HAM-TMC Library
Information Technology Services

Colocation of Critical Services

Project Description

Author: James C. (Chris) Young
Date: October 1, 2005
Revised:

Current Situation

Traditionally, libraries have maintained much of their archival and intellectual property in the form a physical copy, whether paper or film based. These physical assets are indexed by a cataloging system which allows customers to search against keywords and find the location of these assets within the physical library.

In the last decade, there has been a fundamental shift in the delivery and access to information, world wide, due to the adoption of the Internet and the search engines used to find data. As fewer customers are physically entering the premise, many more are accessing the electronic resources. Libraries are rapidly changing to meet the needs of customers including electronic access to journals, databases, books and other materials via authenticated access to publisher's content.

With the advent of electronic delivery of information there is a partial vacuum in the capabilities of library information technology infrastructures, particularly in backup power capabilities and off-site, on-line systems. There are several possible solutions to address these issues. Using models from other industries where the data and underlying infrastructure has to be secure and for the most part, "always on" (banking, telecommunication, healthcare), there have been on-going developments in power generation, network redundancy and placing systems and services in collocation facilities in either a hot-standby or primary services mode.

The HAM-TMC Library ("the Library") has significant public facing services that would be compromised in any event, natural of otherwise, that would compromise the power to the building. The primary concern is for the availability of the Library's web services as an index to available facilities and to the proxy services that allow the customers to access publisher's content via the Library's license.

Definitions Pertinent to Discussion

Business Continuation Plan – a pre-defined plan to insure that critical services are always available. This differs from a disaster recovery plan in that it is proactive and the plans execution is unaffected by adverse conditions.

Disaster Recovery Plan – a pre-defined plan to insure that critical data and services are protected during catastrophic failures due to weather, fire or power events. A disaster plan attempts to minimize the downtime and loss of data but does not provide for or assume consistent availability. This differs from a business continuation plan because in terms of availability, is reactive to the adverse condition.

Critical Services – Those applications, systems, and network elements that must remain available to provide a specific service.

Collocation- A data center, operated by a third party, which contains appropriate power, power conditioning, air conditioning, physical security, fire protection and network connectivity to support diverse and multiple customers with the goal of realizing economies of scale for small and medium businesses. In the long run, should provide better service at a lower cost than the business could provide for themselves.

Rack or Cabinet – An enclosure made to house electronic equipment including servers, network devices and wire management systems.

Rack unit (ru) – a measurement of the space available in a rack. 1 ru = 2 1/2 inches or 3 rack mounting holes.

"helping hands" or "remote hands" – at off-site facilities, it is sometimes necessary to perform an action that requires a physical interaction with a device, like turning the power on or off. "Remote hands" is a generic term for the temporary employment of collocation personnel to execute physical actions to devices.

Service Level Agreement (SLA) – is a contractual agreement that specify goals for availability and performance of a particular service. The contract also specifies and generally limits the provider's liability should the service fail to meet the goals.

Power grid – the interconnection of power generation and delivery systems

Network Backbone – Data networking service provided by a tier 1 network provider.

Network neutral collocation provider – a collocation provider that has available multiple tier 1 and tier 2 providers in their local location. As the name implies, they do not force the choice of a specific network provider in order to use the collocation facilities.

Tier 1 provider – In the United States, the "Internet" is comprised of seven major Internet Protocol (IP) backbones with interconnections between each. Tier 1 providers are one of the seven. Tier 1 providers have special privileges on the internet with requirements for redundancy, availability and speed of connection. Tier 1 providers have the highest likelihood of meeting five nines availability. Being "close to the backbone means that you are directly connected to a tier 1 provider. Tier 1 network providers are AT&T, Sprint, Level(3), MCI, Qwest, Verio and Genuity (now part of Level(3)).

Tier 2 provider – A smaller network provider that aggregates traffic for presentation to a tier 1 provider. Typically these providers also provide collocation and multiple interconnects to multiple providers locally.

Redundancy elements - in both power, network connectivity, and hardware means that there is no physical single point of failure between the source of the service and the consumer of the service. This includes building entrance, internal wiring to the rack or other colo facility type, and redundant support electronics.

Scheduled downtime – a pre-announced or standing timeframe that maintenance on systems can be performed. Does not count against availability statistics.

"Five Nines" (99.999%) reliability – and often quoted goal for uptime. Equates to no more than 5 minutes 16 seconds per calendar year.

> "In short, the metrics for five-nines include performance for some, but not all system elements and components. It includes the following:
>
> - Hardware components.
> - Power supplies.
> - Switch matrix.
> - Any other hardware component that can cause a total failure.

It does *not* include:

- Shut-down of the operating system.
- Loss of electrical power.
- Network loss.
- Time out for application software upgrades and fixes (can be 1–3 hours per month).
- Preventive maintenance (hours per month).

The fact that some call servers must be shut down when line cards, trunk cards or gateways are installed.
Complete server shutdown to install operating system changes or new releases. "
from the May 2002 issue of Business Communications Review, pp. 22–27
by Gary Audin,

"Four Nines" (99.99%) reliability – and often quoted goal for uptime. Equates to no more than 52 minutes 36 seconds per calendar year.

N+1 – redundancy model, in this case for power distribution. N+1 has the number of units required for operation plus one for redundancy.

Power Redundancy

There are several schools of thought covering the design of disaster recover/business continuation models. Regardless of whether the physical equipment resides on site or at a collocation facility, the fundamental elements of all of the plans include: redundant power in the form of generators, Un-interruptible Power Supplies and connectivity to multiple power grids. A typical layout is found below

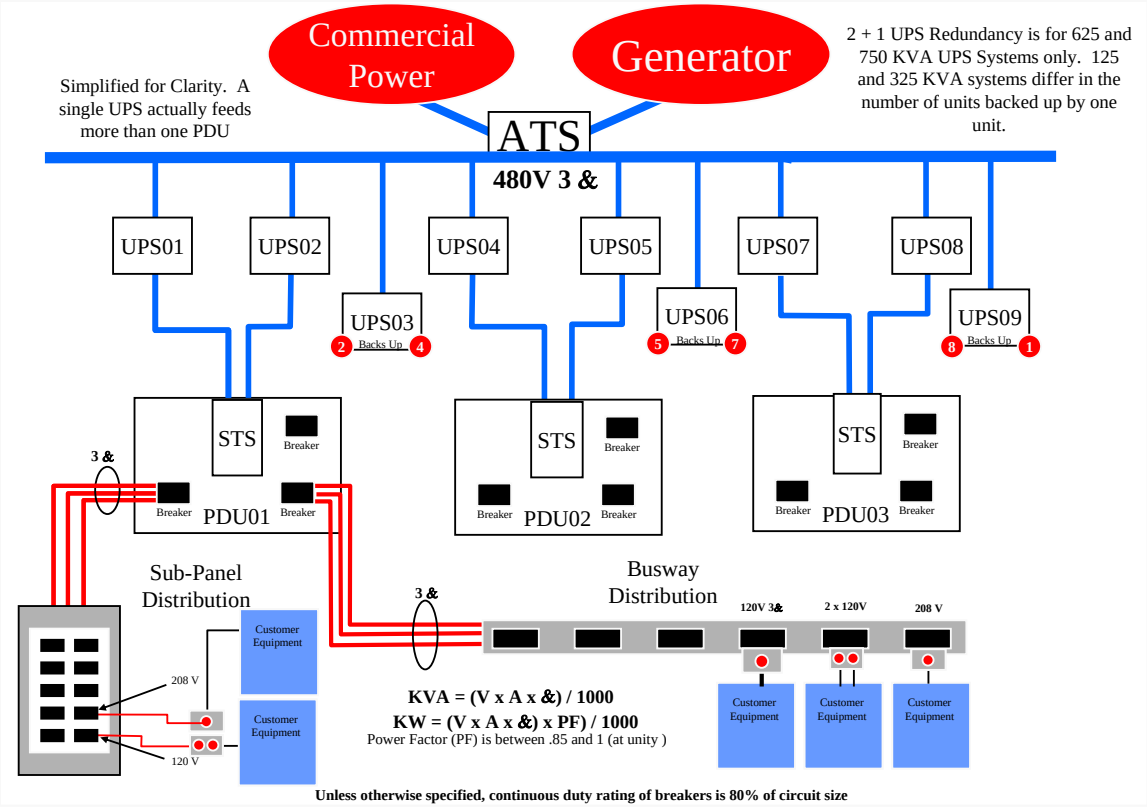

This is a typical n+1 power distribution model. Power enters the building either from the commercial power grid or a generator through an automated transfer switch (ATS). The ATS is responsible for determining if the commercial power is stable. If not, the ATS will start the generator, determine that the power from the generator is stable, match the phase of the commercial power if it is available, introduce the power to the internal power distribution and shunt the commercial power feed.

In this model, uninterruptible power supplies (UPS) pull power off of the building power buss. The UPS functions as a spike dissipation device as well as providing for power to the buildings devices until the power from the ATS is stable.

A power distribution unit (PDU) is placed regionally throughout the facility and feed by multiple UPS's in an n+1 arrangement for redundancy. The PDU feeds the individual cabinets/rack with power to run the collocation devices and internal network elements.

Network Redundancy

In keeping with n+1 redundancy found in the power distribution model, network connectivity also requires physical path and equipment to achieve high availability. In a collocation facility, the network connectivity is much more complicated than a typical enterprise network because of the multiple types of media that must be aggregated into the tier 1 networks from the customer cabinets. A typical network design that incorporates all speeds of Ethernet, serial transport connections, and modem routing is shown below.

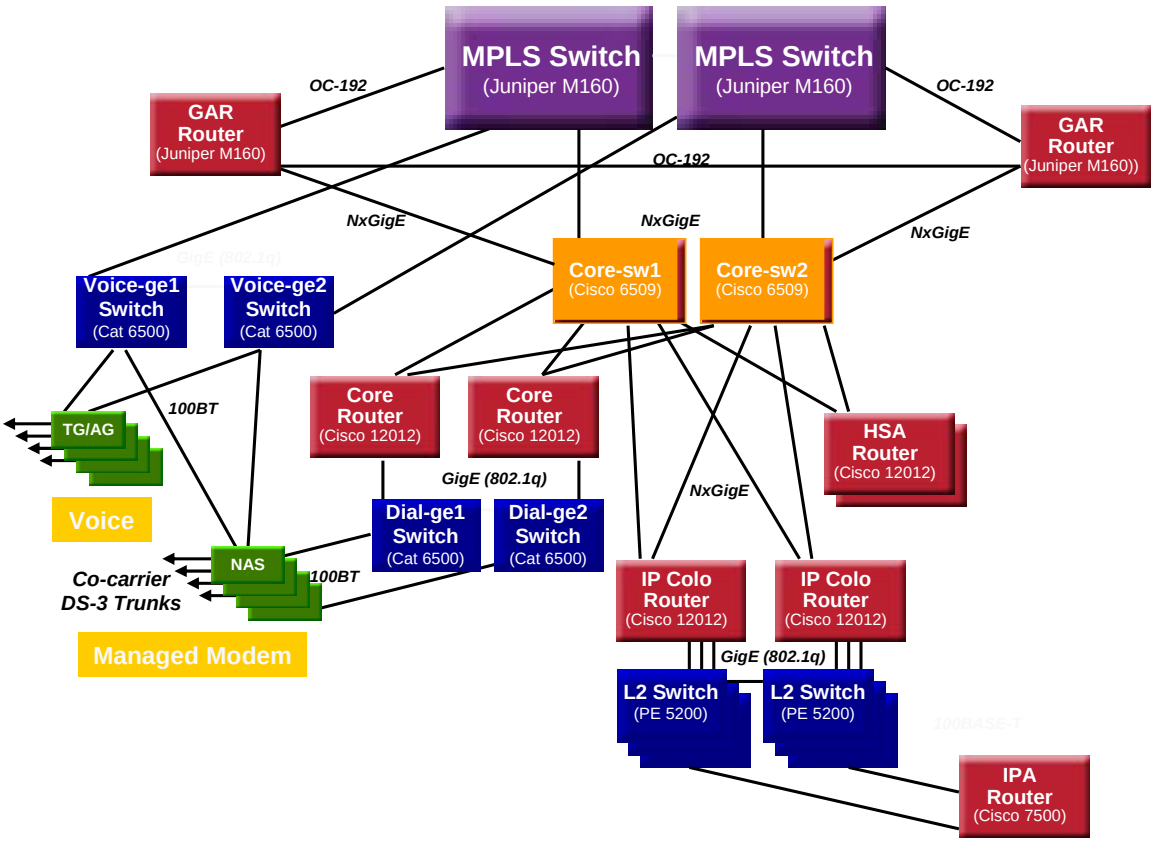

Network Redundancy

Hot-standby Site

In traditional disaster recovery models, a hot-standby site is a duplicate of your critical systems at another site. Network software configurations are set replicate all data for all systems and should the primary site become unavailable for any reason, the preconfigured routing preferences defer all traffic to the standby site. This option assumes that the primary site is designed to be as resilient as possible with all of the redundancy elements in place and a near duplicate site is available off-site.

Primary service off-site

Where the redundancy elements are not available, or cost prohibitive at the primary site, it if preferential to place the service at a site that does have the redundancy elements and use the weaker site as it's backup.

Critical Systems at the Library

By survey, the critical services of the Library are the access to "remote access" electronic resources, the Library catalog. Although there are systems that are critical to the internal business of the library, there are only a handful that are customer facing

Proposed Situation

In the final analysis, in order to provide the most resilient delivery of services to the Library's customers, the use of a off-site collocation facility as the primary site for the HAM-TMC Library's main web server, electronic access authentication system (Shiboleth/EzProxy) and requisite support systems be implemented as the primary site for these services. Implementing network and servers with existing hardware will save money but will need to be replaced if traffic significantly increases to the services.

Network Design

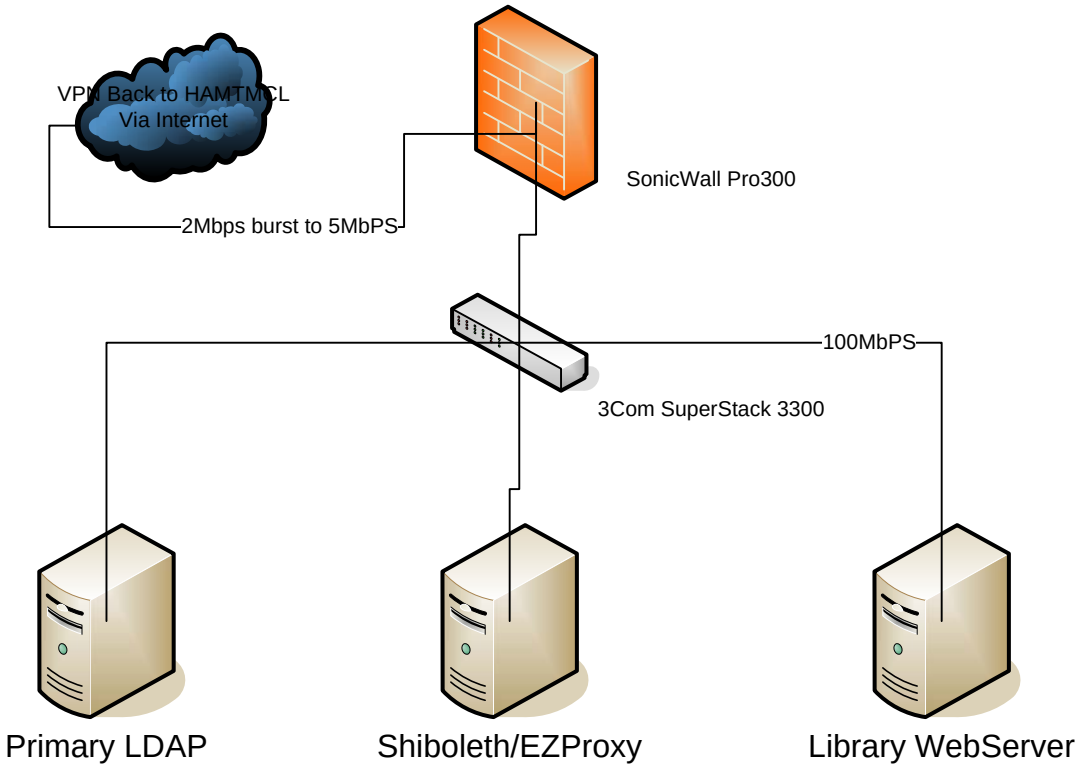

Schedule

[Schedule Deleted because of proprietary nature]

Cost

[Costs Deleted because of proprietary nature]

Benefits

As shown during the Fall 2005 hurricane season, access to the Library's electronic journals and evidence based medicine titles are critical for delivery of patient care, especially during times when doctor's may be providing services that are outside of their specialty.

Having the HAM-TMC Library critical services housed off-site, in a data center class facility with carrier class connectivity will allow users of Library resources to regardless of the availability of power within our data center.

Since the quality of the power and connectivity of the selected site is far superior than the power and the connectivity available at the Library. By placing the off-site services as primary and the Library as backup gives a statistical advantage as it is more likely that the

Library will lose connectivity or power for an extended period of time than a collocation facility.