines" or 99.999%. Unfortunately, the substantial investment a business frequently makes to achieve Five Nines in its computer hardware and software platforms is likely to be insufficient unless matched with a complementary site infrastructure that can support their availability goals. The site infrastructure includes 16 power, cooling, and other critical physical layer environmental sub-systems that must work together as a tightly integrated uptime system.

Tier History

The Uptime Institute, Inc.[®] (*Institute*) developed a four tiered classification approach to site infrastructure functionality that addresses the need for a common benchmarking standard. The *Institute's* system has been in use since 1995 and has become the default standard for the uninterruptible uptime industry. An early-1990s Tier predecessor outlined seven ways of distributing critical power to the computer equipment, but was not simple and all inclusive. A broader standard was required.

Creation of the *Institute's* original Tier definition was stimulated by multiple industry requests. Senior management decision makers needed a simple and effective non-technical means of conveying the differences in data center investments. Since the original pioneering work done more than 10 years ago, the Tier concept has been further developed and validated by broad industry use. The *Institute's* objective performance-based standard is very useful in ensuring a consistent framework to compare various alternatives companies may consider for obtaining data center space. These include such options as owned, leased, third party providers, and so on.



Site Availability As Actually Experienced By Information Technology

The following tier commentary includes actual measured results for site availability ranging from 99.67% to more than 99.99%.

These figures are not predictive of future site results, but do reflect actual operating experience at a specific list of sites representing the four Tiers of functionality. It is important to note that this range of availability is substantially less than the current Information Technology (IT) expectations of Five Nines. This leads to the conclusion that site availability limits overall IT availability.

Four Tier Levels Reflect Evolution of Data Center Uptime Objectives

Over the last 40 years, data center infrastructure designs have evolved through at least four distinct stages, which are captured in the *Institute's* classification system. Historically, Tier I first appeared in the early 1960s, Tier II in the 1970s, Tier III in the late 1980s and early '90s, and Tier IV in 1994. The *Institute* participated in the development of Tier III concepts and pioneered in the creation of Tier IV. Tier IV electrical power distribution systems were made possible, in part, by Ken Brill, Executive Director of the *Institute*. In 1991, he envisioned a future when all computer hardware would come with dual power inputs. This became US Patent 6,150,736. United Parcel Service's 1994 Windward data center project was the first Tier IV design. During construction of the Windward project, United Parcel Service worked with IBM and other computer hardware manufacturers to provide dual-powered computer hardware¹.

Tier IV technology requires having at least two completely independent electrical systems. These dual systems supply power through diverse power paths to the computer equipment. This effectively moves the last point of electrical redundancy from the Uninterruptible Power Supply (UPS) system downstream to a point inside the computer hardware itself. Brill's intuitive conclusion has since been confirmed by *Institute* research that has determined that of the vast majority of site infrastructure electrical failures occur between the UPS and the computer equipment. Since completion of the Windward project in 1994, System plus SystemSM (S+S) Tier IV electrical designs have become common and the number of computer devices with dual inputs has grown dramatically. There are exact parallels in the mechanical systems design.

The advent of dual-powered computer hardware in tandem with Tier IV electrical and mechanical infrastructure is an example of site infrastructure design and computer equipment design working together to achieve higher availability. Even with the significant improvements in computer hardware design made over the past 10 years, many data centers constructed in the last 5 years, and even today, claim Tier IV functionality, but actually deliver only Tier I, II, or III. This constrains their capability to match the availability required by the information technology they support. The purpose of this paper is to outline what it takes to consistently meet the requirements of the different tier levels.

The Need for Third-Party Certification Is a Growing Self-Preservation Requirement

In site infrastructure design and operation, the "devil is in the details" and the truth about a particular design topology will ultimately come out, but all too often after the warranty period has expired. When this happens, it can be a career ending event. Forensic investigation by the *Institute* into thousands of Abnormal Incidents over the last 12 years indicates that at least five and often seven interacting problems are required before a downtime failure occurs. The database upon which this analysis is built is unique in the world.

Increasingly, senior executives desire to have their critical sites independently certified as being compliant to the Tier standards. This provides a validation that the technical details of what the designer designed and the contractor built is actually what the owner wanted. When project designers "self proclaim" a site meets a certain tier level or capacity, it is all too often inaccurate or only partly factual. The results can often be tragic involving unnecessary downtime and tens of millions in unforeseen upgrade expense.

Certification is a service performed by The Uptime Institute, who is uniquely qualified to interpret and apply the standards since the *Institute* created the underlying technology concepts that allowed the standards to develop in the first place. In addition, the *Institute* also brings awareness of emerging downtime problems and trends at least three to five years before they are commonly recognized and addressed by the rest of the industry.

Site Certification by The Uptime Institute involves two separate, interrelated activities. The first is verification of the design topology and how it complies with the

¹ There are 13 technical requirements to describe what is commonly called "dual power." The actual details and additional dual power information can be found in the *Institute's* white paper *Fault-Tolerant Power Certification Is Essential When Buying Products for High-Availability* which may be found at www.uptime.com/whitepapers.



Tier standards. The second phase is verification of site sustainability. While a particular topology design may meet the literal requirement of a Tier level, the lifecycle effectiveness of that design may be extremely limiting — typically less than five years. Sustainability includes site selection; lifecycle effectiveness of the design topology and its transparent flexibility/scalability; ease of use; staffing level and coverage, training, and skills development; management procedures and processes; metrics and dashboards; commissioning and maintenance practices; and the integration of the site infrastructure with the IT architecture. Human factors are important because 70% or more of all site failures involve people. Of these failures, 2/3 are management error and 1/3 is human error. Human sustainability factors will largely determine the actual level of site availability achieved.

Previous Tier Level Information Is Now Divided into “TIER PERFORMANCE STANDARDS” and “COMMENTARY” Sections

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

- The TIER PERFORMANCE STANDARDS are now in a totally separate section, similar to many 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 confirmation tests and operational results. This is totally different than a prescriptive approach or a specific list of equipment not guaranteeing a performance outcome.
- The TIER COMMENTARY focuses on examples of the various ways to design and configure each Tier. In addition, the commentary section includes discussion and examples to aid in Tier understanding and information on common design topology failures. A comparison table of typical Tier attributes, availability and cost are provided. The commentary section also offers guidance in the comprehension, design, implementation, and the use of the Tier definitions.

Definition of Terms Used in the TIER PERFORMANCE STANDARDS and TIER COMMENTARY Sections

- Computer equipment: This is a broad phrase encompassing all information technology equipment required at a data center to perform the information processing work. It includes servers, storage, network, and all other information technology components.
- Redundant capacity components: The components beyond the number of capacity units required to support the computer equipment are 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 or N+2 are commonly applied.

- Useable capacity: This is the maximum amount of load that can be applied to the “N” level of capacity. Typically, the maximum amount of useable load is less than the non-redundant capacity to allow for component aging, installation errors, and to provide a contingency for unexpected demands.
- Useable capacity: This is the maximum amount of load that can be applied to the “N” level of capacity. Typically, the maximum amount of useable load is less than the non-redundant capacity to allow for component aging, installation errors, and to provide a contingency for unexpected demands.
- Site infrastructure: This comprises all of the site facility that includes the central plant plus the equipment that supports the power and cooling in the computer room. It is important to remember that a typical data center site is composed of at least 20 major mechanical, electrical, fire protection, security and other systems. Each has additional subsystems and components.
- Fault tolerant: This means that a system can sustain a worst case, unplanned event and not disrupt the end user. The fault tolerant concept originated in the IT environment. In the site infrastructure world, it means that the computer equipment will not be impacted by a facility failure. This requires multiple sources and multiple 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. Computer equipment that does not meet that specification requires additional components, such as a point-of-use switch. During site infrastructure maintenance activity, the risk of disruption may be elevated.
- Concurrent maintainability: Originally, this was also an IT term. It means any work can be performed on a planned basis without impacting the end user. In the site infrastructure world, this means that ANY capacity component or distribution element can be repaired, replaced, serviced, tested, etc., without impacting the computer equipment.

TIER PERFORMANCE STANDARD

Tier I: Basic Site Infrastructure

The fundamental requirement

- A Tier I basic data center has non-redundant capacity components and single non-redundant path distribution paths serving the site’s computer equipment.



The performance confirmation test(s)

- Any capacity component or distribution path failure will impact the computer systems.
- Planned work will require most or all of the systems to be shut down, impacting the computer systems.

The operational impact

- The site is susceptible to disruption from both planned and unplanned activities.
- 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 perform this maintenance work increases the risk of unplanned disruption as well as the severity of the consequential failure.
- Operation errors or spontaneous failures of site infrastructure components will cause a data center disruption.

Tier II: Redundant Capacity Components Site Infrastructure

The fundamental requirement

- A Tier II data center has redundant capacity components and single non-redundant distribution paths serving the site's computer equipment.

The performance confirmation test(s)

- A capacity component failure may impact the computer equipment.
- A distribution path failure will cause the computer equipment to shut down.

The operational impact

- The site is susceptible to disruption from both planned activities and unplanned events.
- Redundant UPS modules and engine generators are required.
- 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 perform this maintenance work increases the risk of unplanned disruption as well as the severity of the consequential failure.
- Operation errors or spontaneous failures of site infrastructure components may cause a data center disruption.

Tier III: Concurrently Maintainable Site Infrastructure

The fundamental requirement

- A concurrently maintainable data center has redundant capacity components and multiple distribution paths

serving the site's computer equipment. Generally, only one distribution path serves the computer equipment at any time.

The performance confirmation test

- Each and every capacity component and element of the distribution paths can be removed from service on a planned basis without causing any of the computer equipment to be shut down.

The operational impact

- The site is susceptible to disruption from unplanned activities.
- 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 Specifications Version 2. This document can be found at http://www.upsite.com/TUIpages/tuifault_spec_2-0.html. 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.
- Operation errors or spontaneous failures of site infrastructure components may cause a data center disruption.

Tier IV: Fault Tolerant Site Infrastructure

The fundamental requirement

- A fault tolerant data center has redundant capacity systems and multiple distribution paths simultaneously serving the site's computer equipment.
- All IT equipment is dual powered and installed properly to be compatible with the topology of the site's architecture.

The performance confirmation test(s)

- A single worst-case failure of any capacity system, capacity component or distribution element will not impact the computer equipment.
- Each and every capacity component and element of the distribution paths must be able to be removed from service on a planned basis without causing any of the computers to be shut down.
- In order to establish fault tolerance and concurrent maintainability of the critical power distribution system between the UPS and the computer equipment, Tier IV sites require all computer hardware have dual power



inputs as defined by the *Institute's* Fault Tolerant Power Compliance Specifications Version 2. This document can be found at http://www.uptime.com/TUIpages/tuifault_spec_2-0.html. Devices such as point-of-use switches must be incorporated for computer equipment that does not meet this specification.

- Complementary systems and distribution paths must be physically separated (compartmentalized) to prevent any single event from impacting both systems or paths simultaneously.

The operational impact

- The site is not susceptible to disruption from a single unplanned worst-case event.
- The site is not susceptible to disruption from any planned work activities.
- 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.
- Operation of the fire alarm, fire suppression, or the emergency power off (EPO) feature may cause a data center disruption.

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. Figure 1 graphically illustrates the tier performance standards. For discussion of the standards, see the following commentary section.

Simply put, the Tier rating for an entire site is limited to 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 is driven by the need to manage perception in senior management, as well as to factually report actual site capabilities. If a site is advertised within an organization as being fault tolerant and concurrently maintainable (Tier IV), it is intolerable to shut the site down at any time in the future—regardless of what subsystem may have required the shut down.

There are no partial or fractional Tier ratings. The site's Tier rating is not the average of the ratings for the 16 critical site infrastructure subsystems. The site's tier rating is the **LOWEST** of the individual subsystem ratings.

Similarly, the "Tier" cannot be imputed by using calculated Mean Time Between Failure (MTBF) component statistical reliability to generate a predictive availability and then using that number to "match" the actual measured availability results shown later in Figure 2. Even if statistically valid component values existed (and they don't because product life cycles are getting shorter and shorter and no independent, industry-wide database exists to collect failures), this approach fails to include people which consistently are involved in 70% of all site failures. A calculated reliability of 0.9999 which ignores human interaction does **NOT** define a site as being Tier IV. The only way to determine Tier Level is to

Figure 1:
Performance Standards by Tier Level

Tier Requirement	Tier I	Tier II	Tier III	Tier IV
Source	System	System	System	System + System
System Component Redundancy	N	N+1	N+1	Minimum of N+1
Distribution Paths	1	1	1 normal and 1 alternate	2 simultaneously active
Compartmentalization	No	No	No	Yes
Concurrently Maintainable	No	No	Yes	Yes
Fault Tolerance (single event)	No	No	No	Yes



objectively determine a site's ability to respond to planned and unplanned events.

TIER COMMENTARY

The *Institute's* STANDARDS Are Outcome Based

The requirements used in the *Institute's* Tier Performance Standard are necessarily and intentionally very broad to allow innovation 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 defining the four Tiers of site infrastructure are very straight forward. Recent initiatives by several groups to replace the *Institute's* Tier concepts with component counts and checklists has lost focus that ultimately counts is uptime performance. Most designs that will pass a checklist approach will absolutely fail a performance requirements approach. What this means is that there is still considerable "art" to the science of uptime and how sub-systems are integrated (or not integrated).

Tier Functionality Progression

Tier I solutions acknowledge the owner/operator's desire for dedicated site infrastructure to support IT systems. Tier I infrastructure provides an improved environment compared to an office setting and includes: a dedicated space for IT systems; a UPS to filter power spikes, sags and momentary outages; dedicated cooling equipment that won't get 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 from site infrastructure equipment failures. The redundant components are typically an extra UPS modules, cooling units, chillers, pumps, and engine generators. Loss of the capacity component may be due malfunction or to normal maintenance.

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 *tactical* solutions, usually driven by first-cost and time-to-market more so than life cycle cost and uptime (or availability) requirements. Rigorous uptime requirements and long-term viability usually lead to the *strategic* solutions found 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 to Tier I and Tier II solutions. Concurrent maintenance means that any component necessary to support the IT processing environment can be maintained without impact on 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 *any* and *every* system or component that supports IT operations must be able to be taken offline for scheduled maintenance without impact on 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, and others.

Tier IV site infrastructure builds on Tier III, adding the concept of fault tolerance to the site infrastructure topology. Just like concurrent maintenance concepts, fault tolerance extends to *any* 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. While the Tier IV definition is limited to consideration of a single system failure, Tier IV requires that the effect of such a failure is considered on other site infrastructure systems and components. For example, the loss of a single switchboard will affect the operation of all the equipment fed from that switchboard: UPS systems, computer room cooling equipment, controls, etc.

The progressive nature of functionality from Tier I through Tier II and Tier III to Tier IV is demonstrated in the schematic illustrations found at the end of this paper. The examples show the addition of components and distribution paths, as described above. Although the illustrations shown are not recommended design solutions for any particular set of requirements, the four electrical topologies are illustrative of the Tier classification concepts. Mechanical system functionally



progresses through the increasing Tiers similarly. 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.

Over the last few years, site infrastructure has been occasionally described by others in the industry in terms of fractional tiers (i.e. Tier 2.5), or incremental Tiers (Tier III +, or Enhanced Tier III, or Tier IV light). Fractional or incremental descriptions for site infrastructure are not appropriate. A site that has an extra UPS module, but needs all the installed computer room air handlers running to keep the UPS room temperature within limits does not meet *site* redundancy requirements for Tier II. A switchboard that cannot be shutdown without affecting more than the redundant number of secondary chilled water pumps is not concurrently maintainable (Tier III).

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 sub-systems. Data centers are dependent upon the successful operation of over 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 perspective owners and designers must consider in making tradeoffs is what impact the decision has on the integrated impact of the site infrastructure on the IT environment in the computer room.

The *Institute* has measured the actual availability, or performance, of 16 data centers having site infrastructure topologies meeting the four Tier definitions and has established availability values representative of each classification. In practice, representative site availability, stated as a percentage of annual operating time, is associated with each of the *Institute's* standard Tier classifications. These empirically determined values include sustainability and human factors over a period of up to 10 years with uptime measured from the perspective of the IT client's operations in the computer room. This "real world" site availability is strikingly different than the probability of system failure that is often calculated 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% (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."

The *Institute* defines site availability from the perspective of a user of IT. Any site incident or event that affects information availability as experienced by end users detracts from site infrastructure availability. The site downtime clock starts running from the moment IT operations were first affected until they are fully restored. Thus, site downtime is not the 15 seconds of a utility power failure, but the total time users were down until IT availability was 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.

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

Tier IV provides robust, fault tolerant site infrastructure, so that facility events affecting the raised floor 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% availability.

The representative availability percentages are a characteristic of the operating experience of multiple sites within each Tier classification. A site with a measured infrastructure availability of 99.90%—midway between Tier II (99.75%) and Tier III (99.98%)—has an operating experience consistent with sites having Tier II topology, but does not achieve the availability of Tier III sites. Availability does 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.

Independent of site infrastructure experience, IT organizations often describe data center availability objectives as Five Nines, or 99.999% of uptime. This is a very aggressive goal, especially if compared to the observed consequences of a single site outage. While the site outage is assumed to be promptly restored (which requires “24 by 7” 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 a data base 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 power outage results in a 4-hour end-user disruption, how relevant is an objective of 99.999% availability? Based on a single site outage of 4 hours, it will take 45.6 years of 100% uptime to restore cumulative site availability back to the 99.999% 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%) uptime. The best a Tier IV site hope for 100% uptime for a string of multiple years. Figure 2 of Typical Tier Attributes uses 99.995% for representative Tier IV site availability, but this assumes a site outage occurs not more than once every 5 years. With a properly designed Tier IV configuration, the single event exposures that can result in a site failure are the results of a fire alarm or the unintended operation of the EPO feature. Only the top 10 percent of Tier IV sites will achieve this

level of performance. Unless human activity issues are continually and rigorously addressed, at least one failure is likely over 5 years.

Typical Tier Attributes

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 1990s. Tier II and Tier III facilities were widespread in the 1980s; Tier III is the most common site infrastructure currently being implemented although most are designed for future transparent upgrade to Tier IV. Most owners find it fairly difficult to upgrade by more than one tier level from what they previously had. 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, then to consistently design, build, and operate the site to conform to those needs.

The following chart (Figure 2) depicts various attributes commonly associated with a particular Tier classification, but the attributes are not requirements of the Tier definitions. For example, the presence of a raised floor or any particular floor height are not criteria for any Tier. (The recommended height of raised floors, when used, is most directly correlated to power density.)

Integration of IT Architecture and Topology with Site Architecture and Topology Helps to Ensure Achieving Uptime Objectives

There are many opportunities within the Information Technology architecture to reduce or minimize the impacts of these unfortunate site infrastructure failures. These steps may include placing the redundant parts of the IT computing infrastructure in compartments served by different site infrastructure systems so that a single event cannot simultaneously affect all IT systems. Another alternative is focusing special effort on business-critical and mission-critical applications so they do not require 4 hours to restore. These operational issues can improve the availability offered by any data center and are particularly important in a “Four Nines” data center housing IT equipment that requires “Five Nines” availability.

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 (Redundant Components) solution, or on two distribution



Figure 2:
Typical Tier Attributes

	Tier 1	Tier II	Tier III	Tier IV
Building Type	Tenant	Tenant	Stand-alone	Stand-alone
Staffing	None	1 Shift	1+Shifts	"24 by Forever"
Useable for Critical Load	100% N	100% N	90% N	90% N
Initial Build-out Gross Watts per Square Foot (W/ft ²) (typical)	20-30	40-50	40-60	50-80
Ultimate Gross W/ft ² (typical)	20-30	40-50	100-150 ^{1,2,3}	150+ ^{1,2}
Class A Uninterruptible Cooling	No	No	Maybe	Yes
Support Space to Raised Floor Ratio	20%	30%	80-90+% ²	100+%
Raised Floor Height (typical)	12"	18"	30-36" ²	30-36" ²
Floor Loading lbs/ft ² (typical)	85	100	150	150+
Utility Voltage (typical)	208, 480	208, 480	12-15 kV ²	12-15 kV ²
Single Points-of-Failure	Many + human error	Many + human error	Some + human error	None + fire and EPO
Annual Site Caused IT Downtime (actual field data)	28.8 hours	22.0 hours	1.6 hours	0.8 hours
Representative Site Availability	99.67%	99.75%	99.98%	99.99%
Typical Months to Implement	3	3-6	15-20	15-20
Year first deployed	1965	1970	1985	1995
Construction Cost (+ 30%) ^{1,2,3,4,5}				
Raised Floor	\$220/ft ²	\$220/ft ²	\$220/ft ²	\$220/ft ²
Useable UPS Output	\$10,000/kW	\$11,000/kW	\$20,000/kW	\$22,000/kW

¹ 100 W/ft² maximum for air-cooling over large areas, water or alternate cooling methods greater than 100 W/ft² (added cost excluded).

² Greater W/ft² densities require greater support space (100% at 100 W/ft² and up to 2 or more times at greater densities), higher raised floor, and, if required over large areas, medium voltage service entrance.

³ Excludes land; unique architectural requirements, permits and other fees; interest; and abnormal civil costs. These can be several million dollars. Assumes minimum of 15,000 ft² of raised floor, architecturally plain, one-story building, with power backbone sized to achieve ultimate capacity with installation of additional components or systems. Make adjustments for NYC, Chicago, and other high cost areas.

⁴ Costs are based on 2005 data. Future year costs should be adjusted using ENR indexes.

⁵ See *Institute White Paper entitled Dollars per kW plus Dollars per Square Foot Is a Better Data Center Cost Model than Dollars per Square Foot Alone* for additional information on this cost model.



paths that may result in a Tier III (Concurrently Maintainable) solution. Compare the Tier II and Tier III diagrams 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 define the Tier III example.

Applying the Standards

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 manufactures—as long as the result can meet the target Tier classification. 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, 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. The increasing power densities of IT equipment required other considerations than the redundancy in the power and cooling systems. Project elements like video surveillance and gaseous fire suppression are frequently necessary to meet an owner's regulatory or insurance requirements, 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.

Consideration of cost, risk tolerance, and Best Practices clearly point to a wider number of site infrastructure characteristics than Tier classification, alone. Experience with the Tier Standard since its inception indicates that Sustainability characteristics become an important factor over time. Investments in Sustainability characteristics account for much of the variance within individual Tier solutions, often leading to increased availability. Typically, Sustainability characteristics decrease the cost or risk of completing maintenance, or

speed the recovery from site infrastructure incidents. 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

Some examples can illustrate site infrastructure characteristics that impact 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 running 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.
- Installing engine generators and switchgear inside the facility (with adequate access space) eliminates the effects of weather and time-of-day on safe maintenance and repair activities.
- In order to improve stability, the combined load on a critical system is often limited to 90% of non-redundant nameplate over a sustained period of time.
- 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") occurred in a Tier III site, the site could be disrupted. However, if the maintenance path is physically separated from the normal path, compartmentalization would permit the site to rapidly recover on a power path through a completely different space than where the fire occurred.



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. Some considerations for selecting an appropriate site infrastructure Tier are:

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 startup companies without quality of service commitments

These companies typically do not have an established revenue stream or 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

- 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”
- Commercial research and development firms, such as software, who do not typically have “on-line” 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 no meaningful impact of disruption due to 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 schedule short periods when limited service is acceptable
- Businesses whose information technology resources support automated business processes, so client impacts of system shutdowns is manageable
- Companies spanning multiple time zones with clients and employees spanning regional areas

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 upgraded to Tier IV.

Tier IV is justified most often for

- Companies with an international market presence delivering 24x365 services in a highly competitive client-facing market space
- Businesses based on E-commerce, market transactions, or financial settlement processes
- Large, global companies spanning multiple time zones where client access to applications and employee exploitation of information technology is a competitive advantage

Companies who have extremely high-availability requirements for ongoing business, or for whom there is a profound 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.



Making the Appropriate Tier Selection Should Be Based on Business Requirements

Selecting the site infrastructure solution based on the availability objectives required to sustain well-defined business processes with substantial financial consequences for downtime provides the best foundation for investment in data center facilities. The owners' focus during the data center design and delivery process should be the consistent application of the Tier Performance Standard, rather than allowing recurring debate over every characteristic or attribute that makes up the data center's site infrastructure.

Including criteria from a higher Tier classification, or an attribute leading to increased availability, does not increase the overall Tier classification. Moreover, deviation from the Tier standard in any subsystem will prevent a site from classification at that Tier. For example, a UPS system patterned after a Tier IV system within a site featuring 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 utilize a Tier II mechanical system that cannot be maintained without interrupting computer room operations. This results in the overall site achieving a Tier II rating. Most often the mechanical system fails concurrent maintenance criteria because of inadequate isolation valves in the chilled water distribution path.

Another common oversight is the effect of shutting down electrical panels on the mechanical system the panel feeds. If more than the redundant number of chillers, towers, or pumps is de-energized for electrical maintenance, computer room cooling is impacted.

Occasionally, electrical systems fail to achieve Tier III or Tier IV criteria due to the UPS power distribution path. Topologies that include static transfer switches that cannot be maintained without affecting computer room power, fail the concurrent maintenance criteria. UPS configurations that utilize common input or output switchgear are almost always often unmaintainable without computer room outages and fail the Tier III requirements even after spending many hundreds of thousands of dollars.

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 Figure 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 effect on IT availability. A well-executed data center master plan or strategy should consistently resolve the entire spectrum of IT and facility requirements.

Figure 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 Generators Dual Power S + S
Mechanical Infrastructure	Redundant Components Fans and Pumps on UPS	Thermal Storage	Dual Pipe Thermal Storage
Facility Operations	Passive Automation Change Management MAPS/Certification Simulation	24 by "Forever" Staffing Compartmentalization Failure Bypass Options On-Site Spares	Work Performed during Regular Hours In-House Knowledge In-House Supervision



Lifecycle Planning

It is disappointing to observe brand new sites that received very little thought during initial design to future operations. Valves were located in inaccessible places, the access path for the addition of future components was not thought out, or sufficient capacity to simultaneously test new systems while sustaining the critical load was not provided. These details could have been addressed for no additional cost during design. This failure limits both investment value and site performance right from its initial occupancy. A more sustainable site will project future requirements and anticipate them during the initial design and construction.

Sites should be designed to anticipate increasing power requirements and tier levels. These sites provide future locations for necessary site infrastructure equipment as well as a planned means to commission them and then connect them transparently to operational systems.

Institute Site Topology and Sustainability Certification

The *Institute* exclusively reserves the right to determine Tier ranking and to certify sites as meeting Tier requirements as summarily described in this white paper. This comprehensive process involves additional criteria beyond the information provided herein. The process is similar to that used for ISO 900X certification. The ISO standard is set and maintained by the International Standards Organization who trains and certifies field inspection agencies in different parts of the world. These field inspectors inspect and validate conformance to the ISO standard before certification is granted for a limited time period. The *Institute* has licensed ComputerSite Engineering Inc., a separate but related company, to perform inspection and validation utilizing the *Institute's* Tier Performance Standards and the *Institute's* comprehensive database of emerging industry problems and best design practices. Sites reviewed and certified by the *Institute* can be seen at www.uptimeinstitute.org/tui_certification.html.

Conclusion

Data center owners have the responsibility to determine 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 should be uniformly implemented.

About the Authors

Mr. Turner is a Distinguished Fellow and Senior Certification Authority for the *Institute* and a Principal of ComputerSite Engineering, Inc. in Santa Fe, NM.

Mr. Seader is a Distinguished Fellow and Certification Authority for the *Institute* and a Principal of ComputerSite Engineering, Inc. in Santa Fe, NM.

Mr. Brill is the founder of the *Institute* and is its Executive Director. He is a Principal of ComputerSite Engineering.

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 85 members of the *Institute's* Site Uptime® Network are committed to achieving the highest levels of availability with many being Fortune 100 companies. They interactively learn from each other as well as from *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. The *Institute* also conducts sponsored research and offers insightful seminars and training in site infrastructure management.

© 1996, 2001-2006 The Uptime Institute, Inc.



Building 100
2904 Rodeo Park Drive East • Santa Fe, NM 87505
Fax (505) 982-8484 • Phone (505) 986-3900
tui@uptimeinstitute.org • www.uptimeinstitute.org


Illustrative Electrical System Topology - Tier I



Note: This diagram illustrates basic Tier I electrical distribution concepts. This diagram shall not be interpreted to represent a standard or compliant electrical system topology, or a solution fulfilling any particular set of requirements.

Site certification requires consistent application of Tier concepts to all 16 critical subsystems that comprise data center site infrastructure.


Illustrative Electrical System Topology - Tier II



Note: This diagram illustrates a basic Tier II electrical distribution concept. This diagram shall not be interpreted to represent a standard or compliant electrical system topology, or a solution fulfilling any particular set of requirements.

Site certification requires consistent application of Tier concepts to all 16 critical subsystems that comprise data center site infrastructure.


Illustrative Electrical System Topology - Tier III



Note: This diagram illustrates a Tier III electrical distribution concept. This diagram shall not be interpreted to represent a standard or compliant electrical system topology, or a solution fulfilling any particular set of requirements.

Site certification requires consistent application of Tier concepts to all 16 critical subsystems that comprise data center site infrastructure.

Illustrative Electrical System Topology - Tier IV



Note: This diagram illustrates a Tier IV electrical distribution concept. This diagram shall not be interpreted to represent a standard or compliant electrical system topology, or a solution fulfilling any particular set of requirements.

Site certification requires consistent application of Tier concepts to all 16 critical subsystems that comprise data center site infrastructure.