



Network System Architects
Engineering Solutions For Business Problems

PLANNING A DATA CENTER

**By Steven Kieffer, Will Spencer, Art Schmidt & Steve Lyszyk
Network System Architects, Inc. (NSAi)**

**© Network System Architects, Inc. (NSAi)
February 2003**

COPYRIGHT NOTICE

Network System Architects, Inc. (NSAi)

This document and the information contained here is the subject of copyright and intellectual property rights under international convention. All rights reserved. No part of this document may be reproduced or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, or stored in any retrieval system of any nature, without the prior written permission of Network System Architects, Inc. (NSAi).

Network System Architects, Inc. (NSAi)

1550 Larimer St., Suite 222
Denver, CO 80202

TABLE OF CONTENTS

1.	INTRODUCTION	5
2.	FLOOR SPACE - SIZE AND LAYOUT.....	6
2.1.	Aisle, Ramp, Racks and Cooling Dimensions	6
3.	RAISED FLOORING	7
3.1.	Floor Load Planning	7
3.1.1.	Entire Floor Load.....	7
3.1.2.	Tile Load Capacity	7
4.	POWER	8
4.1.	UPS and Backup Generators	8
4.2.	Power Conditioners	9
4.3.	Surge Arrestors	9
4.4.	Emergency Power Off	9
5.	HVAC (HEATING, VENTILATING, AIR CONDITIONING)	10
6.	NETWORK	11
6.1.	Network Cabling	11
6.2.	External Connectivity	11
6.3.	Network Devices.....	12
6.4.	Servers	12
6.5.	Storage.....	13
7.	SERIAL CONSOLES AND KVM (KEYBOARD, VIDEO, MOUSE) SWITCHES	14
8.	NOC (NETWORK OPERATIONS CENTER)	15
9.	BUSINESS CONTINUITY / DISASTER RECOVERY	16
9.1.	Data Backups	16
9.2.	Standby Sites	16
9.2.1.	Hot Standby Site	16
9.2.2.	Warm Standby Site	16
9.2.3.	Cold Standby Site	16
9.3.	Global Load-Balancing	16
10.	FIRE PROTECTION	18
10.1.	FM200	18
10.2.	Sprinkler System	18

11.	SECURITY	19
11.1.	Physical Security	19
11.2.	Logical Security	19
11.2.1.	Computer Security	19
11.2.2.	Network Security	19
11.2.3.	Intrusion Detection	20
12.	SUMMARY	21

1. INTRODUCTION

A data center is a very large investment for most companies, one that will provide a return on investment over a very long period of time. To ensure that ROI (Return on Investment), a great deal of forethought and planning is required.

Some of the key considerations for a data center that should be addressed as early as possible are:

- Adequate floor space and layout
- Power
- Cooling
- Network connectivity
- Support/Maintenance strategy
- Business Continuity/Disaster Recovery
- Fire Suppression
- Security – Physical, Computer, Network

Each of these items is critical to a functioning data center and will be detailed in this document. Other important items are also covered.

2. FLOOR SPACE - SIZE AND LAYOUT

Floor space in a data center is normally divided between space for servers and network gear, and space for operators and managers.

Modern computer and networking equipment is much smaller than legacy equipment, so data center floor space requirements have tended to shrink over the last few years. However, your situation will require careful analysis. It's always better to have too much space in your data center than too little.

2.1. Aisle, Ramp, Racks and Cooling Dimensions

In calculating the floor space required for the data center you not only have to calculate the dimensions of all the racks and equipment to be placed in the room; you must plan space for aisles, ramps and the cooling requirements of the racks.

Aisle space must be wide enough for racks to easily be brought in and out without touching or moving other racks in the room. Careful consideration needs to be taken with any special clearance requirements for electrical panels, fire suppression systems, cooling devices, rack door openings (for enclosed racks) and room to perform maintenance on the devices within the racks as some are designed to slide out into the aisle for maintenance access. There should also be several breaks in the rows across the room to allow for the efficient maneuvering of racks and people throughout the data center. Depending on equipment type used, the need for dock access and sufficiently sized entry points needs to be evaluated. If the data center is to reside above the main floor, elevator size and opening dimensions will play an important part in any equipment delivery and usage.

In addition to the rack's width and depth dimensions, the cooling requirements of the enclosed rack must be considered. Most racks pass air through them for the cooling of the equipment they contain. Some racks pass cool air in from one side and exhaust out the other. This type of rack typically requires several inches of clearance on either side to accommodate proper airflow. Many racks pass air from front to back or from bottom to top, allowing the racks to be placed right next to each other. This is a more space efficient design and should be considered when choosing a rack.

Different sizes of racks are also available, varying in width and depth depending on the type of equipment they will contain. To avoid wasted space, equipment that is too large to fit in the standard 19" rack (22" W x 80" H x 30" D) should be placed together in the same racks and area of the data center. Equipment that is not rack-mountable, such as large disk arrays, DLTs, and stand-alone server and network devices, should be similarly placed.

A simple diagram, done on a grid of your floor-tile layout, will go a long way to helping position your equipment for the best possible use of floor space.

Typical free floor space for a well-designed data center is from 40 to 50 percent of its total square footage.

3. RAISED FLOORING

Raised flooring is, by far, the most popular choice in data center design for many reasons. It gives you the most flexibility for network and electrical cabling and HVAC (Heating, Ventilating and Air Conditioning) air flow system design. Raised floor cooling systems are also typically cheaper and require less electricity than open-air cooling systems, as only the air beneath the floor and within the racks needs to be cooled. With an open-air cooling system, the temperature of the entire room must be controlled, requiring more capacity in the air-handling system.

A raised floor is typically built 24 inches off the ground with the surface made up of two-foot square tiles. These tiles can be removed and replaced to allow access to wiring that may lie beneath. Strategically placed grated or perforated tiles allow for controlled airflow to cool the racks.

Cable trays for all the cabling that will run under the floor need to be well thought out and planned in advance. They should be installed above the main flooring to keep the cables out of any potential flooding and meet any special requirements there may be for the type of cables being run, such as fiber.

When all the tiles are in place, you can't easily see what is going on underneath. One concern is the possibility of flooding. Moisture/water sensors must be placed under the flooring to detect this. These sensors usually run SNMP and can tie into your monitoring solution. Fire suppression systems should also be deployed under the raised floor, as the smoke from a fire could go undetected by ceiling-mounted smoke detectors, or possibly consumed directly into the air handling system.

3.1. Floor Load Planning

Weight is a very important factor to consider when designing a data center. Incorrect or incomplete calculations could spell disaster.

3.1.1. Entire Floor Load

The first factor to consider is the total weight of the entire floor. The weight of all the racks and other equipment, current and future, must be calculated. This is especially important to consider when the data center is not going to be on the ground floor of the building. You must consult with your building's design to determine its maximum weight capacities. Don't forget to calculate the weight of the raised floor itself when figuring floor load; floor tiles and the load-bearing framework that hold them are very heavy.

3.1.2. Tile Load Capacity

There are different types of tiles, each with a different load rating. The capacity of the tile must exceed the weight placed upon it. With racks that have casters, it is important to think of the point load capacity of the tile when planning. The point load of a rack is its weight at any one of its casters. For example, if a rack weighs 2000 pounds, the weight at any one of its four casters should be 500 pounds. Because it is possible that two casters, from two different racks, could rest on the same tile, you need to double the point load value when figuring tile load capacity requirements. The load capacity of the tile that is chosen should exceed one half the weight of the room's heaviest rack.

4. POWER

Planning power for your data center normally requires planning these four components of a power plan:

- Electrical power from the local utility company
- Power filtering and monitoring
- Backup UPSs (Uninterruptible Power Supplies)
- Backup generators

The main planning task involving electrical power from the local utility company concerns power grids. It is best if your data center can be connected to multiple power grids. If power to one grid is lost, all or at least some portion of your data center will continue to operate normally. Bear in mind that multiple power grids will require multiple sets of power transformers, circuit breaker panels, and battery backup units.

Power requirements can be determined by researching manufacturer specifications for each piece of equipment you wish to place in your data center and tallying the results. However, the average lifespan for a server is four years and hopefully your data center will survive many iterations of new server technology. Therefore, you have to plan with generous rough estimates and add a safety factor (20-50%) for growth. Keep in mind that if the devices have redundant power supplies, you must figure in the combined wattage drain.

A current industry standard for estimating power requirements is to estimate 60 watts of power per square foot. As network and server devices are becoming smaller and smaller, the standard estimation is quickly rising to 100 watts per square foot. Another equation that yields similar, but not identical, results is to estimate 4 kilowatts of power per 19" rack. Just a few years ago, that rule of thumb was 2 kilowatts per 19" rack.

When calculating the total power requirement to feed your data center, you must figure in the power required by the HVAC system. A standard way to measure this is to add up the total equipment power requirement and add 70 percent.

Power typically needs to be available in both 115v and 230v to every location within the data center. The proper jacks must be installed to match the type the equipment you will be connecting. You must also determine if your equipment will need single or three phase power. Some equipment will require special adapters for connection to single or three-phase power; ensure that you have reviewed the technical specifications for any device that does not plug into a standard 110VAC 60 cycle receptacle.

Power to devices with redundant power supplies should come from separate circuits, providing an additional layer of redundancy and stability.

4.1. UPS and Backup Generators

UPS sizing should take into account the amount of total power required to operate the data center and the length of time required to get the backup generator into service.

Generator sizing should take into account the total power required to operate the data center and the total length of time you want the data center to be able to operate on its own power. It does little good for your generator to provide 2,000amps if it runs out of fuel in 30 minutes. Make sure you include the HVAC and emergency lighting power requirements in your calculations.

UPS capacity is typically measured in Volt-Amp's (VA). VA is a unit for measuring power. Like the Watt (W), it describes a quantity of electrical power. To convert from W to VA, multiply by 1.4. To convert from VA to W, multiply by 0.714. Make sure you size the UPS with your equipment's

peak power load in mind. A device might draw 1000 watts during normal operation but when the device is turned on it might draw 1500 for startup. You must size a UPS that will handle this peak load.

Both UPSs and backup generators should be tested at regular intervals, at least once every six months. Depending upon your business requirements, redundant UPS and backup generators may be appropriate.

4.2. Power Conditioners

The electronic equipment that the data center will house can be extremely sensitive to “dirty” power. “Dirty” power is that which has high frequency noise in the line, varying voltages, surges, and other electrical impurities. These electrical impurities can disrupt and even ruin sensitive electronic equipment. The electrical system should be tested for quality of power. If not found to be within acceptable tolerances, power conditioners can be installed to “clean” the power and protect the data center equipment.

4.3. Surge Arrestors

Voltage spikes can disrupt or even destroy data center equipment. Surge arresting equipment should be included as part of the electrical system. Most Power conditioners and UPSs perform surge arresting as part of their feature set. If you are not using a power conditioner or UPS, a separate surge arresting unit must be installed.

4.4. Emergency Power Off

According to NFPA (National Fire Protection Association) 75 “Standard for the Protection of Electronic Computer/Data Processing Equipment”, an emergency power off switch must be placed at every point of entry to the data center. This switch must cut power to every computer system, UPS and HVAC. It must be clearly labeled and unobstructed.

5. HVAC (HEATING, VENTILATING, AIR CONDITIONING)

Modern computing and network equipment is much smaller than equipment utilized in even the recent past. A small number of large servers have been replaced with a large number of small servers, and those small rack-mount servers are giving way to blade servers. More servers in the same space mean more heat that must be dissipated.

An industry standard rule of thumb for estimating cooling requirements is to estimate one BTU (British Thermal Unit) of cooling for every three kVa of power. You will often see cooling described in terms of tonnage. One ton of air is equal to 12,000 BTUs.

Maintaining the data center between 70 and 74 F provides the best balance between safe equipment operating temperature, operator comfort, and acceptable relative humidity levels.

The relative humidity of a data center is optimally kept between 45 and 50 percent range. Relative humidity levels that are too high can cause component corrosion. Levels that are too low can make the occurrences of ESD (electrostatic discharge) dangerously high.

As mentioned earlier in this document, a raised floor system makes it very easy to design and tailor an HVAC airflow system to cool the equipment racks. The HVAC system pumps cooled air into the plenum between the sub floor and the raised floor. Grated or perforated tiles are strategically placed in front of or underneath (depending on the design of the rack) the racks to be cooled. The holes in these special tiles allow the cool air to escape the sub floor and be delivered to the racks. The return air is then drawn back into the HVAC to be cooled and delivered to the racks again.

HVAC equipment should be redundant. It does you little good to have power and network connectivity if you have to shut down your servers because they can't be kept cool.

Critical to the operation of the data center, the air temperature, humidity level, and the health of the cooling systems must be monitored. Most modern HVAC systems offer an SNMP agent that can be tied in to your monitoring system in the NOC.

6. NETWORK

The purpose of building a data center is to house computer equipment. These computers need network connectivity to communicate and make the center functional. A well designed network with either room for growth or a simple upgrade path will allow your data center to run smoothly and perform its function as planned. Otherwise, the network can easily become bottlenecked and degrade performance over time. This is why the network should be as modular as possible. VLAN technology in today's switches makes this fairly easy.

6.1. Network Cabling

The first thing that must be determined is the connectivity requirements. Usually data cable needs to consist of a mix of fiber optic and Cat5e or Cat5E rated cables. Cat5 is no longer recommended for new installs and Cat6 specifications have not been finalized which could render some of the currently manufactured Cat6 obsolete. However, much of the cable choice will be determined by the needs of the equipment you are installing. Fiber optic and Cat5e or Cat5E are used for the data communications, connectivity to your SNMP-enabled air handling and power devices, and to allow connectivity to the Internet and/or to other offices and data centers in your WAN.

Distribution of network equipment is especially important in larger data centers, where cable runs might exceed the cable standard. Orientation of the data center to the building's Telco equipment should be taken into consideration, as well. Placement of the data center's firewall and externally-connecting equipment should be as close to the Telco closet's ingress into the data center as possible, to minimize latency and possible damage to the cabling carrying vital external connectivity.

Most data center designers prefer to run cabling in overhead cable trays versus under the raised flooring. This is because cable runs under the floor can reduce the cooling efficiency of the raised floor if there is not sufficient space for air-flow to occur. In some instances, a mixed design can be used, where longer fiber-optic runs and some power cabling are done under the raised floor (in sealed cable trays) and the bulkier Cat-5, SCSI, and shorter fiber-optic runs are placed in overhead cable trays. If the raised floor is shallower than 24", overhead cable trays are recommended for all cable runs.

6.2. External Connectivity

The main planning issues involving network connectivity for data centers are capacity and redundancy.

Capacity refers to the total network bandwidth into and out of your data center. This is usually expressed in megabits per second.

Leased Line	Capacity
T-1/DS-1	1.544 Mb/s
T-3/DS-3	44.75 Mb/s
OC-3	155 Mb/s

Network redundancy is provided both by selecting network providers that offer highly redundant network connectivity, and by selecting multiple network providers.

Network redundancy can be planned on a symmetrical or asymmetrical basis. An example of symmetrical network redundancy would be provisioning a T-3 circuit from one Tier-1 provider and a T-3 circuit from another Tier-1 provider. An example of asymmetrical network redundancy would be provisioning a T-3 circuit from one Tier-1 provider and a T-1 circuit from another Tier-1 provider.

Network redundancy can be active-active or active-passive. Active-active network redundancy is significantly preferable. Active-active redundancy means that both circuits are constantly up and passing traffic. Active-passive redundancy means that the backup circuit only becomes operational when the primary circuit fails. Active-passive redundancy is used to provide redundancy at a lower cost. Active-active redundancy provides more total bandwidth. Active-active redundancy is also superior because the backup circuit is constantly being tested.

However, there is one major planning step to be aware of when utilizing active-active network redundancy. If your data center network capacity requirements are 2.5Mb/s, and you provide that with two T-1 lines – you do not have network redundancy. In the event of an outage of one of your circuits your total network capacity will be 1.544Mb/s. To provide effective network redundancy, you will need to provision at least three T-1 circuits from three different providers, or four T1s, two each from different providers.

Active-active redundancy is normally enabled through the use of the BGP4 (Border Gateway Protocol 4) routing protocol.

Another important issue in provisioning network capacity for your data center is performance, which is often measured in latency. A 1.544Mb leased line will have much better performance (i.e. lower latency) than a 1.544Mb satellite link. When purchasing bandwidth, the Internet routing model, ISP-to-customer routing model, private and public peering relationships, and type of SONET technology being used by potential providers should all be taken into account.

In addition, you often have to worry about network congestion caused by oversubscription on the part of your network provider. If they are selling T-1s to fifty customers, and they only have a single T-3 upstream to their provider, you may not actually see 1.544Mb of throughput on your T-1 circuit. This information is often difficult to locate; a thorough investigation of news and events about the provider can usually be found online, and be sure to request testimonials from the provider's current customer list.

Finally, when purchasing bandwidth, it is important to negotiate appropriate SLAs (Service Level Agreements) with your providers. The SLA should state promised uptime and latency. It should also discuss rebates and remedies for failure to meet the agreed upon service levels.

6.3. Network Devices

In addition to redundant WAN (Wide Area Network) circuits; your network devices should be redundant at the edge. This means that you should plan for two Internet routers meshed to your internal LAN (Local Area Network).

LAN devices are often not redundant, as an outage in a LAN device usually only affects a small portion of the data center. In addition, LAN devices usually operate only at the Data Link layer of the OSI model, which has the effect of making them highly robust. Finally, LAN devices will be so numerous in most data centers that maintaining a small percentage, typically 10-20%, of spares on-site is usually a better approach than building a completely redundant LAN.

6.4. Servers

Servers have gone through many generations of evolution. The first generation of servers were large mainframes, where you would often see one server per data center. Then minicomputers then came along, where the server would be no larger than a few home refrigerators. Workgroup servers were the next norm, when you could fit from two to eight servers in a single 19" rack.

From there, servers kept shrinking. Many servers are now from 1U to 5U in height. A 'U' is a measurement of rack height that equals 1.75" (44.45mm). Blade servers are now becoming available when many servers share the same small case. In these systems, you may have eight servers in one 5U case.

Higher density server farms also mean higher concentration of power and HVAC consumption; be sure to research these numbers carefully.

6.5. Storage

Storage has matured from mainly DASD (Directly Attached Storage Device) to NAS (Network Attached Storage) and SAN (Storage Area Network) technologies. The effect of these changes has been to move storage from a required server component to a data center wide *service*.

Centralized storage is provided via NAS and SAN technology through a subscription model by various servers and departments as a standard data center resource.

The marketing department, for example, may request 50GB of storage while the accounting department requests 2 terabytes of storage. Both of these storage requirements will be met by the same NAS or SAN device. This centralization enables a much lower TCO for each megabyte of storage space, in addition to enhancing reliability.

NAS devices move data traffic over the same (usually Ethernet) network as the other network traffic in a data center. SAN devices move data traffic over a dedicated (usually Fiber Channel) network. NAS tends to be a less expensive centralized storage solution, while SAN provides greater throughput and scalability.

7. SERIAL CONSOLES AND KVM (KEYBOARD, VIDEO, MOUSE) SWITCHES

Most industrial strength network and computing equipment can be configured via a serial console. This is the preferred method for remotely connecting to a router or a UNIX system that has failed. A serial cable is connected from the router or UNIX system to a centralized terminal server. The terminal server is usually reached by connecting to it from across the network using `telnet`.

PC servers usually do not feature serial consoles. For PCs, you have to connect a keyboard, a video monitor, and a mouse. Putting these items on a cart and rolling them around the data center is inefficient and non-scalable. A KVM switch will allow one or more centralized KVM units to communicate with a number of PC servers. Once again, distance is a limitation and placement of PCs within the data center should be planned accordingly, relative to the KVM device(s) that will be used to access them.

Serial communication servers require one cable per server. KVM switches typically require three cables per server. These can quickly consume limited space in many cable runs.

8. NOC (NETWORK OPERATIONS CENTER)

The NOC is the place from where the network is supervised, monitored, and maintained. It should be the central point for troubleshooting, software distribution and updating, router management, and performance monitoring. This includes monitoring of the climate control systems, UPS and generators. If you have many data centers spread across different time zones, it may be beneficial to incorporate the ability to transfer control of the different data centers to your different NOCs. This affords you a larger pool of engineers for human redundancy with a greater opportunity to avoid late night shifts.

The NOC will be the centralized point for communications into and out of the data center, as well as providing personnel with monitoring and management capabilities for most, if not all, devices within the data center. Ample space for monitoring devices, consoles, workstations, and telephony equipment should be planned. And don't forget space for personnel themselves, such as desks, chairs, and the occasional family portrait.

9. BUSINESS CONTINUITY / DISASTER RECOVERY

Disasters happen. Fire, earthquakes, floods, intentional explosions, and out of control trucks passing through your data center are all things that could happen. With proper planning, you can minimize the impact that a disaster can have on your company's operation. Statistics show that 90% of companies that lose their data in a disaster without means of recovery are out of business within 2 years.

9.1. Data Backups

Tape backups are the most common and least expensive form of backup today. Backups are usually run daily with the tapes being kept offsite for safe storage. This not only protects you from losing the data if the data center is destroyed, but also protects against accidental erasure or corruption on the main storage hard drive.

Data vaulting is another method of data backup which involve a WAN link to a remote backup facility. The data is continuously backed up across the wire and archived off site. This allows you to have up to the minute off site backups compared to the usual daily backups with a tape storage scheme. It also allows for immediate access to your backed up data, should you need it. This method is substantially more expensive than tape backup, so determine your requirements and plan accordingly.

9.2. Standby Sites

Standby sites are facilities at different locations that can take over operations should a main data center be crippled to the point of non operation.

9.2.1. Hot Standby Site

A hot standby site is a one that duplicates the critical systems in your main data center. It has the same systems installed with the same software, current data, and external connectivity ready to take over operations at the drop of a hat. As you might imagine, hot standby sites are a very expensive option.

9.2.2. Warm Standby Site

A warm standby site is similar to a hot site in that it has duplicate critical systems and external connectivity ready to use. It is different in that the systems are not loaded with all the current data needed for full operation. Backup tapes must be brought to the warm facility and loaded onto the backup systems. The drawback to this type of setup is that it takes longer to get up and running than with a hot site.

9.2.3. Cold Standby Site

A cold standby site is a backup site with minimal or no equipment. Replacement systems have to be brought in, configured, and loaded with data, and external connectivity has to be ordered and installed before operations can resume. Cold standby sites are the least expensive because equipment is only purchased if needed but it has the longest time to functional operation.

9.3. Global Load-Balancing

Two or more data centers can carry the active load of your applications through means of global load-balancing. This is a method by which two or more locations on the Internet take in all

connection requests for your application(s) and actively load-balance them between multiple data centers. This method of load-balancing is done in redundant fashion, so that the load-balancer itself does not become a single point of failure for the entire business.

If the business need is to deploy highly-redundant data centers on a continental or global scale, this method can be more cost effective than a Hot Standby data center. With global load-balancing, each load-balanced data center need not be built to carry the full load; instead, three or more data centers can be built in different locations. Each site is sized appropriate to carry $1/N-1$ of the total load of the enterprise, wherein N is equal to the number of data centers being load-balanced.

10. FIRE PROTECTION

The data center should to be equipped with a passive fire suppression system. That is, one that activates automatically with no human intervention. These systems come in two main varieties, chemical suppression systems and sprinkler systems.

There is also a need to have manual gas or sprinkler activation switches, portable fire extinguishers, and floor tile lifters placed throughout the data center. They need to be clearly marked and unobstructed.

10.1. FM200

FM200 is the best option for fire suppression. FM200 employs heptafluoropropane, an invisible gas that draws the heat energy out of the fire and stops combustion. It is not damaging to the equipment or personnel. With the use of this non-harmful agent, the data center can return to operation more quickly after a fire incident. When using a gas suppression system, the law requires that manual abort switches be placed in easily accessible locations throughout the center. Check with local fire codes to determine the specifics.

10.2. Sprinkler System

Sprinkler systems deliver water to the data center. This is the most damaging option for the data center equipment. Sprinkler systems come in two varieties, wet pipe and dry pipe. Wet pipe systems constantly have pressured water in the pipes on the ceiling of the data center. Dry pipe systems do not have water delivered to the ceiling pipes until it is activated. This second option lessens the chance for water leaks.

It is important to integrate a sprinkler system with the data center power cut off so power is severed before the water is delivered. All that water mixed with all that electricity makes for an extremely hazardous situation. It is also important to include a drainage system so the room does not fill up with water.

Some state and local fire codes mandate that sprinkler systems be present regardless of the use of a gas suppression system, like FM200. Check with local fire department or the county clerk's office to obtain the proper regulations for the area(s) you plan to build in.

11. SECURITY

The data center not only contains very expensive equipment that someone might want to steal, but also valuable data that essential to your company's survival. You must make sure that the data center is protected against theft, sabotage, vandalism, and industrial espionage. For that purpose, overall security policies should be put in place that address all issues in a high level fashion, defining guidelines for the more specific physical, computer and network security requirements. The key item to remember, especially in regards to computer and network security, is that nothing is foolproof.

11.1. Physical Security

The NOC can also function as the point of security control for the data center. If you allow access to the data center only through the NOC, you have a single point for monitoring and control.

There are several things to keep in mind when designing physical security into a data center.

- Avoid using an exterior building wall as one of the center's walls.
- If you must use an exterior wall, do not have windows. This is also a concern with cooling; windows allow in excess heat and can add to the workload of your HVAC system.
- All doors to the data center and NOC must be controlled by access devices such as card readers and cipher locks. Access attempts, successful or failed, should be logged.
- Only authorized personnel should be permitted access.
- Video cameras should be placed at strategic locations throughout the facility and be monitored and recorded
- Motion detectors and alarms should be employed and monitored. Deployment under raised flooring and in drop ceilings not only helps deter unauthorized intrusion, but can help maintain a stable environment by detecting, and preventing, unscheduled (and unauthorized) maintenance to take place to cable runs, power, or other critical systems.
- Air ducts must be checked to ensure they cannot be used to gain entry into the room.

11.2. Logical Security

There are things to keep in mind when designing Logical security into a data center.

11.2.1. Computer Security

- Several levels of authorization should exist for administering devices. Engineers should be granted the minimum access necessary to complete their tasks.
- Server console access should be exclusive to a separate administrative network that can only be accessed from the NOC, where possible. If requirements specify access from a central administration area outside of the NOC, network connectivity should be run to that specific location, parallel to the company's internal backbone. At the very least, strong encryption should be used for this kind of access.

11.2.2. Network Security

Network security is best implemented in a tiered fashion.

- Tier 1 is typically edge-access protection via firewalls for safeguarding access into the network.

- Tier 2 may be a firewall that separates publicly accessible devices such as web servers, DNS servers, mail relays, etc. from the rest of the internal network. Many times the Tier 1 and Tier 2 policies reside on the same physical device. In addition, communications that pass confidential or sensitive data to the outside should be strongly encrypted using a method like VPN tunneling that can be setup in parallel to the firewall.
- Tier 3 may also be implemented in environments where highly critical information, databases, or other such assets require additional separation from the rest of the network.

11.2.3. Intrusion Detection

No matter how strong the security is compromises and incidents could occur. It is therefore very important to have a means of detecting such a breach in order to be able to respond to it and handle it in an appropriate manner. This should be implemented and available on both the computer and network level.

12. SUMMARY

The data center is a critical component of the operation of most companies these days. Many factors go into the design of a successful data center. Careful planning during the design phase will ensure a successful implementation resulting in a reliable and scalable data center which will serve then needs of the business for many years.