

Protect Yourself

Justin Bois

April 4, 2002

Abstract:

Many people go to great lengths to secure their network from the outside so that intruders cannot get in. However, if they do get in they are detected, hopefully, and dealt with as quickly as possible. Often times the same people that are such zealots about maintaining firewalls and software updates are incredibly lax about ensuring that the servers themselves are safe from direct attacks by people physically at the machine. Most often, this is because there are little or no physical security measures in place past simple perimeter defenses- locks and bolts for instance.

Designing a facility from the ground up with security in mind is an expensive proposition, but not terribly difficult. The ideas and practices for this can be carried over to pre-existing structures and institutions fairly easily if done on a scaled down level. This paper is intended to demonstrate the design of a building with physical security in mind and how to apply the same theories to existing buildings.

Body:

In this day and age many people worry about having their information stolen from right under their noses by hackers and other ill-mannered technically savvy thieves. This is a very real and valid concern. To counter this system administrators the world over have implemented a wide variety of counter measures. After all, this is a crime costing companies billions of dollars a year, why would they not spend a few hundred thousand dollars a year to ensure that they are up to date with the latest and greatest Intrusion Detection devices, along with Firewalls, honeypots, and tar pits?

They do- they spend millions of dollars to ensure that their networks are protected from the bad guys on the outside. They make sure that they can communicate efficiently and securely via corporate networks in order to maintain their competitive advantage. Then, after they install, bring on line, and tweak everything so it works beautifully, the engineers go home and go to sleep.

While they are sleeping, the janitorial service they hired has a new employee cleaning in their building. It is a larger building with several floors, so the cleaning crew is often split up and out of visual contact with one another. This new employee has the same set of keys that everyone in the crew does- it gets him or her into any spot they want. After all, they need to do a thorough job and earn their pay, right? Being an enterprising young person, they decide that the door over here is a bit dirty and could use a good cleaning. While they are cleaning the door, oops, they just happen to put the key in the lock and turn it. Well, since they are inside, they might as well look around and see what is in here, it might need cleaning too. So the new janitor wanders around the server room, wondering at all the pretty lights. After they realize that there is sensitive equipment in here, they leave and never pay it a second thought.

That is the nice version of the story. If that were all that happened then site security would be a far sight easier. Instead, that same employee could be a bitter former employee, a rival company's employee, or just a malicious hacker. In which case they could have done any number of things to your server room that you just spent large sums of money to protect. They could have done anything from simply disconnecting a computer from the network, to stealing equipment, stealing information, or just straight vandalism and the wanton destruction of your valuable equipment.

There is an obvious double standard at work here. We harden our networks to outside attackers, yet we let people that we do not even know have access to the very thing that we are trying to protect. One would think that if we were trying to protect something, we would protect it from those seeking to do it harm it both near and far. It is almost as if the closer the person gets to the server, the less we do to protect it.

One of the biggest reasons that people fail to secure their equipment physically, is that often the systems administrator is not asked how to run the office, just the network. People fail to realize that the paths the networks take out of the building are not the only ones that attacks can come in on. The fact that the network only uses wires and strands of fiber, does not mean that its only threats come in on the same mediums. As an industry, it is vital that we acknowledge that people seeking to do harm to our information infrastructure will not stop if they cannot get to us via the Internet. If what we hold is important enough, they will resort to physically attacking us. Keeping that in mind, where and how we set up our offices and data centers is vitally

important.

Ideally, we would get to design our facility from the ground up. We could choose where to build it, how to build it, and what to put inside. Were that the case, we would ensure that the place we chose to build was in a stable geographic region- while California is a lovely state, it can be rather prone to earthquakes, not something we want for our networks. Ensuring that the location of the site is relatively hazard free is one of the most basic, and most difficult things to do. The requirements for a building site are long and varied with the most major that are usually considered being:

- The area is not prone to natural disasters.
- The utilities are well established and hopefully redundant in their access to you.
- The climate is temperate and does not experience extremes or large changes.
- The surroundings are not prone to attacks (military, government, and nuclear are all candidates for attack).
- Emergency services are close.

Those will help maintain uptime, but other than the last two, they really do nothing for the security of the site. To really help with actual security from the ground up keep an eye out for these things:

- Clear lines of approach.
- Avoid places that have single methods of getting in or out of the facility.
- Avoid places that have seemingly infinite methods of getting in or out of the facility.

The proximity of emergency services is hopefully an absolute last resort. Hopefully, your fire suppression systems or security dealt with everything before the authorities showed up, and the only thing that they need to do is fill out a report (Friedman).

Many people argue with great fervor about having a separate building so that you are not subject to anybody else's problems should they occur. For example if you are in a large office complex and somebody else has a fire, you would not want to be forced to evacuate because of them, or worse yet have your sprinkler system go off as well. The fact that you are in a separate building makes it much easier to deny access to people that do not belong there. If you have to let people in to get to other offices in the building, you are letting them get that much closer. The object is to keep them as far away as possible. Another very good argument for a separate building lies in your ability to control every aspect about it. You don't rely on a third party for anything and have complete control over any company that you do have to deal with- utilities, alarm, and the myriad of others. An additional benefit of having a separate building can be found when looking at various building requirements. If you move into an office complex, you get what you are given. If you are moving into a building by yourself, you can dictate to a far greater degree just what you want and or need for building specifications.

Other people argue for combined buildings however. The theory runs that the more people and/or businesses are in a given location, the better the response from all involved will be for any problems. For example, when a city loses power who gets their power back first- the large apartment complexes or the single homes? The large apartment complexes of course; there are more paying customers there and more people that need electricity likewise with fire and police assistance. It is not necessarily a bad thing, but it does play into your favor if you are looking for the best possible help (Scheller).

Obviously, these arguments only hold weight if someone is attempting to perform a rather drastic denial of service attack- starting a fire, attacking the building physically, and the like. In actuality, the number of people willing to go to that sort of lengths is relatively few, and the amount of information worth that kind of action relatively concentrated. That level of hostility is usually reserved for military or government installations and those working on their behalf. More often however, the type of physical attack is not one of protecting yourself against armed assailants, but a single person after specific information. That person is generally not going to resort to violence, but is going to attempt to infiltrate the building in a manner that does not reveal his presence. As soon as most people find out that an intruder has been in somewhere they deem private, they usually take stock of things immediately and thoroughly to see what is missing and what is changed. That would make installing a Trojan or similar piece of malware more difficult as the chances of it being spotted are far greater. Because of this, the major concerns in physical site security lies more in your implementation of layout, design, and processes than the actual location of the building.

Layout is a basic concern that is the most easily implemented security measure. How you set up the inside of your office/data center is an important decision. The first and most obvious concern is that of external access to your data. Do not put important equipment in rooms with windows. Windows can be broken, forced, or accidentally left open. A second concern about external access would be ventilation systems. Air conditioning units can be a source of problems. While maintaining a climate-controlled environment is important for a server room, security is more vital. Air conditioning units can be either pulled through the window to the outside or pushed back into the server-room. Either method results in a large hole for an intruder to gain access to what your equipment; that is, if the intruder wants to gain access to the servers. There are reports of people having servers destroyed after the cleaning crew was washing out the air conditioners with hoses. Water ran into the server room and proceeded to run all over the server that was positioned directly below the unit. There is no reason a perpetrator could not do the same thing and flood your server room, cause a loss of both data and equipment as well as causing water damage to everything else in the room that is on the floor. All of this assumes of course that the intruder is going to use a hole in the walls that already exists. If the server room is on the ground floor and has an exterior wall, if the person is desperate enough, there is no reason a car or truck could not be used as a battering ram to go through the wall. While it is messy, inefficient, dangerous, and extremely reckless, it can get the job done and that is usually all the bad guys care about.

Keeping the servers away from outside access is vital. Almost as important, however is protecting the servers from illicit access from inside the building. As previously discussed, often a simple janitor can get into your server room with no outside help and no idea what he or she had access to. Limiting access to the server room is vital. Keep the servers in a separate room; keep them under lock and key. Having the servers in a separate room is useless if the room has no doors on it. At that point, they are merely stuck in a corner- hardly a valid method of defense. I have seen server rooms with two doors. One was kept locked whenever nobody was in there, a good thing. The second door however, fit the frame so poorly that the increased air pressure in the room caused by the ventilation system would actually open the door- the latch hardly even touched the frame. Generally, only one door should be created to access the server room. That door should be kept locked and closed at all times. Even if somebody is inside the room, why make it easier for the wrong person to walk in?

The location of the server room has been discussed as it relates to the exterior- keep it away from exterior walls and windows. What about its relationship to the rest of the interior? Server rooms should be kept in the middle of the office, surrounded by as much "stuff" as possible. The first reason is that if somebody does happen to break into the building, having it in the middle makes it take longer to reach the server room and their target. Anybody trying to walk in like they belong there during the day time has to walk through more people, hopefully at least one of which would ask who the person is and what they are doing. Another benefit of having the server room in the middle of the office surrounded by other offices and people is that it will make it harder for outsiders to monitor it using electronic means.

What kind electronic means would somebody use to spy on a computer that location would actually play a role in? If they cannot get in over the network, and cannot access the computer directly, then they cannot monitor anything about it, one would think. There are however always emanations from electronic equipment known as Electro Magnetic Radiation. Everything that carries a current of some sort produces this- everything from power lines and extension cords, to monitors and CPUs. These emanations are known as TEMPEST emissions. From this radiation it is possible from this radiation to do everything from see what is on people's CRT monitors, to actually monitoring what a computer does remotely. The biggest concern with TEMPEST lies in the fact that because of how the emanations occur, they are usually plain text and therefore once monitored, perfectly legible. This can be a huge concern as it could be extremely detrimental if somebody with the correct equipment, training, and motivation were to start monitoring for TEMPEST emanations. There are a couple of methods for combating TEMPEST emissions however. Just as you can shield against EMR coming in (solar flares and nuclear explosions are common sources), you can shield against that sort of energy being broadcast (Goodman).

The most thorough methods of shielding involve building with copper foil in the walls to make a continuous cocoon around the room. All seams must be soldered together, and doors have to have the same shielding. Door edges need to have special seams to ensure adequate conduction of current can occur, and the floor has to have a layer of steel in it to prevent any signal leakage occurring there. Everything that needs to enter or exit the room (wires, heating/cooling systems, etc.) have to run through special filters to help make sure that only clean signals that are intended to be leaving the facility are. Shielding a whole building for TEMPEST is usually prohibitively expensive. It is far more economical to merely shield a portion of the building- the server room. There is actually an advantage to this as well. As long as all the important information is processed in the server-room and the TEMPEST shielding is doing its job properly, then all the exposed computers and equipment outside the shielding act as an additional barrier to prevent people from monitoring EMI emissions. It would be similar to holding a whispered conversation in a football stadium when the home

team just scored a touchdown. You can hear each other, but all anybody else hears is the roar of the crowd, which contains no meaningful information (Department). What if you could read lips though?

That brings us to another form of TEMPEST, Optical TEMPEST. A recent development has shown that the nice LEDs that are in everything from our desktop PCs, to our Servers, to our switches and router can be monitored for the flickering of network traffic and actually decode when things are being sent and received. Of course, if we observe the previous building/locational criteria, there are no windows or other openings for criminals to spy on the LEDs, and hence renders that method of observation useless (Loughry). Another method of Optical TEMPEST utilizes the light given off from monitors and recreating the image displayed from that. Obviously if one can look directly at a monitor, the recreation process is fairly straightforward, however new research shows that even light that has reflected off of walls can be used to recreate monitor displays. While there are some things that cannot be avoided, placing monitors with their displays facing windows is something that is an obvious security measure that is easy to remedy. Ensuring that the displays of monitors face perpendicular to windows would be most beneficial, as it prevents the monitor from shining off the wall directly opposite a window, in addition it should cut down on glare reducing eye strain (Kuhn).

The final option for TEMPEST security is white noise jamming. The theory is that if you make enough noise, even an unshielded system is safer because of the amount of noise surrounding it. It takes the idea of placing the server-room in the center to help mask the TEMPEST emissions one step further. Instead of just letting other equipment make noise, you generate the noise intentionally, and "louder" to drown out any emissions that your equipment might really be making. Obviously, a combination of all of the above would be the most effective method of TEMPEST protection, however as stated earlier, this can be incredibly expensive, especially for entire buildings (Scheller).

TEMPEST shielded buildings protect against unintentional emissions, what about signals that were intended as broadcast signals that are therefore inherently a vulnerability? Wireless networks, while incredibly convenient to install can be a severe liability in terms of security. In an office situation, there is no reason that, with proper planning and implementation, a completely wired office would not yield the same mobility that a wireless network does. It may be slightly less convenient due to the requisite plugging and unplugging of cables, but it is almost ten times faster, and there are no broadcast signals to be monitored by people who just happen to wander by with a laptop and a wireless network card (Scheller).

One of the measures talked about above to help maintain a secure workplace was to lock the server-room door. Most people's first thought is lock and key in the standard sense of the word. A small chunk of metal inserted into another chunk, that when turned allows entry. While that is a very valid first line of defense, it should very rarely be your only one. Locks can be picked or broken, either of which render it useless. Of course, at this point many would point out that there are also combination locks that cannot be picked, as they require no physical key. Unfortunately, they are just as prone to a coerced entrance and anybody can guess a key sequence. It may take a while and it may cause many, many re-codings (if you change the code after people keep attempting guessed entries), but eventually it can and will be defeated.

To date those are the two major forms of keyed entry used across the globe. Keys are based upon one of three things- what you have, what you know, or what you are. A regular key and lock- like your car door is a perfect example of a "what you have" type of lock. You possess a key that, in theory, is unique to that lock. As long as you are the only one with that key, you are the only one that can operate that vehicle. As evident by the innumerable car thefts in the United States, this is not the case. A key based on what you know is best exemplified in the password you use to check your email. It is something that hopefully only you know. Because you are the only person with the correct password, you are the only one that can check your mail. Of course, passwords are easily stolen, cracked, or written down and lost, so there is a small issue with those as well. A "what you are" based lock is a lock only opens because you are you- in theory the same as the other two locks, but these locks require no passcodes or keys, just you. These are commonly referred to as biometric locks (Kessler).

Biometric locks take a unique identifier from a person and use that identifier to allow/disallow people past whatever the lock is protecting. The most common forms of biometric locks are fingerprint-scans, hand-scans, or retina/iris-scans. The obvious benefit to these types of controls is that it is very hard to forge most of these. Of all the above listed types, hand-scans are the least secure. While hand-scans are fast and easy, the hand geometry typically used does not lend itself enough points of reference to be able to do one-to-many searches (Hand). Fingerprint scans on the other hand are very good at this, however the rate of false rejections goes up due to more misreads on users fingers when attempting to gain entry. Some places have cited nearly a 50% rejection rate of users after having them entered into the system just weeks before (Finger). Iris scans on the other hand has variations in a person's iris factored into its algorithms to reduce false rejection rates. The

possibility of two irises being identical is somewhere around 1 in 10⁵². That is obviously a very, very slim possibility. A retinal scan on the other hand examines the retina, a nerve along the back of the eyeball, a place not visible to the naked eye (Iris). The retinal scans have a false rejection rate of somewhere in the neighborhood of 10%. In addition, there is an estimated 5-10% of the population that just cannot get a retina scan to work. A retinal scan is far more intrusive than any of the above listed methods. They require the subject to be no more than ½ and inch away from the capture device for the scanning to occur (Retina). An iris scan on the other hand requires you to be within 3 feet and takes less than 2 seconds to happen, start to finish. Iris scans are also affected by the "live-ness" of an iris. An iris separated from the body it used to be attached to are not allowed through (Scheller). While biometric locks are very good, they can be prohibitively expensive.

All the locks in the world however do nothing if everyone has a key to get in. One of the biggest issues with keys and similar items is that as soon as a person is hired, they get a set of keys that gets them in everywhere. Why does the new administrative assistant need access to the server room, why does the new intern need keys to the wiring closet that all the cables for the network run through? In short, they do not need and should have not have access. Much like permissions on a network share, give people as little as they need to effectively do their job. Of course, that does not mean that a single person should have the single key and that he or she spend their entire day giving other people access to secured areas. Keep track of who has keys and have a maximum number of keys to be given out at any given time. This helps keep traffic in and out of secured areas to a minimum as well as increases accountability. If only five people have access to a room and something happens, the list of people is far shorter than if everyone in a 50 person office has keys.

The last key control method to help keep a work area secure is to segregate key distribution. Do not give anybody keys to the entire facility. Give a few people some keys to get into A, B, C, and give others keys to X,Y,Z. Do not however give anyone all six keys. That way if anybody did want to get into all the restricted areas they either need an accomplice if working from within, or have to steal keys from multiple targets.

Key control is definitely a concern, however another extremely important thing that computerized locks using biometrics makes much easier, is tracking which key went where and when. Previously we talked about giving out a minimum number of keys with the idea that the fewer keys there were, the fewer people could get access to restricted zones. We now take that one step forward in that we keep track of just who is coming and going at all times. Keep track of who is at work and where they are. If the CEO of the company is in Bermuda working on his tan, there is no reason that his security code should be used to open the server-room door. A practice that goes hand in hand with this idea is the maintenance of up to date employee records. This goes beyond a simple check of who is on payroll, but keep track of employees that get keys and get fired. If you are unable to get that set of keys back, change the locks. Recently a story was related where an accounting company was put in financial straits as a disgruntled employee who never turned his keys in kept entering the building at night, logging on with his still active username and password- after a month of termination- and destroying the data that people were entering in the daytime (Schirling).

What happens if somebody does manage to sneak by all of your security measures and is attempting to exit the building with one of your servers on a handcart? What happens if for some unknown reason, your biometric scanners are not working and you need access to your server room? At this point, an external security force in the form of a real live breathing human being is called for. Actual security guards with patrols and areas to watch are both expensive and rarely needed. Security guards are more a deterrence than something that actually gets used. Most security guards never see anything more exciting than stray animal, however the one time that you do need them they pay for themselves ten-fold.

All the key control, biometrics, and office layouts do nothing if the people you have working for you are the same ones that will be stealing and selling your information and equipment to your competitors and on the street. If you cannot trust your employees then no amount of security is going to help you. The single largest security measure a company can take is unfortunately also the most difficult to do well and on a regular basis. Employee screening is vital if sensitive data is being handled. Spend the money on the employee's background check. If they have access to data that your company feels is vital, why would it not make sense to ensure that you have trustworthy employees? It is a proven fact that 90% of all attacks on a company's data network come from former disgruntled employees (Schirling). If that is the case ensuring that your employees are solid and stable would be more than a good thing to know, it could save your business if anything happens.

Summary:

All of this is well and good for a building whose foundation has yet to be poured, but how much of this is

immediately relevant to an existing office? Quite simply almost all of it. Biometric devices can be installed, key access can be altered to ensure that nobody has too much control, equipment can be moved away from air conditioners and out of rooms with windows, employees can still be screened, rooms can still be shielded. The only difference is scale and how tightly integrated everything is to start with. The first roll out of an iris scanner is sure to cause some hiccups at first, however once implemented, things are flowing just as smoothly as before. The biggest hurdle lies in the initial investment of time, planning, and money.

At first glance, much of this is overwhelming. Biometrics and TEMPEST shielding for a web server that hosts Ma and Pa's website? Probably not strictly necessary. While everything talked about does indeed address how to secure a site physically, it does not talk about when to use what method. This varies on a case by case basis. Some people may only need a regular lock and key to keep their one machine secure, because there is nobody that wants access to it. At the same time however, there may be a data center hosting information for insurance companies that needs armed guards on the premises at all times and iris scans for anybody attempting to enter the server room. Look at what you are trying to protect, and who you are trying to protect it against when considering the amount of security required. Then take it one step further just to be safe.

Works Cited:

Department of the Army. "Engineering and Design - Electromagnetic Pulse (EMP) and Tempest Protection for Facilities." 31 December 1990. US Army Corps of Engineers. 31 March 2002. <<http://www.usace.army.mil/inet/usace-docs/eng-pamphlets/ep1110-3-2/c-12.pdf>>.

"Finger-Scan Accuracy." 2001 finger-scan.com 1 April 2002. <http://finger-scan.com/finger-scan_accuracy.htm>.

Friedman, Seth. "Building the Ideal Web Hosting Facility: A Physical Security Perspective." 10 December 2001. SANS Institute. 31 March 2002. <<http://rr.sans.org/physical/facility.php>>.

Goodman, Cassi. "**An Introduction to TEMPEST.**" 18 April 2001. SANS Institute. 31 March 2002. <<http://rr.sans.org/encryption/TEMPEST.php>> .

Kessler, Gary. "Computer and Internet Security." Champlain College. Burlington VT, Spring 2002.

"Hand Scan Technology." 2001 hand-scan.com 1 April 2002. <<http://hand-scan.com/technology.htm>>

"Iris Recognition: The Technology ." 2001. iris-scan.com 1 April 2002. <http://iris-scan.com/iris_technology.htm>.

Kuhn, Markus G. "Optical Time-Domain Eavesdropping Risks of CRT Displays." 12 March 2002. 2002 IEEE Symposium on Security and Privacy. 31 March 2002. <<http://www.cl.cam.ac.uk/~mgk25/ieee02-optical.pdf>>.

Loughry, Joe and David A. Umphress. "Information Leakage from Optical Emanations." March 2002. ACM Transactions on Information and System Security. 31 March 2002. <http://applied-math.org/optical_tempest.pdf>.

"Retina Scan Technology." 2001 retina-scan.com 1 April 2002. <http://retina-scan.com/retina_scan_technology.htm>.

Schirling, Michael. Champlain College. Burlington VT, 3 April 2002.

Scheller, Matthew. Personal Interview. 3 April 2002