Global Data Center Engineering

GDCE White Paper Series:

Dirty Little Secret of Data Center Standards – A 20 Year History

The data center industry and its core users have become obsessed with "Design Standards". Globally, there are many standards, some of which have come and gone, but the three discussed here are Uptime Institute, TIA, and BICSI. So, first a little recap, and history.

Uptime Institute (UTI):

In the early 1990's Uptime Institute began a study to understand data center performance differentiators. This ultimately led to their publishing of the white paper "Tier Classifications Define Site Infrastructure Performance" (Which later became Tier Standard: Topology). As time progressed, Uptime Institute realized that just designing and building a data center to a Tier standard, was a "one-time event", and did little to predict performance over the life of the data center, and particularly how the data center was maintained. In July of 2010, they released a new standard to address this issue: "Tier Standard: Operational Sustainability".

Primary focus: Power and Cooling (Critical Systems Infrastructure) (Tier Standard: Topology)

Secondary focus: Data Center Operations (Tier Standard: Operational Sustainability

UTI First Release: 1995 "Tier Classifications Define Site Infrastructure Performance"

UTI Updated: 1997

UTI Updated: 2001

UTI Updated: 2005 - Updated, and Name Changed to: "Tier Standard: Topology" (TS:T)

UTI Updated: 2008

UTI Updated: 2010 - "Tier Standard: Topology" (Major design change enabling zero utility power feed for all Tier levels) also introduced "Tier Standard: Operational Sustainability" (TS:OS) as operations standard

UTI Updated: 2012 - No major change, restructure of TS:T to better align terminology and integration of TS:T with TS:OS UTI Updated: 2013 - TS:OS changes removing bonus categories for operational performance criteria (made them mandatory instead for Tier III and IV). Other minor changes, no major modification or directional/policy change.

Telecommunications Infrastructure Association (TIA)

In 2001, a committee was formed by TIA to address the importance of data center cabling standards. TIA looked out into the data center thought leaders and struck up a friendly relationship with Uptime Institute. The intention was to expand on the UTI's works and describe data center cabling (Structured Cabling) standards to provide further depth to data center standards. Part of that process included adopting the UTI Tier definitions, but to differentiate the two, where UTI utilized roman numerals to describe the various "Tier", TIA adopted Arabic numerals. (So UTI is Tier I - IV, where TIA is Tier 1 - 4). In addition, TIA added beyond the "Power and Cooling" (or Mechanical and Electrical) two additional disciplines: Architectural and Telecommunications.

With permission from UTI, they adopted the overall definition of the four Tiers: Basic Infrastructure (Tier I/1), Redundant Capacity Components (Tier II/2), Concurrently Maintainable (Tier III/3) and Fault Tolerant (Tier IV/4). The friendly partnership was very short-lived, and as TIA was based on a standards committee, many committee members didn't agree with the approach that UTI had taken for establishing these definitions. In 2005, TIA released their first version of the "Telecommunications Infrastructure Standard for Data Centers" or what is better known as "TIA-942".

TIA-942 as a standard was widely and rapidly accepted by many in the data center industry. The variation between UTI and TIA in the first release was somewhat significant, with TIA adding extensive "rules" around the mechanical and electrical facilities (some in conflict with themselves, but it was a first version of the standard). UTI objected strongly to TIA's application of their terms and definitions, and the relationship quickly deteriorated. While the standards "main body" chapters focus on cabling architecture, and some energy efficiency practices, what has become known as the "TIA-942 Standard" includes the Annex chapters, where the definition of the standards other critical systems infrastructure, architectural and telecommunications are described in detail (particularly Annex G in early editions, and Annex F in later versions), which provide an extensive (though not meant to be comprehensive) list of requirements that define the various Tiers (Later changed to "Rated" in the 2014 and later releases of TIA-942 as a result of the ongoing dispute with Uptime Institute which was settled in 2014).

Primary focus: Data Center Cabling and Space Planning

Secondary focus: Data Center Critical Systems Infrastructure, Architectural, Site Selection

TIA-942 First Release: 2005 "Telecommunications Infrastructure Standard for Data Centers"

TIA-942-1: 2008 Some changes to acceptable cabling standards for coaxial cable specification and usage, no major changes

TIA-942-2: 2010 Major changes to the "core" of the standard, removing dozens of elements from "Annex G", and adding a few others across all four areas (a messy release as only "changes" were grouped into a small document, resulting in the necessity to have a copy of the 2005 standard and the 2010 standard to actually have a full picture of the standard).

TIA-942-A: 2012 - Incorporated changes form TIA-942, TIA-942-1 and TIA-942-2 into one standard to consolidate the other messy addendum.

TIA-942-A1: 2014 - After a sternly worded letter from UTI insisting that TIA stop using UTI and TS:T as a "base standard" TIA releases a new version changing the word "Tier" to "Rated"

TIA-942-B: 2017 - Further refinement in what's in and what's out. A return of the "fuel" supply runtimes of 8, 24, 72, and 96 hours for the 1 to 4 ratings.

The 2010 release is really where the major "break" occurs in terms of philosophy with Uptime Institute. UTI up to that point had maintained the requirement for utility power supply as the

"Primary Source of Power" on a site, with generators serving only as a "backup". However, at 2010, UTI changed that position, and required for Tier III and IV that continuous run (or de-rated Prime) generators are the primary source of power on site, allowing for 0 utility power feeds (for any Tier) so long as the power generation can be maintained as a constant, and concurrently maintainable for Tiers III and IV.

TIA-942 instead regards utility power as the primary source of power for the site, with generators as a backup (requiring redundancy in the generators sets for Rated 3 & 4 sites). The 2014 release's main change was to drop the term "Tier" and TIA replaced it with "Rated". It also added Data Center fabric as an addendum.

Building Industry Consulting Service International (BICSI)

Recognizing ever growing conflict and market requests for a Data Center standard that aligned to other BICSI standards (a long established standards body within infrastructure and building practitioners), BICSI began to work on their own Data Center standard as well. Their first release was a rush to market in 2010 and has a number of serious shortfalls, (including the absence of an Electrical chapter, or any suggestion around electrical design). While TIA and UTI both started off with "Tier" as their defining term, BICSI described theirs in terms of "Facility Class" or "FC" or in some cases just "C". Having 5 "levels" instead of 4 as UTI and TIA had, their 5th level is actually a "Zero" class, meaning, it doesn't really qualify as a data center at all but deserves some special consideration when designing and implementing than just standard office space. Again, at the core of the definitions of the FC, are the four standard definitions of what each of these represents: Basic Infrastructure, Redundant Components, Concurrent Maintainability, and Fault Tolerance. The BICSI 2010 release reads very much like the TIA-942 2005 release, but instead of "annexing" the definitions of data centers, it embraced them as a full part of the standard (for better or for worse).

BICSI sought as well to expand even further on the standard than TIA had gone. They also added significant depth on site selection and commissioning, as well as some operational considerations. It's clear their intent was to be an "All-embracing data center design and operation best practice guide". It certainly failed to achieve that in the 2010 release, but subsequent releases have improved and expanded upon it. (In fact, the BICSI standard seems to have bloated considerably. The 2010 release was just over 200 pages, 2011 release was 391 pages, and the 2015 release is a whopping 533 pages). It has removed maps and references to the US only, and replaced them with more global data, which is a big improvement in the overall standard, and clear it is attempting to obtain a greater foothold in the global marketplace.

(Called "BICSI 002-2014" the standard was released for sale to the public on January 7, 2015).

Primary focus: Data Center Design, Site Selection

Secondary focus: Operations, Commissioning, Maintenance

BICSI 002-2010: First Release: "Data Center Design and Implementation Best Practices" - an abomination, incomplete, and missing entire chapters

BICSI 002-2011: Second release with all sections complete

BICSI 002-2014: Third release, significant expansion, and new chapter additions

Now: The Dirty Little Secret You've been waiting for...

Take all of these data center standards into consideration. By definition, their "Core" values for each "Level" are the same.

Level 1 (Tier I, Rated-1, Facility Class 1): Basic Components (or just enough to run the site, "N" capacity). Any failure means downstream loss. One thing these all have in common: They all require at least 1 generator for backup. Single pathway for all systems. (Power, water, fuel, cooling, etc.)

Level 2 (Tier II, Rated-2, Facility Class 2): Redundant Components (the main components require redundancy, for both mechanical and electrical systems) (there's an unusual caveat here for TIA, however, as they only require 1 Generator, because by "definition" the generator is a "backup capacity" to the utility, which is primary, and is therefore not a "capacity component" by definition). Single pathway for all systems.

Level 3 (Tier III, Rated-3, Facility Class 3): Concurrently Maintainable. The philosophy here is, that any component or pathway can be planned to be taken out of service, and the site continues to operate without impact to "Production IT Systems". A single component failure will not result in loss of critical load. In most cases, this means at least N+1 for all capacity components, and a minimum of dual diverse redundant pathways for all systems, though one may be active, and the other "standby" with manual changeover between the two allowed. Remember the objective is concurrent maintainability, NOT "survive any failure (particularly of the pathway)".

Level 4 (Tier IV, Rated-4, Facility Class 4): Fault Tolerant. Each accepts that any one component can fail within the system, and the site will not suffer a loss of "Production IT Systems". This is where the greatest "variation" between the standards exists. They are very complex implementations, some requiring automation, some extensive redundancy (2(N+1)), all require compartmentalization at this level, though the implementation of compartmentalization varies in each of the standards. TS:T focuses on redundant components to isolated space (no more than R = Redundancy in any compartment). While TIA uses a more subtle definition, no mix of system A or system B within the same compartment. This includes pathways as well.

The real dirty secret though is this: Take all you have just read, the entire history of the evolution of **published data center standards** from 1995 to 2016, and the **AVAILABILITY of each of the "Levels" across ALL of the standards have not changed!** I really cannot stress that point enough.

Yes, you read that correctly. The availability of a "Level 1" data center in 1995 is the same as it is in 2016. Regardless of which standard you use. So FC-1 has the same availability as Rated-1 has the same availability as Tier I. The availability of a "Level 3" data center in 1995 is the same as it is in 2016! (et al You get the point). There is ZERO evidence to support that the availability at any of these levels has improved (and in fact, all but BICSI have dropped any "claim" to DC availability % as expressed by the "Five 9's").

What has happened is the cost to implement across each of the various "Levels" has decreased. Each of the standards has recognized that creating stranded capacity is just wasted cost, that does nothing to improve the availability.

Mixing of standards is an even bigger hazard because they lose cohesion. The advice, in this case, is: Pick a standard and stick to it. Don't design to one and certify to another. They have different ways of achieving the same outcome. If you're heavily invested in BICSI, then BICSI may be the best choice for you. If you've been engaged with TIA for years (in other areas) then go with TIA. Of all of the standards, Uptime Institute will require you to "look beyond" to other standards to "complete" the areas they do not cover. And of course, above all, **do not forget LOCAL CODE**. Regardless of the standards mentioned above (all of which are voluntary) local codes are not voluntary. You must adhere to them, as a minimum. If the standard you have chosen exceeds the local codes, (in most cases), then you are fine. When in doubt, consult with a knowledgeable expert on the local code compliance (also known as "Authority Having Jurisdiction" or "AHJ").

Why Isn't EN 50600 included in this discussion...

Two reasons. First and foremost, it's still incomplete. So without it being a completed standard, I don't wish to comment on what is is, is not, or may be.

Second, and just as important, it applies just the same... They are still dividing it into 4 "levels" and there will be reasons to justify their definition as well. From what I have seen of this standard it is 90% TIA-942 from 2010 edition with some "Europe Specific" elements included, which is really not necessary. (Not necessary in that you can just as easily write them to apply anywhere, and why the compulsion to exclude is at least unwarranted and at best, unclear.

In Closing

Are data center certifications useful? Yes.
Why? Because they show that some guided process has been followed to improve the probability of desired outcomes.
Which one should you use? The one your data center was designed for. If it wasn't designed using a standard, then the one it most closely matches.

DC design standard changes have not reduced the probability or frequency of outages in their 20-year history (at their individual levels). However, the elements required to meet their objectives have been challenged, and that has resulted in lower cost to implement (removing stranded capacity in the design requirements, removal of some incremental "bigger is better", more flexibility in new technologies so there is less need for the standard to "catch up" to current innovation).

Don't forget Reliability Engineering. The key principle to remember is reliability engineering teaches us: Bigger and Fewer. That translates to, use the fewest number of components possible to provide the level of redundancy you need. Ideally, that would always be 2 components, but the world doesn't always work that way. Understand that in many cases if you have 2N system, you have a higher probability of a component failure than if you have N+1, particularly when exceeding 2 components. It's simple math, if you have 2 10MW transformers (N+1) or 6 2MW transformers (N+1 at the same 10MW), your likelihood of a failure in the 6 versus 2 is 3 times higher.
Use the standards wisely, and know why you are using any of these (or other DC standards, including corporate standards), and remember to design for reliability and redundancy that is justified by the business case.

Biography: Scott Payton is technical director for Global Data Center Engineering. He has a critical systems background spanning more than 30 years including engineering in ICBM Missile silos with the US Air Force. He holds a Master of Science in Information Systems Management, and is currently a doctoral candidate, where his thesis topic is Human Error Management in the Data Center. Scott has previously worked for Dell as well as other data center consulting firms, and is a frequent speaker and writer on a range of data center topics including Computational Fluid Dynamics (CFD), Data Center certification, and Data Center Design and Operations standards.

Acknowledgments:

GDCE would like to graciously acknowledge the following people who had contributory input into this white paper, making it exceptional instead of ordinary: All those who commented on this living article from LinkedIn over the past several months.