TRG Datacenters

Datacenter Infrastructure for HPC: Site selection & special operational considerations.

Also a post mortem on Hurricane Harvey.

A brief overview of electrical design in datacenters

- Industry standard is tiering system
  - Evaluates many points of design in facilities to include concurrent maintainability, fault tolerance, physical isolation, and others
  - Nuanced evaluation process, we will discuss at a high level in a simplified format
- This is managed by a private management consultant company
  - Many datacenters endeavor to follow these guidelines, some do, and quite a few have no idea what they are talking about!

Uptime® Institute

https://uptimeinstitute.com/

Datacenters are complex systems in terms of probability of failure from a design perspective

- Adding components in series increases complexity and probability of failure
  - The least reliable component in series generally is your “bottleneck” mathematically
- Adding components in parallel reduces probability of failure drastically, for that component

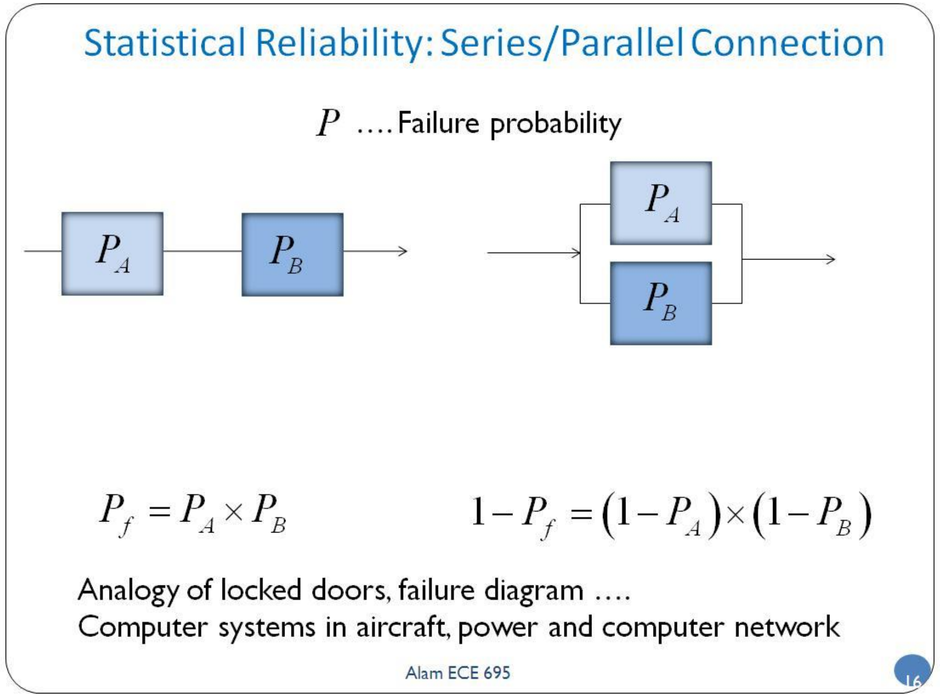

https://nanohub.org/app/site/resources/2013/01/16583/slides/016.03.jpg

"N" Design

- Most often Tier II
  - Provides single source backup and UPS
  - Mechanical Distribution is often not backed up by UPS
  - Overhead is calculated as 1x

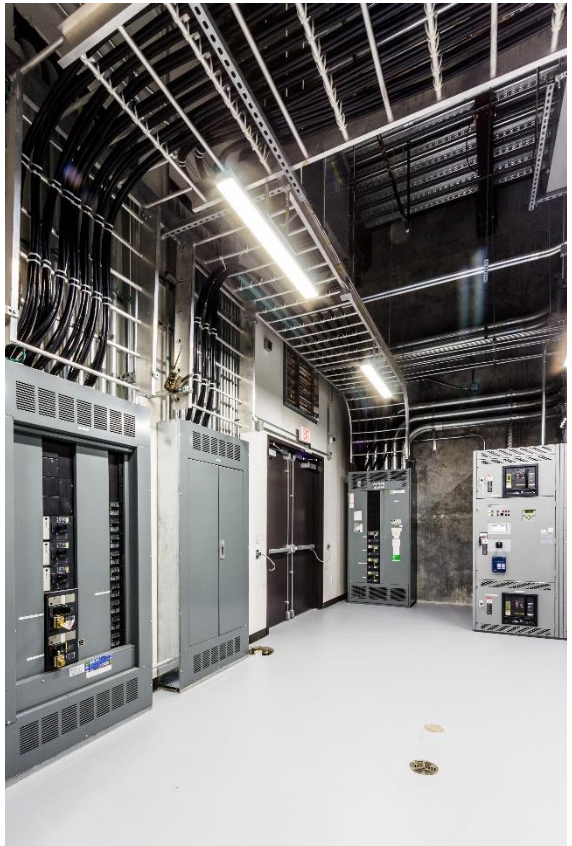

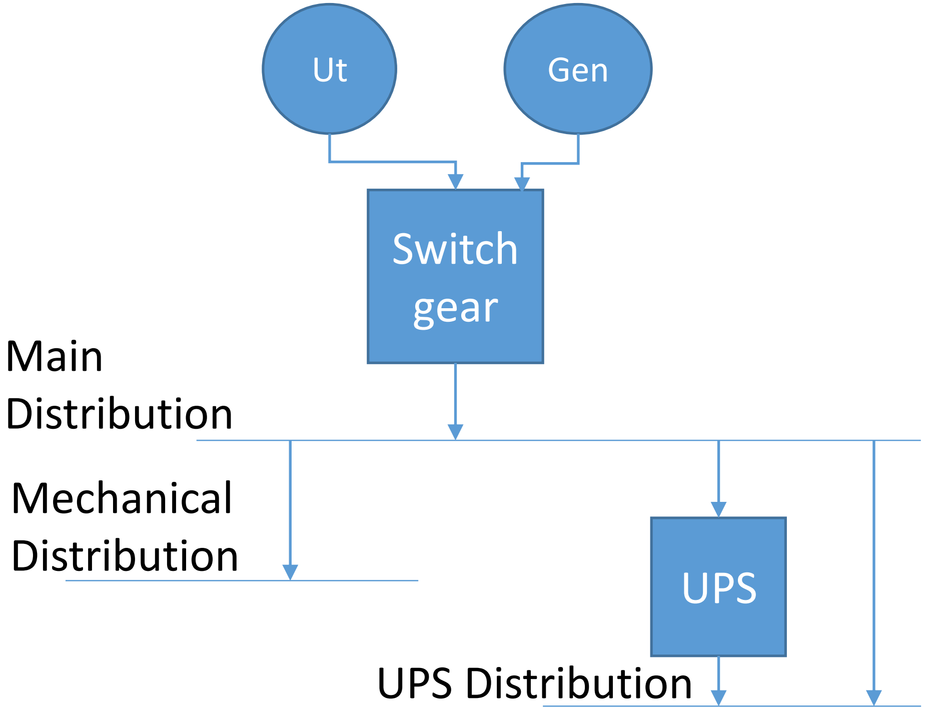

Flavors Of “N+1” (Typical Tier III Attempts)

- Not all designs are created equal
- The worst Ive seen: “Dual Corded” Feeds
- Who sees what is wrong here?

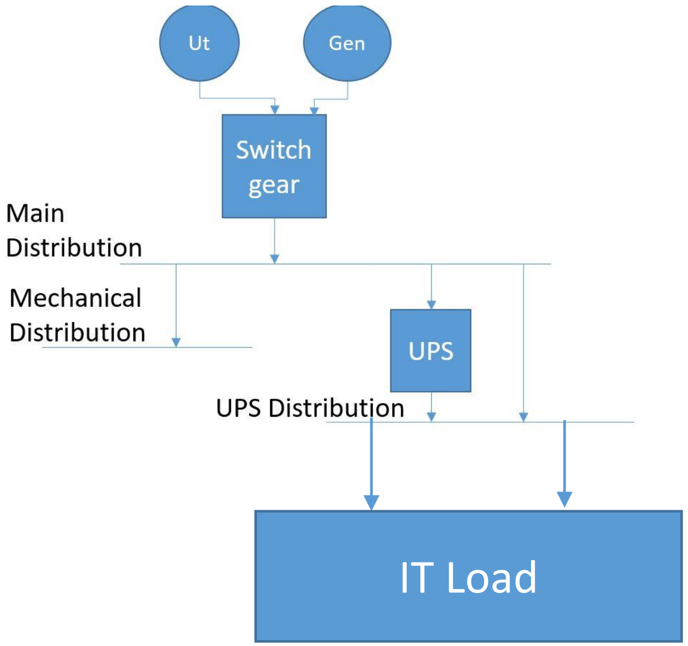

TRG Datacenters

Other "Flavors" Of N+1 (Typical Tier III Attempts)

- "Redundant UPS"
- "Redundant Generator"
- System reliability hovers as reliability of Switchgear, Mechanical Distribution Breaker, Etc.

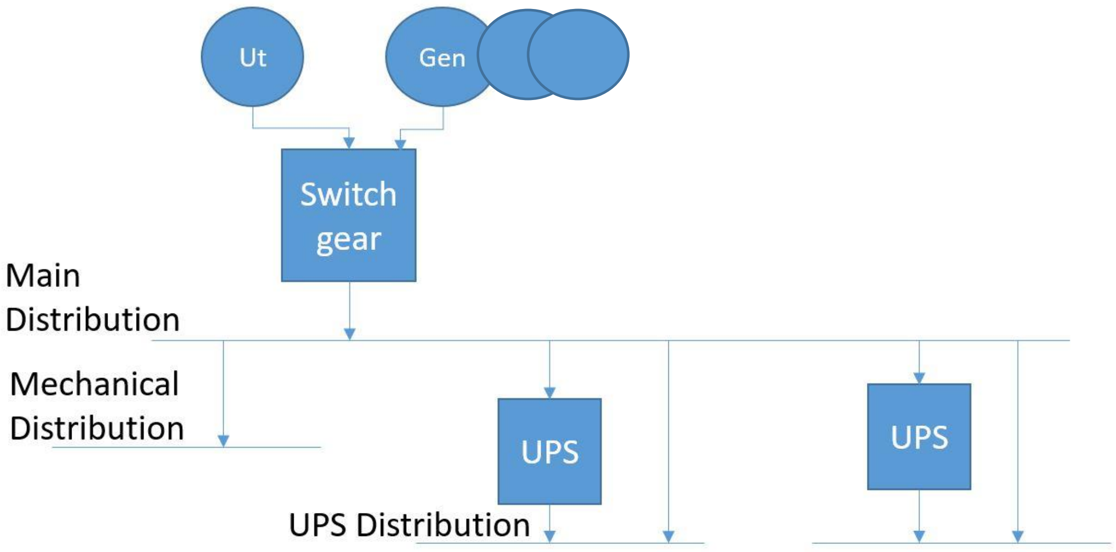

True "N+1"

- Often lacks physical / electrical isolation required for Tier IV, Typically Tier III
- Catcher systems also utilized with static transition switches to achieve Tier "N+1"
- Ranges from 1.33x to 1.5x Overhead

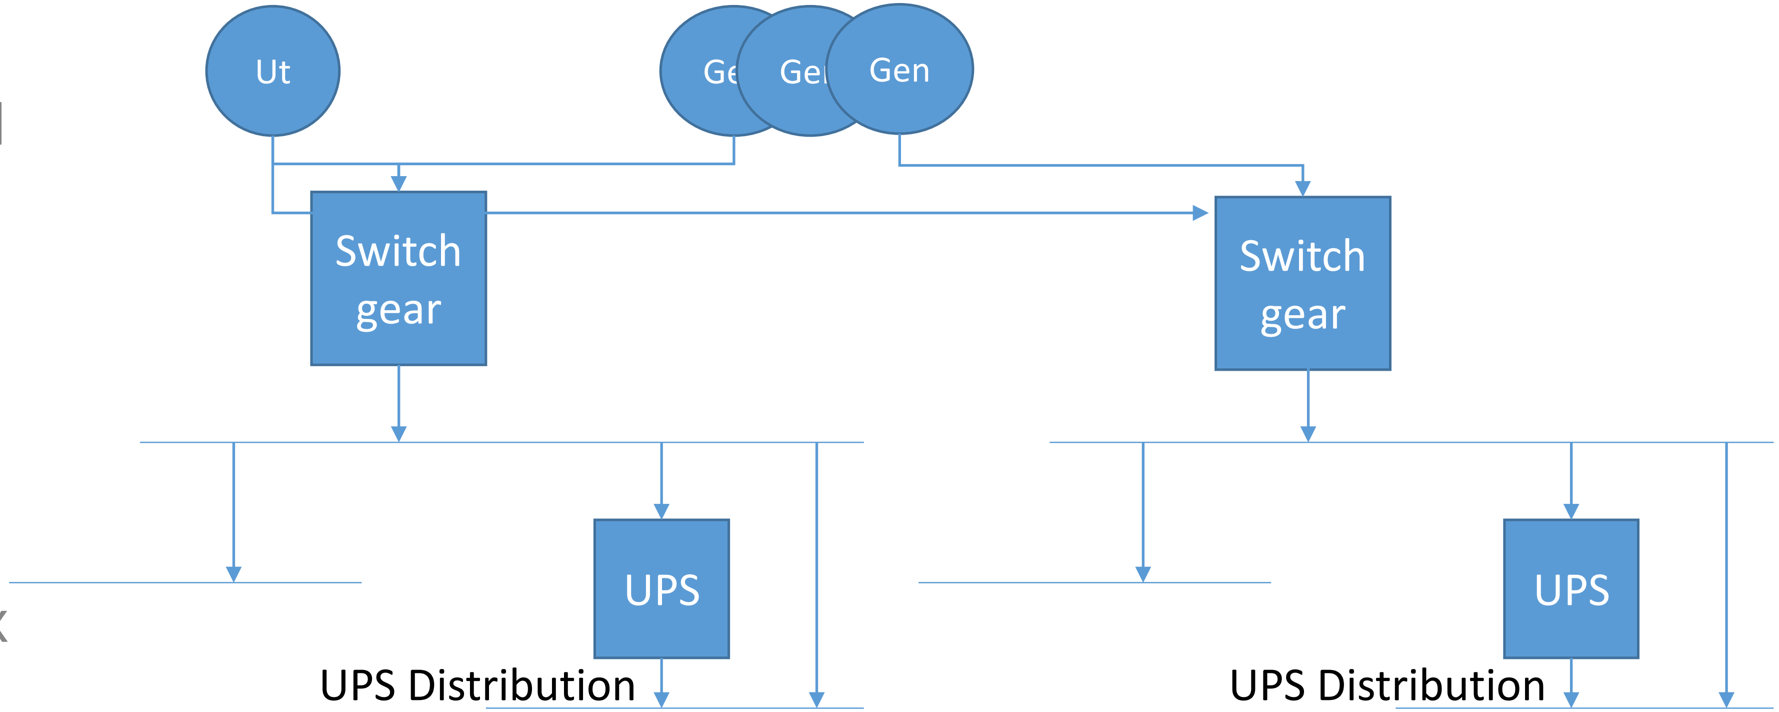

Pros and Cons of N+1 Reliability

- Typical Equipment Overhead: 1.33x – 1.5x (Where N in 3 and 2 respectively)
- Typically best achieved in larger facilities due to cost of modularity
- Cost of modularity: Example – 3+1 for a 600 kVA system means installing 4x separate utility systems that are 200 kVA each.
  - Labor Cost Is Increased Significantly for conduit and wiring
  - Smaller units typically cost more from a $/kVA perspective
- Easy to screw up somewhere!
- For smaller facilities, a 2N infrastructure may approach cost parity of properly achieving N+1

“2N, The Ultimate?”

- Not necessarily, but many times yes.
- Pros:
  - Simple to layout
  - Redundant and reliable
- Cons:
  - 2x Equipment overhead/cost

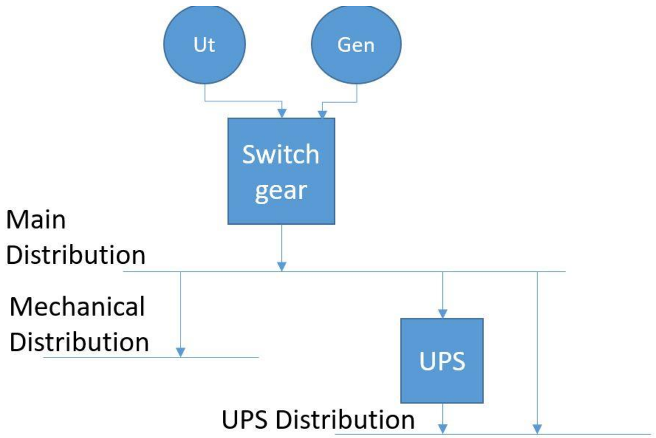

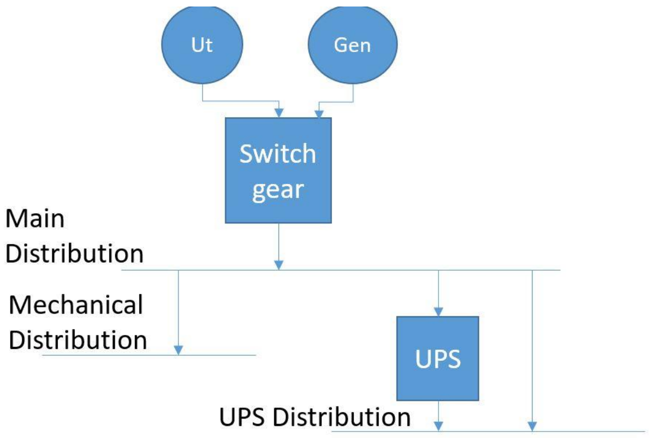

Distributed Redundancy

- 4N/3 and 3N/2 Popular
- Combinatorically combine load in a 2N fashion from multiple sources
  - 4N/3 Has 6 Combinations
  - 3N/2 Has 3 Combinations

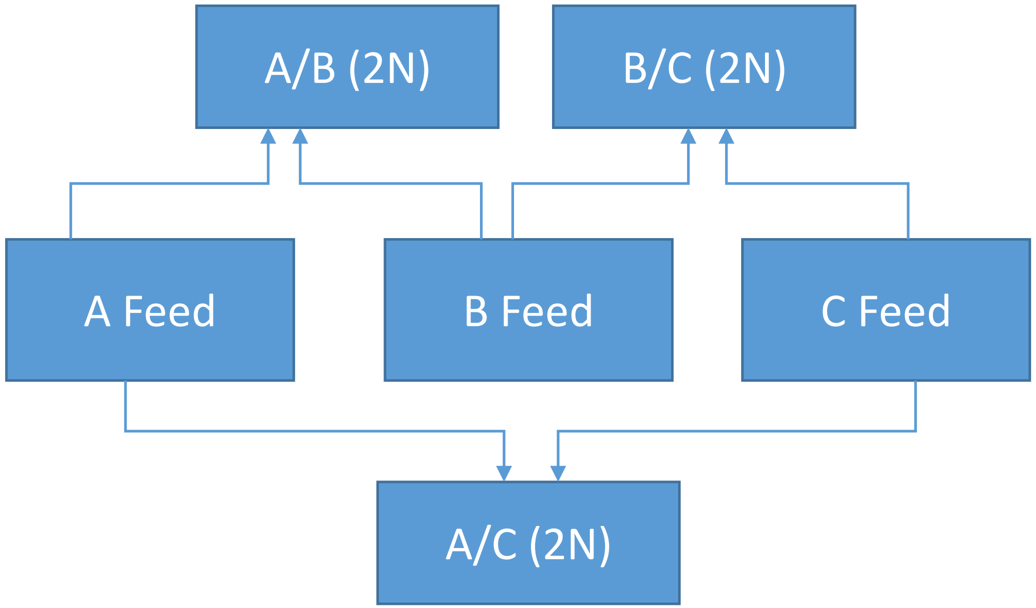

TRG Datacenters

Pros and Cons of Distributed Redundancy

- Cons
  - Complex planning required – it gets complicated fast!
  - Same challenges of cost of modularity as N+1
- Pros
  - Approaching 2N level of reliability at N+1 cost for proper scale
  - Ability to provision N level and 2N redundancy within the same space without any wasted overhead
  - Ability to do layered Distributed redundancy on a larger scale (4N/3 x4 + 1 for example) that can achieve 2N levels of redundancy at N+1 cost

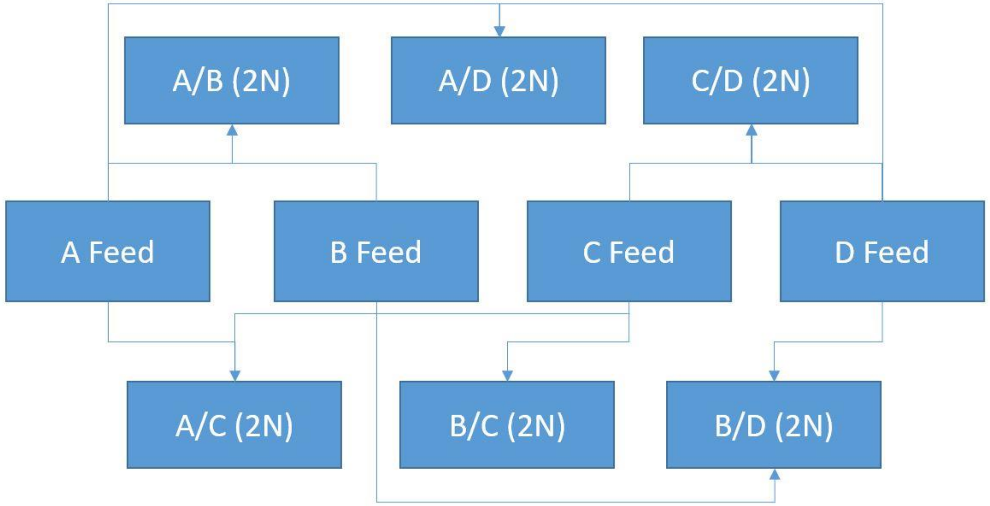

TRG Datacenters

Summary of Electrical Design, and How It Applies To HPC

- HPC typically fits a mixed resiliency model, where storage networking and other applications are best used with 2N redundancy, while batch based processing components are often best suited for N levels of redundancy
- 2N (for smaller designs), and Distributed Redundancy at scale is recommended to best suite the ability to provide N and 2N as necessary
  - Note: 2N is technically distributed redundancy (2N/1)

HVAC and Reliability Considerations for HPC

- “Where the rubber hits the road”
  - High density and thermal load requires us to think about transient failure conditions
  - Often times, the electrical infrastructure is not properly married to the mechanical infrastructure to account for operating under failure conditions as well as transient conditions

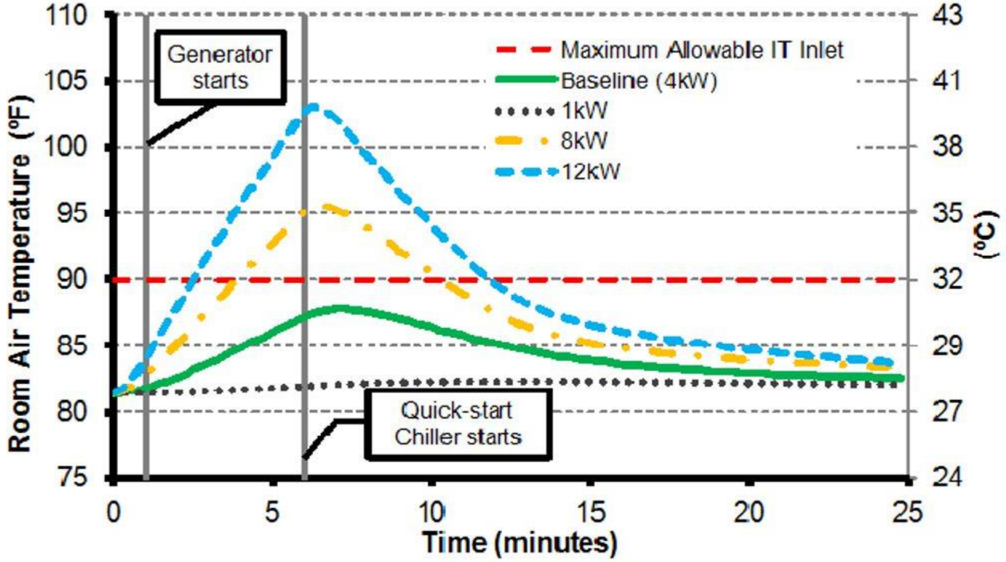

www.apc.com/salestools/DBOY-7CDJNW/DBOY-7CDJNW_R1_EN.pdf

TRG Datacenters

High Level Types of Cooling

- Chilled Water Distribution
  - Used for in rack/HPC/liquid cooling applications often
  - Piping adds complexity fast, especially with redundancy requirements (often relegated to 2N)
  - Restart time on chillers often very long, must supplement with UPS backup to chillers, or chilled water reservoirs
- DX Based
  - Very easy, modular, and scalable
  - Can only provide cooling at a room level, not suited for liquid cooling
- Air Economization
  - Trim cooling often needed in “challenging environments”

Thinking about the worst case scenario

- The best thing to do is ask, under a full utility failure with full loading on the hottest day of the year, what happens to my HVAC?
  - With chiller systems, also ask under a total pipe failure, what happens.
  - Also need to account for the effect of maintenance during a failure
- We implement distributed redundancy across 4x utilities where under worst case scenario we are at N+2 redundancy (to account for maintenance during these situations, concurrent maintainability)
- This is an especially important question to ask in mixed resiliency environments, making sure your failure modes of lower resiliency systems do not impact your higher resiliency systems

For Traditional Systems (Air Cooled), Large Plenums and Hot Air Containment Do Well

- Allows for ambient room to feel comfortable with a high deltaT of of systems for the hot air containment areas
- With proper hot air containment, we are able to think about cooling in terms of CFM in and CFM out

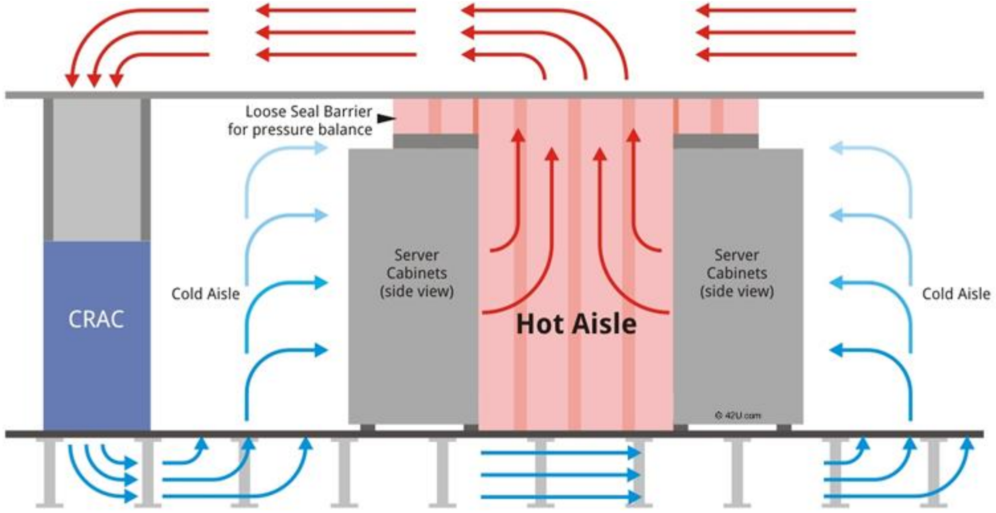

TRG Datacenters

Moving a system into a datacenter

- Technical:
  - System weight
    - What floor is the facility on
    - Onloading/Offloading considerations, proximity of loading dock to datacenter floor
    - Weight load capacity of flooring
  - Environmental controls, especially chilled water temperature if doing in rack cooling
  - The roofing system, and leak protection
  - Properly scoping resiliency
    - Batch based processing, long term outages, critical components
  - Resiliency and marriage of mechanical/electrical
  - Survivability
    - What is the failure phenomenology of your region?
    - What type of failures are you planning for?

Roofing Systems Risk profile

- Age of roof
- Slope of roof
  - Flat roof with tapered insulation and roof drains? Plan for ponding and eventual leaks
- Number of penetrations in the roof
- Redundancy in roof, what happens if an unplanned penetration occurs?
- Wind load rating of roof
- Method of removing water from roof
- Single slope roofs with high wind load rating, 2 layers of leak protection, and zero roof penetrations expected to perform the best

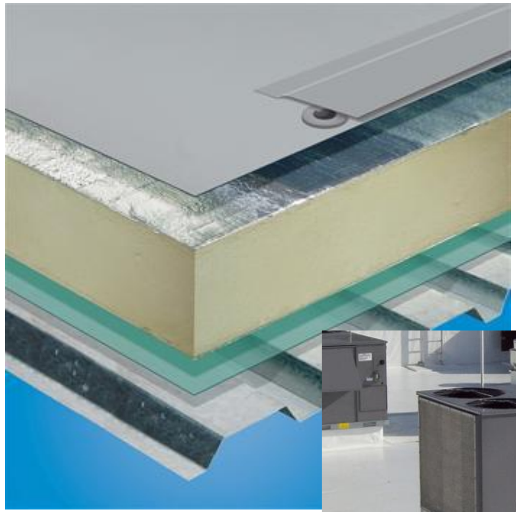

Hurricane Harvey, and Site Selection In General

- Disaster Phenomenology Profile For Houston
  - High Wind Load
    - Long term utility outages
    - Facility failures for equipment and buildings
  - Flooding
    - Limited accessibility/Resupply
    - Substation flooding
  - Trains and Hazmat transport (Street, Pipeline, and Rail)
    - Mandatory Evacuation
  - Construction (buried cable cuts)
  - Temperature
- Considering all these factors can be challenging for an individual, utilizing expert help can often be the best solution for evaluating a site
  - Example: CRESA MCS iSiteSelection tool

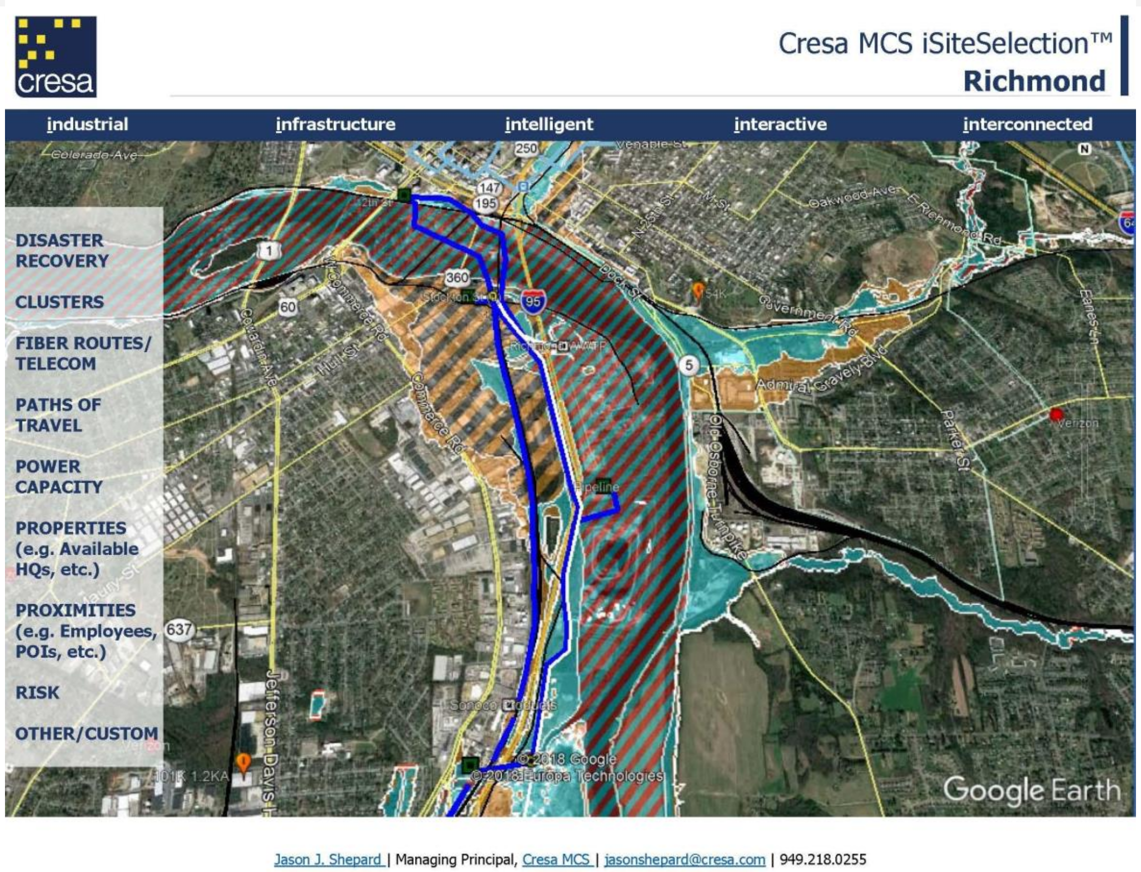

Flooding

- Checking flood maps for your particular property is a great place to start
  - Testing route traversability to fuel supply depots during major storms is also important
  - Checking accessibility from your office or home can also be important, or asking questions about remote hands services in the event of an emergency
  - http://www.harriscountyfemt.org/
- To deal with mitigating resupply requirements, and increase off the grid “survivability”, bi-fuel conversion on generators run with up to 66% natural gas,
  - Increase your onsite diesel run time by 3x
  - Reduce your resupply hourly requirements by 3x
- Substation flooding
  - Substation standard design for coastal regions is above 100 year floodplain. Not always helpful for Houston. Checking and driving your local substations is helpful

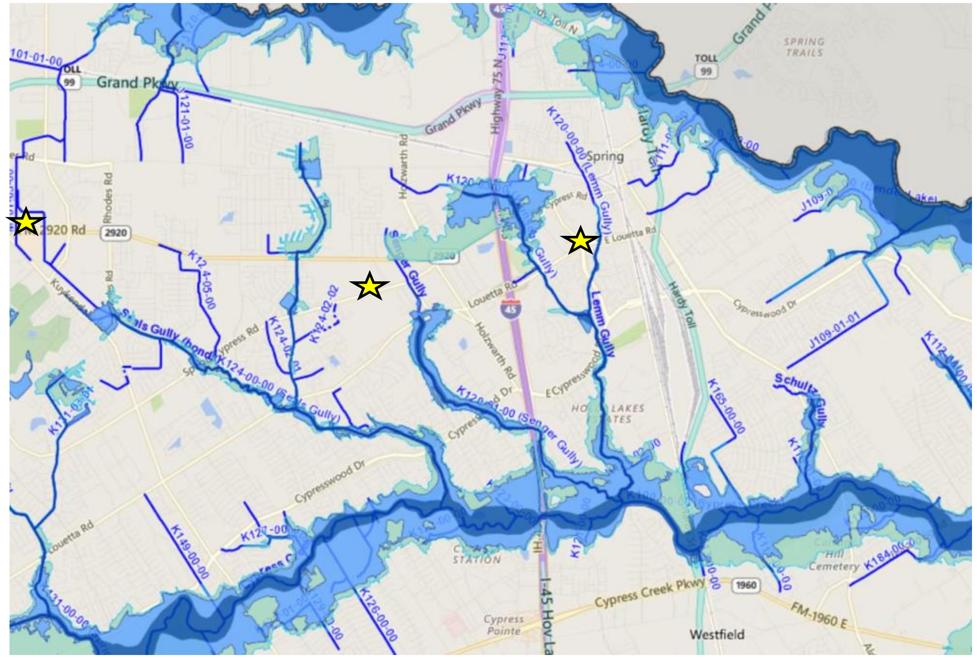

TRG Datacenters

High Wind Load

- Generators
  - Generator radiator exhaust being properly designed to prevent high static pressure situations during hurricane (turning the air up)
- Resupply due to long term outages
  - See prior comments regarding fuel storage and alternatives
- Outdoor equipment
  - Building rated for a wind load, is your outdoor equipment rated properly?
- Utility distribution failures
  - Short term failures at 60-70MPH+ from trees
  - Mitigate with proximity to sub stations
- Transmission Line Failures
  - Long term transmission failures at 135MPH+
  - Regardless of your proximity to substations, you transmission level distribution will be affected by these types of wind loads
- Building failures
  - X MPH design?

$$P = \frac{1}{2}\rho v^2$$

P – pressure, MPa

ρ – density of wind T/mm^3

v – velocity of wind mm/s

Wind (MPH)	150
P (Mpa)	0.00275412

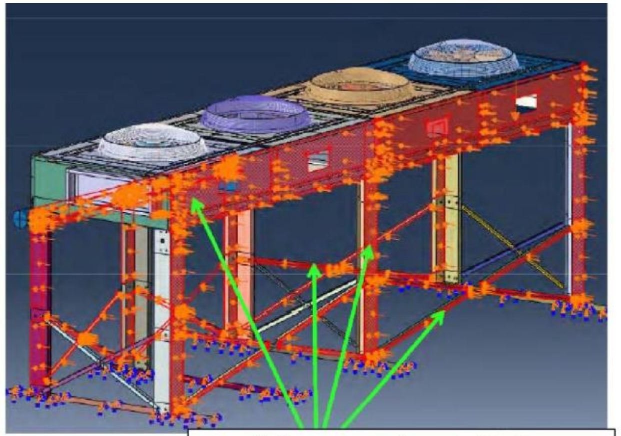

Uniform pressure from wind loading.

Trains and Hazmat Transport

- DOT (Department of Transportation) ERG and NIH have established guidelines for evacuation and threat zones
  - https://chemm.nlm.nih.gov/threatzone.htm
  - https://environmentalchemistry.com/yogi/hazmat/erg/
  - Generally at least .5 Miles, may be up to 2 miles
- Must consider
  - Distance from rail lines
  - Distance from major thoroughfares
    - Crash risk as well
  - Distance from fuel supply depot
  - Distance from major pipelines

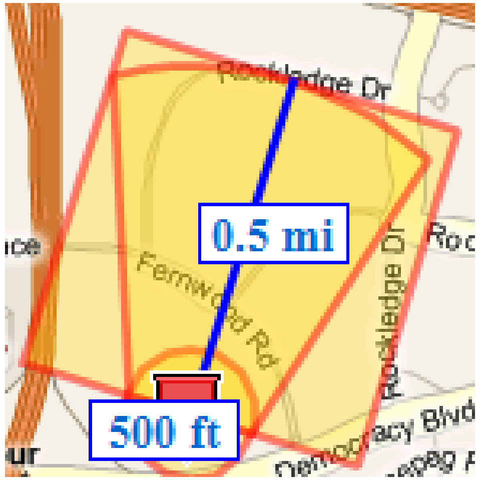

Construction

- For fiber lines and power, construction may be a risk to cut lines, especially to buried lines
  - Consider route diversity and carrier diversity with fiber
- Consider diversifying above ground/below ground fiber, especially if true route diversity is not available
  - Need to look at termination points as well

https://www.youtube.com/watch?v=UJQ-A9CsK2I

Temperature

- High temperatures in Houston must be considered
  - Generator high ambient radiator package
  - Most advertised HVAC and Condenser packages show capacities at 95 degree ambient
  - Standard ASHRAE design conditions often don’t consider true worst case scenario, which is important for a high uptime environment
- Low temperatures as well?
  - Refrigerant tuned for higher temperatures can actually have challenges when temperatures drop too low
  - Some chillers standard packages have performance issues below 20 degrees F, recently several enterprise datacenters suffered outages from the cold due to manufacturer settings not being changed
  - Checking settings and low ambient conditions is important

Matching a decision to your concept of operations

- Datacenters are a distinctly different animal than IT
  - Primary disciplines: HVAC, Electrical, Security, Structural, Civil
- Does your staff have the wherewithal or interest in managing their own?
- Nature of your load
  - Contract based vs long term load
    - Datacenters are long term assets (15 years or greater) vs. IT that is 3-8 years
    - Varying load requirements will lead to stranded assets long term that is not properly managed
- Smaller facilities carry much of the same operating expenses that larger facilities do for basic levels of security and operations

Summary

- Walk through and logically check your assumptions on your infrastructure around many built worst case scenarios
  - Not all “dual corded” load is built equal!
- Each region has its unique set of challenges that must be considered with design and site selection
- With HPC, protecting your equipment during transient outages and various thermal conditions is extremely important. Ask about these transient conditions vs. just looking at steady state.
- With regards to Harvey from a datacenter perspective, Houston got lucky with very little wind to compound the issues. Don’t rest on your facilities performance.

Questions?

TRG Datacenters
Houston One
2626 Spring Cypress Road Spring, TX 77388
www.trgdatacenters.com

Chris Hinkle
CTO
TRG Datacenters
832-663-6028