UofT Data Centre Design Considerations

January, 2021

Executive Summary

Properly designed and managed data centres comprise more than just a room with server cabinets and air conditioning. While it may be tempting for a department to build out a room to host computer servers and storage in order to maintain local access to and control of the hardware, there are significant risks associated with not following industry best practices. These risks include prolonged service outages caused by any of the many single points of failure that are common in departmental data centres, such as: IT hardware, power and cooling infrastructure. Equally important but often overlooked is ensuring that the facility is supported appropriately by staff with specialized skills. A comprehensive maintenance and testing program is also necessary to ensure that all of the many sub-systems associated with reliable data centre operations are properly serviced.

All of these considerations combined represent a substantial investment for even a small data centre. These costs are typically out of the financial reach for most university departments. Rather than building many small, distributed, inadequately designed and resourced "data centres" it is strongly recommended that departments consider other options such as virtual hosting in UofT's central ITS data centre, or public cloud environments such as Amazon Web Services and Microsoft Azure.

If a data centre is to be built, it is essential that the team who will ultimately be responsible for the ongoing management and maintenance of the facility are heavily involved from the start, including the detailed design, tendering and construction phases.

This document aims to detail the many considerations involved in building a data centre.

Background

In its simplest description, a data centre is a room full of server cabinets filled with computers that generate heat and consume electricity, in addition to whatever else those computers may do. How many cabinets are included, and therefore how much power and cooling are required, are considerations central to the process of Data Centre design.

The four basic questions that need to be answered in any data centre design are these:

1) What is the desired end-state IT load to be accommodated within the facility?
2) What is the desired average power density per cabinet?
3) Is a generator required?
4) Is a NOVEC 1230 fire suppression system required?

All other activities and costs stem from these four questions.

Reliability, Redundancy, and Serviceability

Obviously, we would like our computers to operate without interruption, indefinitely. However, every component in the networks of components that keep those computers operational will, at some point, fail. Even cloud-based services suffer interruptions of service due to physical component failures. In order to keep a physical computer operational, it must have a continuous supply of power and cooling. Every source of power and cooling requires routine maintenance.

Reliability

The Uptime Institute has a classification and certification system for the reliability of Data Centres, both in the design and in construction. Tier I is the lowest reliability and Tier IV is the highest. BMO Financial's facility in Barrie and the Province of Ontario's in Guelph are the only two facilities in Canada having been certified Tier IV by the Uptime Institute as constructed. A brief description of the four Tiers is copied here from the Uptime Institute's web site.

The Tiers (I-IV) are progressive; each Tier incorporates the requirements of all the lower Tiers.

> **Tier I: Basic Capacity** A Tier I data center provides dedicated site infrastructure to support information technology beyond an office setting. Tier I infrastructure includes a dedicated space for IT systems; an uninterruptible power supply (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: Redundant Capacity Components** Tier II facilities include redundant critical power and cooling components to provide select maintenance opportunities and an increased margin of safety against IT process disruptions that would result from site infrastructure equipment failures. The redundant components include power and cooling equipment such as UPS modules, chillers or pumps, and engine generators.
>
> **Tier III: Concurrently Maintainable** A Tier III data center requires no shutdowns for equipment replacement and maintenance. A redundant delivery path for power and cooling is added to the redundant critical components of Tier II so that each and every component needed to support the IT processing environment can be shut down and maintained without impact on the IT operation.
>
> **Tier IV: Fault Tolerance** Tier IV site infrastructure builds on Tier III, adding the concept of Fault Tolerance to the site infrastructure topology. Fault Tolerance means that when individual equipment failures or distribution path interruptions occur, the effects of the events are stopped short of the IT operations.

Whether or not one intends to have a data centre design certified before construction, or have the constructed facility certified after construction, being mindful of those brief descriptions is essential for all stakeholders. Which Tier definition one strives to meet will have the greatest impact on the data centre construction budget.

In the commercial collocation space, it is a common practice to promote a facility as "a Tier III Certified design". When pressed, many operators are forced to admit that only their design was certified, not their facility as constructed. Often, provision has been made for redundancy in the design phase, but the execution has been deferred until a future construction phase when there is money and demand for it. These facilities may never reach their end-state design, but at least a plan exists to add those redundancy components as needed.

Redundancy

If you have any single physical computer that must stay up 7x24x365, then you have a larger problem than can be addressed in this document. However, your servers can be specified and purchased in configurations that leverage the redundancy features of your data centre.

Assuming your servers are air-cooled, then so long as the supply of cold air is not interrupted, your primary concern becomes clean stable power; and that is the crux of both problems - keeping your server running and getting rid of your waste heat.

For your servers' power, the answer is simple – buy servers that have redundant hot-swappable power supply units (PSUs). When one PSU fails, the other continues to provide power until the failed unit can be replaced. Of course, if both PSUs are connected to the same circuit, the upstream breaker is likely to trip when the first PSU fails. There are open-circuit and short-circuit failures. When a PSU short-circuits, it usually trips its nearest upstream over-current device. So, each PSU needs to be on a separate circuit breaker.

For the facility operator, the issue is a little more complicated. Power Distribution Units (PDUs) within server cabinets provide power to the servers. Ideally, each cabinet is outfitted with redundant PDUs, each fed from a different source. At least one of those sources must be protected by an Uninterruptible Power System (UPS), ideally both are protected by independent UPSes. The design of the upstream electrical service must have a permanent generator if the data centre is to remain at full capacity during anything longer than a momentary (i.e. less than 5 minute) power failure. At a minimum, provision for a temporary mobile generator should be included so that the facility, including its cooling plant, can continue to operate during scheduled annual or biennial building electrical shutdowns.

Serviceability

There is a joke amongst electrical engineers – "If you forget about electricity long enough, it goes away." Everything requires maintenance. Deferred annual maintenance is maintenance foregone. Foregone maintenance results in component failure and unplanned outages.

Every circuit breaker and every valve in a data centre should be operated annually. That means that during the annual maintenance cycle, at some point each PDU in every cabinet will, in turn, be de-energized; so will every UPS, every electrical distribution panel, and every component in the cooling plant. The sequencing of these maintenance operations must be addressed during the data centre design phase.

Compromises will almost always have to be made, but it is important to understand and document the reasons for every decision that creates a condition where concurrent maintainability is to be sacrificed. Sometimes there is no money and no space for a permanent generator, so the facility will never reach Tier I capability. Sometimes there is no space for a second chiller or a second chilled-water loop, so the facility can never achieve Tier III. Often, the space required for a permanent generator or a second chilled-water loop makes it impossible to build a suitable data centre in an existing building.

Concurrent maintainability is a fundamental component of continuous and reliable data centre operation. Without it, annual complete facility shutdowns are the only way to accomplish the some of the routine electrical and mechanical maintenance tasks.

Monitoring and Alerting

There are many things that can, and do, go wrong in any facility. In a data centre, failing to notice when things go wrong usually leads to a loss of service. UPS batteries will eventually overheat, circuit breakers will trip, server cabinets will overheat, and equipment will fail. Timely alerts to equipment failures, sometimes even before they happen, may give facility operators time to react before the situation escalates to a full-scale loss of service.

At a minimum, the supply and return air and water temperatures should be monitored, with out-of-band alerting mechanisms, (SMS text and/or alternate email) to tell on-call staff what has occurred. PDUs should have metering for overall power draw, per phase, per bank, and perhaps even per outlet current monitoring. Alerting thresholds for these devices should be at 40% and 80% of their rated limits. Failure to alert, when loads shift and exhaust resources, can lead to cascade failures when a breaker trips and the load on the remaining device doubles unexpectedly.

Entry into a high-security area or cabinet must always produce an audit log and should send an alert to the central monitoring site.

Annual Maintenance and Testing Program

If you can't turn it off, you can't perform maintenance. Building a data centre or server room that must be shut down annually for preventative maintenance creates great difficulties for customers and facility operators. This is especially true in colocation facilities. If customer agreement is required before a maintenance window can be scheduled, there quickly becomes no time that is agreeable to all customers and someone will always be disappointed.

At a minimum, the cooling plant and the electrical plant need to be de-energized and, because the cooling plant maintenance cannot occur without the electrical service being fully operational, a minimum of two scheduled facility shutdowns are usually required when there are no backup systems to leverage. The following is a list of maintenance tasks that occur regularly in the ITS Administrative Data Centre (DCB):

5-Year Major Generator Service with Standby Rental Generator
Biennial Building Major Electrical Maintenance Shutdown
Annual Generator Major Service,2-Hour Load-bank Test, & Engine oil analysis
Annual UPS maintenance (UPS-1 & UPS-2)
Annual FLIR Scan of Electrical Equipment
Annual Glycol Testing
Annual Humidifier + CRAH Maintenance
Annual ATS Maintenance
Annual VFD Maintenance
Annual Fuel System Maintenance
Annual Electrical Busway & Dropbox Visual Inspection
Semi-Annual Fire Prevention/VEDA Testing
Quarterly Pump Inspection
Quarterly Chiller & CRAC Preventative Maintenance
Monthly Generator Test
Monthly ATS Test
Monthly Chiller Inspection
Monthly Heat-Exchanger Test (backup cooling changeover)
Daily Visual Inspection of IT Equipment

Energy Efficiency

Calling a data centre "green" is a stretch by any measure. However, there are simple design principles that can certainly improve the performance of the facility and reduce energy consumption, but energy savings are most often tied to economies of scale. It is very hard to make a small server room energy efficient.

Efficiencies found on the air side of the facility are called Air-Side Economizing; those on the water side are called Water-Side Economizing.

Air-side economizing depends almost entirely on preventing hot and cold air from mixing within the data hall. Three common methods for achieving this are, Hot-Aisle Containment (HAC), Cold-Aisle Containment (CAC), and Ducted Cabinets. In all three cases, hot exhaust air from the cabinets is sequestered and returned to the air handlers in an overhead plenum. To make that possible, each cabinet must be its own separate duct in the overall facility airflow system, such that cold air is prevented from entering each cabinet except as a consequence of being sucked through the IT equipment being cooled. This is achieved by installing blanking panels in all gaps on the front of the cabinets and ensuring that all equipment is mounted with front-to-rear airflow. Air must also be prevented from moving between cabinets by ensuring that every cabinet has its side panels installed.

By sequestering hot air and preventing the mixing of cold supply air with hot exhaust air, cooling equipment is more efficient. Having a larger temperature delta between the hot air being returned to the air handler and the cold supply air that it must deliver reduces energy consumption.

Once heat has been transferred into the chilled-water loop, rejection or reclamation of the waste heat are both options. Again, economies of scale limit the ability to reclaim waste heat in smaller server rooms and a low temperature differential between the supply and return water temperatures can make heat reclamation impractical.

Physical Security & Other Information Security considerations

There is a multitude of security controls that should be consistently implemented against servers or systems hosted in data centers. In the context of the data centre design, the physical security requirements are the most relevant controls that should be taken in consideration:

Security Requirements * **(What needs to be done)**	**Controls**** **(How requirements can be met)**
Limit physical access to organizational systems, equipment, and the respective operating environments to authorized individuals.	• Develops, approves, and maintains a list of individuals with authorized access to the facility where the information system resides; • Issues authorization credentials for facility access; • Reviews the access list detailing authorized facility access by individuals. • Removes individuals from the facility access list when access is no longer required.
Control and manage physical access devices.	• Enforces physical access authorizations at the University's data centre entry/exit points to the facility where the information system resides. • Verifying individual access authorizations before granting access to the facility; and • Maintains physical access audit logs for the entry/exit points of the University's data centres. • Escorts visitors and monitors visitor activity to University's data centres. • Secures keys, combinations, and other physical access devices; • Keeps inventory of physical access devices to control data centre access and review them on a regular basis. • Changes combinations and keys of physical access devices on a regular basis and/or when keys are lost, combinations are compromised, or individuals are transferred or terminated.

*Security requirements are derived directly from the UofT standards published in May 2020. The requirements are based on the NIST 800.171 special publication.

**Controls are derived from the NIST 800.53 standard.

Conclusion

There is a great temptation to keep physical servers in a convenient location close to people's offices, but convenience usually comes at some cost in both information security risks and inevitable adverse availability outcomes.

Offices are not data centres and converting office space into ersatz server rooms does not achieve substantially higher availability than simply leaving servers under a desk in a locked office. The features of a properly designed data centre that are likely financially out of reach for most individual departments could be achieved in a shared facility built with modern information security and BICSI 002/2019 Data Centre Design Best Practices in mind. Research computing, specifically, could benefit from a shared facility, certified to an appropriate standard and managed by professional data centre staff.

Server virtualization and cloud computing should always be considered first, especially for applications that do not require GPUs and do not have extraordinary CPU or RAM requirements. The ITS Private Cloud is a great option for many departments that have troublesome server rooms and limited hardware budgets.