Physical Security in Mission Critical Facilities

By Suzanne Niles

White Paper #82

Revision 1

Executive Summary

Physical security — controlling personnel access to facilities — is critical to achieving data center availability goals. As new technologies such as biometric identification and remote management of security data become more widely available, traditional card-and-guard security is being supplanted by security systems that can provide positive identification and tracking of human activity in and around the data center. Before investing in equipment, IT managers must carefully evaluate their specific security needs and determine the most appropriate and cost-effective security measures for their facility. This paper presents an overview of the principles of personnel identification and describes the basic elements and procedures used in security systems.

Introduction

People: A Risk to be Managed

When data center security is mentioned, the first thing likely to come to mind is protection from sabotage, espionage, or data theft. While the need is obvious for protection against intruders and the intentional harm they could cause, the hazards from ordinary activity of personnel working in the data center present a greater day-to-day risk in most facilities.

People are essential to the operation of a data center, yet studies consistently show that people are directly responsible for 60% of data center downtime through accidents and mistakes — improper procedures, mislabeled equipment, things dropped or spilled, mistyped commands, and other unforeseen mishaps large and small. With human error an unavoidable consequence of human presence, minimizing and controlling personnel access to facilities is a critical element of risk management even when concern about malicious activity is slight.

Identification technology is changing as fast as the facilities, information, and communication it protects. With the constant appearance of new equipment and techniques, it's easy to forget that the age-old problem this technology is trying to solve is neither technical nor complicated: keeping unauthorized or ill-intentioned people out of places where they don't belong. And while the first step, mapping out the secure areas of the facility and defining access rules, may produce a layered and complex blueprint, it isn't intuitively difficult — IT managers generally know who should be allowed where. The challenge lies in the second step: deciding how best to apply less-than-perfect technologies to implement the plan.

Network-Critical Physical Infrastructure

Physical security is part of *Network-Critical Physical Infrastructure* (NCPI) because it plays a direct role in maximizing system availability ("uptime"). It does this by reducing downtime from accidents or sabotage due to the presence of unnecessary or malicious people.

Other NCPI elements are power, cooling, racks, cabling, and fire suppression.

Who Are You, and Why Are You Here?

While emerging security technologies may appear exotic and inscrutable — fingerprint and hand scans, eye scans, smart cards, facial geometry — the underlying security objective, unchanged since people first started having things to protect, is uncomplicated and familiar to all of us: getting a reliable answer to the question "Who are you, and why are you here?"

The first question — "Who are you?" — causes most of the trouble in designing automated security systems. Current technologies all attempt to assess identity one way or another, with varying levels of certainty — at correspondingly varying cost. For example, a swipe card is inexpensive and provides uncertain identity (you can't be sure who's using the card); an iris scanner is very expensive and provides very certain identity. Finding an acceptable compromise between certainty and expense lies at the heart of security system design.

The answer to the second question, "Why are you here?" — in other words, what is your business at this access point — might be implicit once identity has been established ("It's Alice Wilson, our cabling specialist, she works on the cables — let her in"), or it can be implemented in a variety of ways: A person's "who" and "why" can be combined — in the information on a swipe-card's magnetic strip, for example; a person's identity could call up information in a computer file listing allowable access; or there could be different access methods for various parts of the facility, designed to allow access for different purposes. Sometimes "Why are you here?" is the only question, and "Who are you?" doesn't really matter — as for repair or cleaning personnel.

Combining Expertise to Find the Solution

IT managers know the "who and why" of security for their installation, but they may not be conversant in the details of current methodologies or the techniques for applying them — nor should they need to be. They know their budget constraints, and they know the risks inherent in various types of security breach at their facility.

The security system consultant, on the other hand, doesn't know the particulars of the facility, but knows the capabilities, drawbacks, and cost of current methodologies. He or she also has experience in the design of other security systems, and so can help clarify, refine, or simplify the "who and why" requirements by asking the right questions.

With their combined expertise, a system can be designed that balances access requirements, acceptable risk, available methods, and budget constraints.

Defining the Problem

Secure Areas: What Needs Protecting?

The first step in mapping out a security plan is just that — drawing a map of the physical facility and identifying the areas and entry points that need different rules of access, or **levels of security**.

These areas might have concentric boundaries:

- Site perimeter
- Building perimeter
- Computer area
- Computer rooms
- Equipment racks

Or side-by-side boundaries:

- Visitor areas
- Offices
- Utility rooms

"Physical security" can also mean…

Physical security can also refer to protection from catastrophic damage (fire, flood, earthquake, bombing) or utility malfunction (power loss, HVAC failure).

Here it refers only to protection from on-site human intrusion.

Concentric areas can have different or increasingly stringent access methods, providing added protection called **depth of security**. With depth of security, an inner area is protected both by its own access methods and by those of the areas that enclose it. In addition, any breach of an outer area can be met with another access challenge at a perimeter further in.

Figure 1 *– Security Map Showing "Depth of Security"*

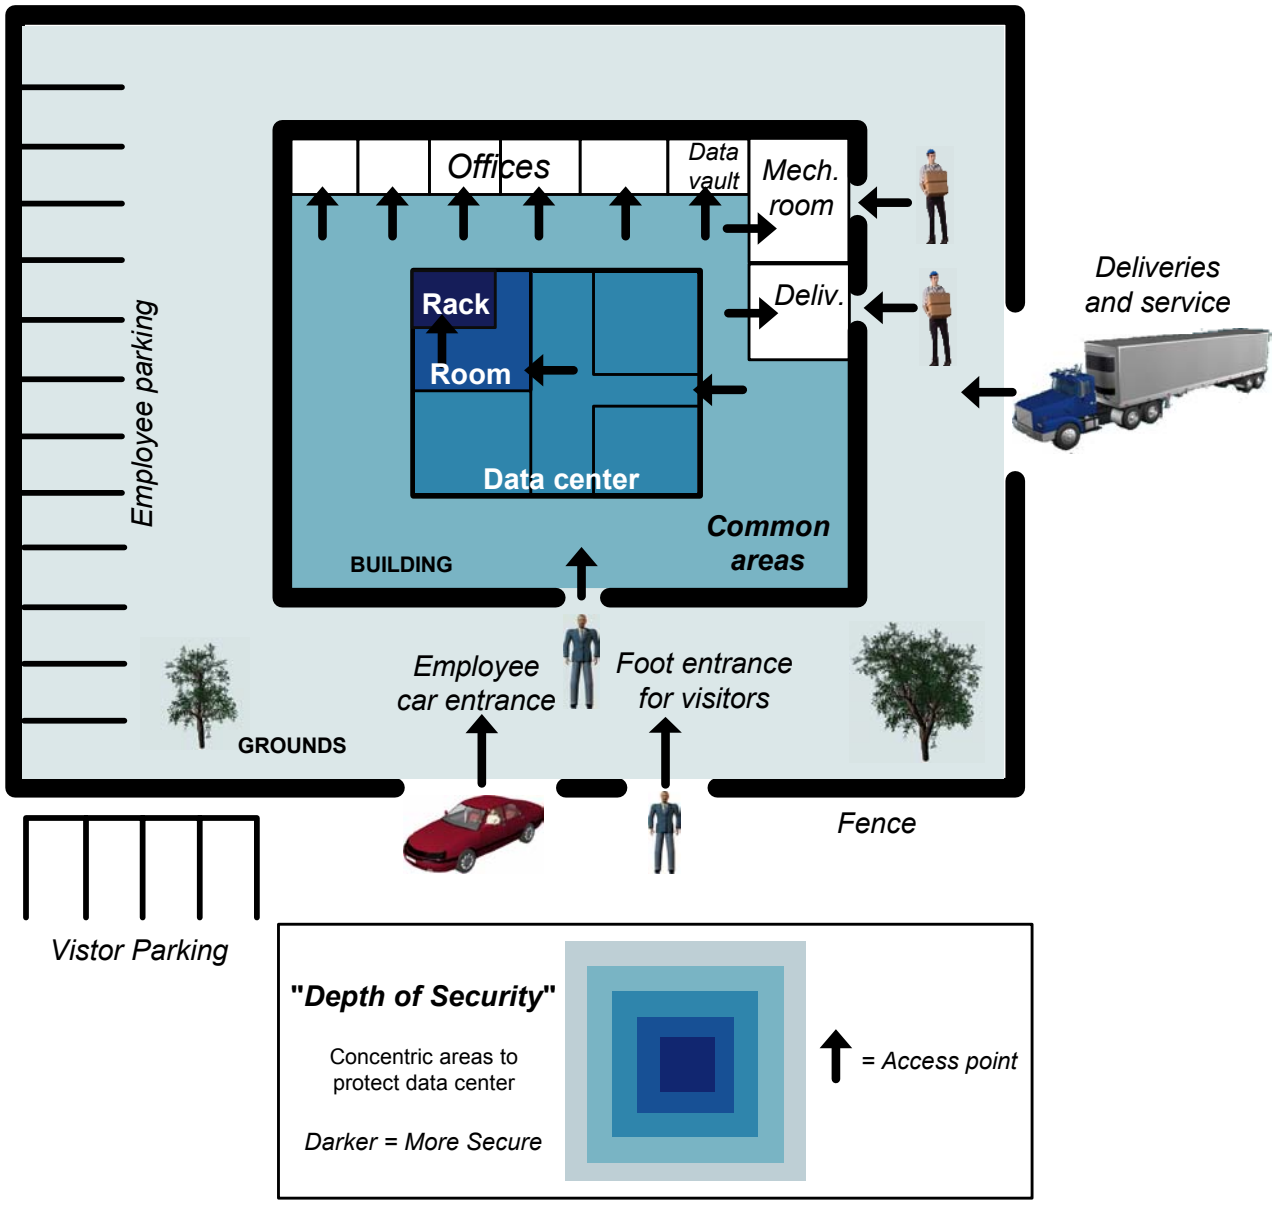

Rack-Level Security At the innermost "depth of security" layer — further in than the data room itself — is the *rack*. Rack locks are not in common use (yet), but if used they serve as the last defense against unauthorized access to critical equipment. It would be unusual for everyone in a room full of racks to have the need to access every rack; rack locks can ensure that only server people have access to servers, only telecommunications people have access to telecommunications gear, and so on. "Manageable" rack locks that can be remotely configured to allow access only when needed — to specific people at specific times — reduce the risk of an accident, sabotage, or unauthorized installation of additional gear that could cause a potentially damaging rise in power consumption and rack temperature.

Infrastructure Security It is important to include in the security map not only areas containing the functional IT equipment of the facility, but also areas containing elements of the physical infrastructure which, if compromised, could result in downtime. For example, HVAC equipment could be accidentally or deliberately shut down, generator starting batteries could be stolen, or a system management console could be fooled into thinking the fire sprinklers should be activated.

Access Criteria: Who is Allowed Where?

A person's authority for access to a secure area can be based on different things. Besides the usual ones — identity and purpose, the first two listed below — there may be additional categories requiring special treatment, such as "need to know."

Personal identity Certain individuals who are known to the facility need access to the areas relevant to their position. For example, the security director will have access to most of the facility but not to client data stored at the installation. The head of computer operations might have access to computer rooms and operating systems, but not the mechanical rooms that house power and HVAC facilities. The CEO of the company might have access to the offices of the security director and IT staff and the public areas, but not the computer rooms or mechanical rooms.

Separate the issues

Don't let the details of identification technologies intrude upon the initial mapping out of security requirements. First define the areas and the access criteria for your facility, *then* attack the cost/effectiveness/risk analysis, consider compromises, and figure out the best implementation of technology.

Reason to be there A utility repair person, regardless of whether it's Joe Smith or Mary Jones, might have access only to mechanical rooms and public areas. The cleaning crew, whose roster could change from day to day, might have access to common areas but nowhere else. A network switch expert might have access only to racks with switching equipment, and not racks with servers or storage devices. At a web server facility, a client's system maintenance personnel might have access only to a "client access room" where there are connections to their personal server for administrative purposes.

Need to know Access to extremely sensitive areas can be granted to specific people for a specific purpose — that is, if they "need to know," and only for as long as they have that need.

Applying the Technology

Methods of Identification: Reliability vs. Cost

Methods of identifying people fall into three general categories of increasing reliability — and increasing equipment cost:

- **What you have**
- **What you know**
- **Who you are**

What you have *Least reliable (can be shared or stolen)*

What you have is something you wear or carry — a key, a card, or a small object (a **token**) that can be worn or attached to a key ring. It can be as "dumb" as an old fashioned metal key or as "smart" as a card having an onboard processor that exchanges information with a reader (a **smart card**). It can be a card with a magnetic strip of information about you (such as the familiar ATM card); it can be a card or token having a transmitter and/or receiver that communicates with the reader from a short distance (a **proximity card** or **proximity token** — Mobil Speedpass® is an example).

What you have is the least reliable form of identification, since there is no guarantee it is being used by the correct person — it can be shared, stolen, or lost and found.

What you know *More reliable (can't be stolen, but can be shared or written down)*

What you know is a password, code, or procedure for something such as opening a coded lock, verification at a card reader, or keyboard access to a computer. A password/code presents a security dilemma: if it's easy to remember, it will likely be easy to guess; if it's hard to remember, it will likely be hard to guess — but it will also likely be written down, reducing its security.

What you know is more reliable than *What you have*, but passwords and codes can still be shared, and if written down they carry the risk of discovery.

Who you are *Most reliable (based on something physically unique to you)*

Who you are refers to identification by recognition of unique physical characteristics — this is the natural way people identify one another with nearly total certainty. When accomplished (or attempted) by technological means, it's called **biometrics**. Biometric scanning techniques have been developed for a number of human features that lend themselves to quantitative scrutiny and analysis:

- Fingerprint
- Iris (pattern of colors)
- Retina (pattern of blood vessels)
- Voice
- Hand (shape of fingers and thickness of hand)
- Face (relative position of eyes, nose, and mouth)
- Handwriting (dynamics of the pen as it moves)

Figure 2 *– What You have, What You Know, Who You Are*

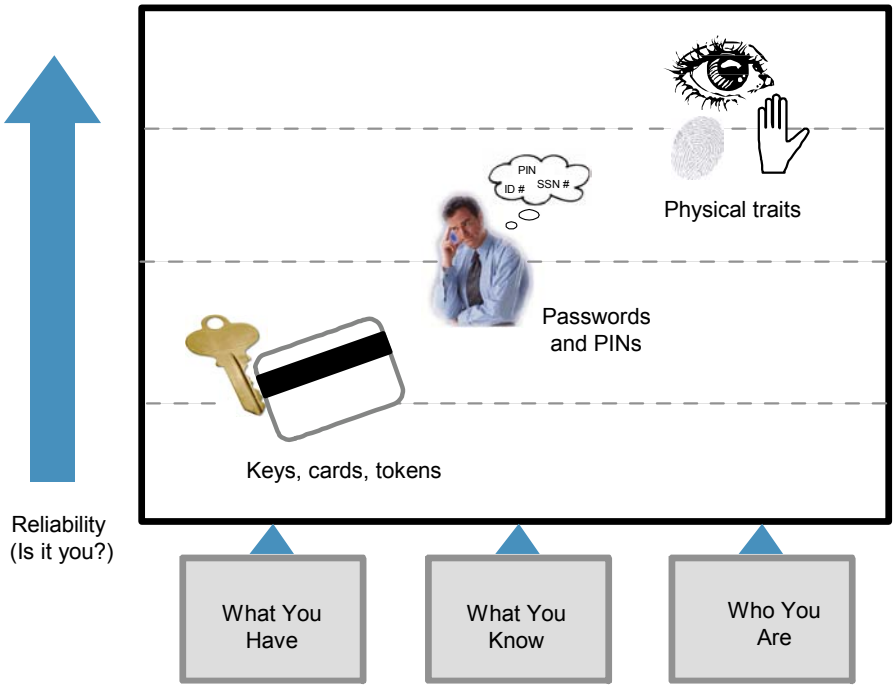

Biometric devices are generally very reliable, if recognition is achieved — that is, if the device thinks it recognizes you, then it almost certainly *is* you. The main source of unreliability for biometrics is not incorrect recognition or spoofing by an imposter, but the possibility that a legitimate user may fail to be recognized ("false rejection").

Combining Methods to Increase Reliability

A typical security scheme uses methods of increasing reliability — and expense — in progressing from the outermost (least sensitive) areas to the innermost (most sensitive) areas. For example, entry into the building might require a combination of swipe card plus PIN; entry to the computer room might require a keypad code plus a biometric. Combining methods at an entry point increases reliability at that point; using different methods for each level significantly increases security at inner levels, since each is secured by its own methods plus those of outer levels that must be entered first.

Why is it so complicated?

The reason security system design seems so complicated is this: We do not have the technology to quickly, easily, and cheaply determine a person's identity with certainty. What we have is an assortment of methods of varying effectiveness, convenience, and expense, resulting in difficult cost/effectiveness/risk analysis and the necessity of combining technologies or implementing concentric security perimeters for backup.

Security System Management

Some access control devices — card readers and biometric scanners, for example — can capture the data from access events, such as the identity of people who pass through and their time of entry. If network-enabled, these devices can provide this information to a remote management system for monitoring and logging (who's coming and going), device control (configuring a lock to allow access to certain people at certain times), and alarm (notification of repeated unsuccessful attempts or device failure).

Access Control Devices

Cards and Tokens: "What You Have"

Several types of cards and tokens are currently being used for access control, from simple to sophisticated, offering a range of performance on various dimensions:

- Ability to be reprogrammed
- Resistance to counterfeiting
- Type of interaction with card reader: swipe, insert, flat contact, no contact ("proximity")
- Convenience: physical form and how carried/worn
- Amount of data carried
- Computational ability
- Cost of cards
- Cost of reader

Regardless of how secure and reliable they may be due to their technology, the security provided by these physical "things" is limited by the fact that there is no guarantee the correct person is using them. It is therefore common to combine them with one or more additional methods of confirming identity, such as a password or even a biometric.

The **magnetic stripe card** is the most common type of card, with a simple magnetic strip of identifying data. When the card is swiped in a reader the information is read and looked up in a database. This system is inexpensive and convenient; its drawback is that it is relatively easy to duplicate the cards or to read the information stored on them.

The **barium ferrite card** (also called a "magnetic spot card") is similar to the magnetic stripe card but offers more security without adding significant cost. It contains a thin sheet of magnetic material with round spots arranged in a pattern. Rather than scanning or swiping, the card is simply touched to the reader.

The **Weigand card** is a variation of the magnetic stripe card. A series of specially treated wires with a unique magnetic signature is embedded in the card. When the card is swiped through the reader, a sensing coil detects the signature and converts it to a string of bits. The advantage of this complex card design is that the cards cannot be duplicated; the disadvantage is they cannot be reprogrammed either. With this technology the card need not be in direct contact with the reader; the head of the reader can therefore be encapsulated, making it suitable for outdoor installation. Unlike readers for proximity cards and magnetic-stripe cards, Weigand readers are not affected by radio frequency interference (RFI) or electromagnetic fields (EMF). The robustness of the reader combined with the difficulty in duplicating the card makes the Weigand system extremely secure (within the limits of a "what you have" method), but also more expensive.

The **bar-code card** carries a bar code, which is read when the card is swiped in the reader. This system is very low-cost, but easy to fool — an ordinary copy machine can duplicate a bar code well enough to fool a bar-code reader. Bar-code cards are good for minimum-security requirements, especially those requiring a large number of readers throughout the facility or a large volume of traffic traversing a given access point. This is not so much a security system as it is an inexpensive access *monitoring* method. (It has been said that bar-code access only serves to "keep out the honest people.")

The **infrared shadow card** improves upon the poor security of the bar-code card by placing the bar code between layers of PVC plastic. The reader passes infrared light through the card, and the shadow of the bar code is read by sensors on the other side.

The **proximity card** (sometimes called a "prox card") is a step up in convenience from cards that must be swiped or touched to the reader. As the name implies, the card only needs to be in "proximity" with the reader. This is accomplished using RFID (radio frequency identification) technology, with power supplied to the card by the card reader's electromagnetic field. The most popular design works within a distance of about 10 cm. (four inches) from the reader; another design — called a **vicinity card** —works up to about a meter (three feet) away.

The **smart card**, the most recent development in access control cards, is rapidly becoming the method of choice for new installations. It is a card with a built-in silicon chip for onboard data storage and/or computation. Data is exchanged with the reader either by touching the chip to the reader (*contact* smart card) or by interacting with the reader from a distance, using the same technology as proximity and vicinity cards (*contactless* or *proximity* smart card). The chip, which is about a half inch in diameter, doesn't necessarily have to be on a card — it can be attached to a photo ID, mounted on a key chain, or worn as a button or jewelry (such as the iButton® token). The general term for objects that carry such a chip is *smart media.*

Smart cards offer a wide range of flexibility in access control. For example, the chip can be attached to older types of cards to upgrade and integrate with pre-existing systems, or the cardholder's fingerprint or iris scan can be stored on the chip for biometric verification at the card reader — thereby elevating the level of identification from "what you have" to "who you are." Contactless smart cards having the "vicinity" range offer nearly ultimate user convenience: half-second transaction time with the card never leaving the wallet.

Keypads and Coded Locks: "What You Know"

Keypads and **coded locks** are in wide use as a method of access control. They are reliable and very user-friendly, but their security is limited by the sharable and guessable nature of passwords. They have familiar phone-like buttons where users punch in a code — if the code is unique to each user it's called a personal access code (PAC) or personal identification number (PIN). *Keypad* generally implies the ability to accept multiple codes, one for each user; *coded lock* usually refers to a device having only one code that everyone uses.

The security level of keypads and coded locks can be increased by periodically changing codes, which requires a system for informing users and disseminating new codes. Coded locks that don't have their code changed will need to have their keypad changed periodically if a detectable pattern of wear develops on the keys. As with access cards, keypad security can be increased by adding a biometric to confirm user identity.

Biometrics: "Who You Are"

Biometric technology is developing fast, getting better and cheaper. High confidence affordable biometric verification — especially fingerprint recognition — is entering the mainstream of security solutions. Many vendors now supply a wide range of biometric devices, and when combined with traditional "what you have" and "what you know" methods, biometrics can complement existing security measures to become best practice for access control.

Why not use *just* a biometric?

Q: If an entry point uses card, PIN, plus biometric, why not use just the biometric alone if biometrics are so reliable?
A: Because (1) Biometric processing time can be unacceptable if a large database of user scans must be searched instead of comparing to the scan of single user, and (2) The risk of biometric false rejection or acceptance can be reduced if the scan is compared to only one user in the database.

While biometric traits are nearly impossible to forge, there is still the risk of incorrect matches by the technology.

Biometric identification is typically used not to *recognize* identity by searching a database of users for a match, but rather to *verify* identity that is first established by a "what you have" or "what you know" method — for example, a card/PIN is first used, then a fingerprint scan verifies the result. As performance and confidence in biometric technology increase, it may eventually become a stand-alone method of *recognizing* identity, eliminating the need to carry a card or remember a password.

There are two types of failures in biometric identification:

> ***False rejection*** — Failure to recognize a legitimate user. While it could be argued that this has the effect of keeping the protected area extra secure, it is an intolerable frustration to legitimate users who are refused access because the scanner doesn't recognize them.
>
> ***False acceptance*** — Erroneous recognition, either by confusing one user with another, or by accepting an imposter as a legitimate user.

Failure rates can be adjusted by changing the threshold ("how close is close enough") for declaring a match, but decreasing one failure rate will increase the other.

Figure 3 *– Hand Scanner*

Considerations in choosing a biometric capability are equipment cost, failure rates (both false rejection and false acceptance), and *user acceptance*, which means how intrusive, inconvenient, or even dangerous the procedure is perceived to be. For example, retinal scanners are generally considered to have low user acceptance because the eye has to be 1-2 inches from the scanner with an LED directed into the eye.

Other Security System Elements

Security system design focuses on devices to identify and screen individuals at entry points — "access control" — which is all you would need *if* there were 100% reliability of identification, total trustworthiness of the intentions of people admitted, and the physical perfection of unbreakable walls, doors, windows, locks, and ceilings. To cover for inevitable failings due to flaws or sabotage, security systems ordinarily incorporate additional methods of protection, monitoring, and recovery.

Building Design

When building a new facility or renovating an old one, physical security can be addressed from the ground up by incorporating architectural and construction features that discourage or thwart intrusion. Security considerations in the structure and layout of a building generally relate to potential entry and escape routes, access to critical infrastructure elements such as HVAC and wiring, and potential sources of concealment for intruders. See the appendix for a list of some of these design considerations.

Piggybacking and Tailgating: Mantraps

A common and frustrating loophole in otherwise secure access control systems can be the ability of an unauthorized person to follow through a checkpoint behind an authorized person (called **piggybacking** when the authorized person is complicit — i.e., holds the door — or **tailgating** if the unauthorized person slips through undetected). The traditional solution is an airlock-style arrangement called a **mantrap** having doors at entry and exit, with room for only one person in the space between the doors. Mantraps can be designed with access control for both entry and exit, or for exit only — in which case a failed attempt to exit the enclosure causes the entry door to lock and an alert to be issued indicating that an intruder has been caught. A footstep-detecting floor can be added to confirm there is only one person passing through.

A new technology for solving this problem uses an overhead camera for optical tracking and tagging of individuals as they pass, issuing an alert if it detects more than one person per authorized entry.

Camera Surveillance

Still cameras can be used for such things as recording license plates at vehicle entry points, or in conjunction with footstep sensors to record people at critical locations.

Closed circuit TV (CCTV) cameras — hidden or visible — can provide interior or exterior monitoring, deterrence, and post-incident review. Several types of camera views can be used — fixed, rotating, or remotely controlled. Some things to consider when placing cameras:

- Is it important that a person in camera view be easily identifiable?
- Is it only necessary to determine if the room is occupied?
- Are you watching to see if assets are being removed?
- Is the camera simply to serve as a deterrent?

If CCTV signals are recorded, there must be procedures in place to address the following issues:

- How will tapes be indexed and cataloged for easy retrieval?
- Will the tapes be stored on site or off site?
- Who will have access to the tapes?
- What is the procedure for accessing tapes?
- How long will the tapes be kept before being destroyed?

New technology is in development to automate a job traditionally done by security guards — watching TV monitors — by software detection of changes (movement) in the image on the screen

Security Guards

Despite all the technological advancements in the field of physical security, experts agree that a quality staff of protection officers tops the list of methods for backing up and supporting access control. Guards provide the surveillance capability of all the human senses, plus the ability to respond with mobility and intelligence to suspicious, unusual, or disastrous events.

The International Foundation for Protection Officers (IFPO) is a non-profit organization founded for the purpose of facilitating standardized training and certification of protection officers. Their *Security Supervisor Training Manual* is a reference guide for protection officers and their employers.

Sensors and Alarms

Everyone is familiar with traditional house and building alarm systems and their sensors — motion sensors, heat sensors, contact (door-closed) sensors, and the like. Data center alarm systems might use additional kinds of sensors as well — laser beam barriers, footstep sensors, touch sensors, vibration sensors. Data centers might also have some areas where a silent alarm is preferred over an audible one in order to catch perpetrators “in the act.”

If the sensors are network-enabled, they can be monitored and controlled remotely by a management system, which could also include personnel movement data from access-control devices (see earlier section, **Security System Management**.)

Visitors

Handling of visitors must be considered in any security system design. Typical solutions are to issue temporary badges or cards for low-security areas, and to require escorting for high security areas. The presence of mantraps (to prevent two people from passing an entry point with one authorization) would require a provision for a temporary override or for issuance of visitor credentials to allow passage.

The Human Element

Technology can't do the job all by itself, particularly since we are calling upon it to perform what is essentially a very human task: assessing the identity and intent of people. While people are a significant part of the security problem, they are also part of the solution — the abilities and fallibilities of people uniquely qualify them to be not only the weakest link, but also the strongest backup.

People: The Weakest Link

In addition to mistakes and accidents, there is inherent risk in the natural human tendency toward friendliness and trust. A known person entering the facility could be a disgruntled employee or a turncoat; the temptation to bend rules or skip procedures for a familiar face could have disastrous consequences; a significant category of security breach is the "inside job." Even strangers can have surprising success overcoming security — the ability of a clever stranger to use ordinary guile and deceit to gain access is so well documented that it has a name: **social engineering**. Anyone in an area where harm could be done must be well trained not only in operational and security protocols, but also in resistance to creative social engineering techniques.

People: The Strongest Backup

Protection from a security breach often comes down to the recognition and interpretation of unexpected factors — a skill in which technology is no match for alert people. Add an unwavering resistance to manipulation and shortcuts, and human presence can be a priceless adjunct to technology.

Beyond an alert staff, the incomparable value of human eyes, ears, brains, and mobility also qualifies people for consideration as a dedicated element in a security plan — the old-fashioned security guard. The presence of guards at entry points and roving guards on the grounds and inside the building, while expensive, can save the day when there is failure or hacking of technological security. The quick response of an alert guard when something "isn't right" may be the last defense against a potentially disastrous security breach.

In protecting against both accidental and deliberate harm, the human contribution is the same: constant vigilance and strict adherence to protocols. Having kept out all but those essential to the operation of the facility, the remaining staff — well trained, following well-designed practices and procedures — are the final firewall of an effective physical security system.

Choosing the Right Solution: Risk Tolerance vs. Cost

The right security system is a best-guess compromise that balances the risk and potential damage from people being in the wrong place against the expense and nuisance of security measures to keep them out.

Potential Cost of a Security Breach

While each data center has its own unique characteristics and potential for loss, most will have something to consider in these general categories:

Physical loss — Damage to rooms and equipment from accidents, sabotage, or outright theft.

IT productivity loss — Diversion of staff from primary duties while equipment is repaired or replaced, data is reconstructed, or systems are cleared of problems.

Corporate productivity loss — Interruption of business due to downtime.

Information loss — Loss, corruption, or theft of data.

Loss of reputation and customer goodwill — Consequences from serious or repeated security breaches: loss of business, drop in stock value, lawsuits.

Considerations in Security System Design

Security system design can be a complicated equation with many variables. While specific strategies for security system design are beyond the scope of this paper, any design will likely consider these issues:

Cost of equipment — Budget constraints ordinarily limit the extensive use of high-confidence identification equipment. The usual approach is to deploy a range of techniques appropriate to various security levels.

Combining of technologies — The reliability of identification at any level can be increased by combining lower-cost technologies, with the innermost level enjoying the combined protection of all the outer concentric perimeters that contain it.

User acceptance — (The "nuisance" factor.) Ease of use and reliability of identification are important in preventing the system from becoming a source of frustration and a temptation for subversion.

Scalability — Can the design be implemented incrementally as necessity, funding, and confidence in the technology increase?

Backwards compatibility — Is the new design compatible with elements of an older system already in place? Keeping all or part of an existing system can significantly reduce deployment cost.

You can't buy your way out

Even if expense were of no concern, blanketing the facility with highest security would, in most cases, be unacceptably intrusive and inconvenient. Each area to be protected must be realistically assessed for security needs based on what's in it and who needs access.

Figure 4 *– Balancing potential loss against known cost of security*

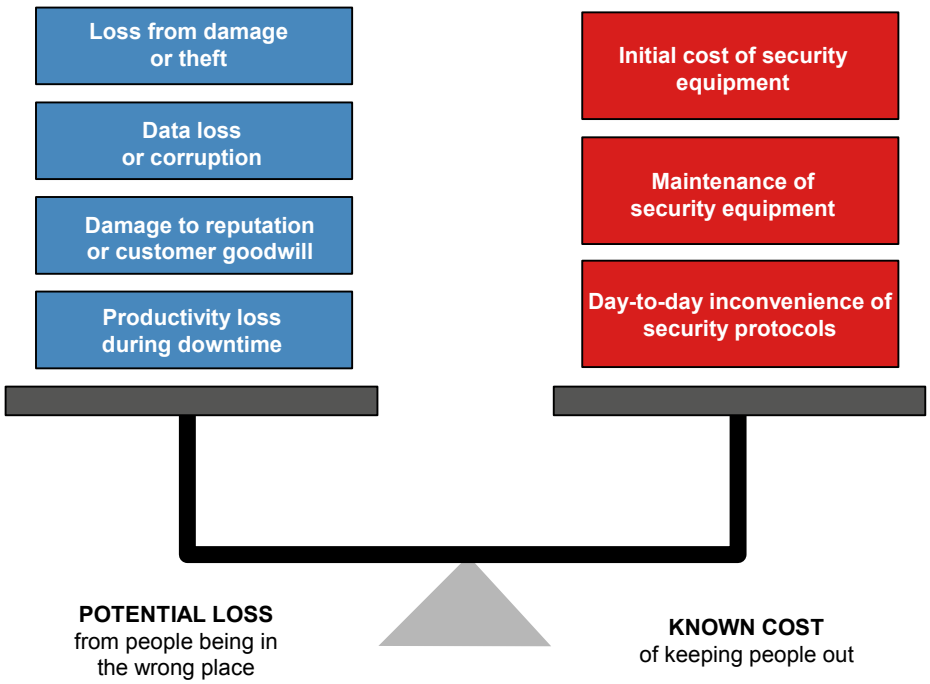

Conclusion

As data centers and web hosting sites proliferate, the need for physical security at the facility is every bit as great as the need for cybersecurity of networks. Intruders who falsify their identity or intentions can cause enormous damage, from physically disabling critical equipment to launching a software attack at an unsecured keyboard. Even the ordinary mistakes of well-intentioned staff pose a significant daily threat to operations, and can be minimized by restricting access to only the most essential personnel.

Technologies are in place, and getting less expensive, to implement broad range solutions based on the identification principles of **What you have**, **What you know**, and **Who you are**. By combining an assessment of risk tolerance with an analysis of access requirements and available technologies, an effective security system can be designed to provide a realistic balance of protection and cost.

About the Author

Suzanne Niles is a white paper author at APC's Engineering Design Center in Billerica, Massachusetts. She studied mathematics at Wellesley College before receiving a Bachelor's degree in computer science from MIT, with a thesis on handwritten character recognition. Suzanne has been documenting and explaining to diverse audiences for over 25 years, in a variety of media from software manuals to photography and children's songs. Prior to joining APC in 2004, she was book editor for The Village Group, where she edited Wes Kussmaul's new book, *Quiet Enjoyment*, about security and identity in the Internet age.

Appendix

Security Considerations in Building Design

When building a new facility or renovating an old one, physical security can be addressed from the ground up by incorporating architectural and construction features that discourage or thwart intrusion. Security considerations in the structure and layout of a building generally relate to potential entry and escape routes, access to critical infrastructure elements such as HVAC and wiring, and potential sources of concealment for intruders.

For security considerations in site selection, see APC White Paper #81, "Site Selection for Mission Critical Facilities."

- Position the data center door in such a way that only traffic intended for the data center is near the door.
- Use steel doors and frames, with solid doors instead of hollow-core. Make sure that hinges cannot be removed from the outside.
- Data center walls should use materials sturdier than the typical sheet rock used for interior walls. Sensors can be imbedded in the walls to detect tampering.
- The room used for the data center should not abut any outside walls.
- Allow long and clear lines of sight for any security stations or cameras within the data center.
- Make use of barriers to obstruct views of the entrances and other areas of concern from the outside world. This prevents visual inspection by people who wish to study the building layout or its security measures.
- Be aware of the placement of ventilation ducts, service hatches, vents, service elevators and other possible openings that could be used to gain access. Tamper-proof grills should be installed on all such openings that exceed 12 inches in width, to prevent human entry.
- Avoid creating spaces that can be used to hide people or things. For example, the space beneath raised floors could be a hiding place. Make sure that potential hiding places are secured and not easily noticed by someone walking through the facility.
- Install locks and door alarms to all roof access points so that security is notified immediately upon attempted access. Avoid points of entry on the roof whenever possible.
- Take note of all external plumbing, wiring, HVAC, etc., and provide appropriate protection. If left in plain site or unprotected, these infrastructure components can be used to sabotage the facility without having to disable security measures.

- Eliminate access to internal runs of wire, plumbing and ventilation ducts inside the facility. You may have a data center thoroughly secured, but if a person walking down a corridor can gain access to a run of power cabling or data cabling, the data center is compromised.

- Consider the placement of the data center within the building when retrofitting an existing facility or constructing a new data center within an existing structure. Avoid vulnerable locations or man-made risks. For example, avoid placing a data center underneath or adjacent to kitchen facilities, manufacturing areas with large machinery, parking lots, or any area with frequent traffic or vehicular access. Anything from kitchen fires to car bombs to traffic accidents can pose a threat.

- Protect the central security monitoring station by enclosing it with bulletproof glass.

- If the data center is housed in its own building, keep the exterior of the building plain. Do not use identifying marks such as company names or logos that would imply that a data center lies within.

- Use concrete bollards or other obstructions to prevent unwanted vehicles from getting any closer than a predetermined distance from the building.

Glossary

Terms that appear in **bold** are defined in this glossary.

access control

Controlling entry of people into buildings, rooms, and racks, and controlling the use of keyboards and equipment, by the use of automated devices that either read information stored on an object such a card (**what you have**), receive a code or password (**what you know**), or recognize a physical trait by biometric analysis (**what you are**).

access point

A place along the perimeter of a secure area where there is a door and some type of **access control** method to screen users attempting entry to the area.

availability

A calculated prediction of a network's percentage of "uptime." For mission-critical facilities, the goal is "five nines" or 99.999% – less than 5 minutes of downtime per year.

bar-code card

A type of **access control** card that uses a bar code to store information; read by swiping through a reader.

barium ferrite card

A type of **access control** card that uses a pattern of magnetic spots to store information; read by laying flat on a reader. Also called a "magnetic spot card."

biometric lock

A lock that is controlled by a biometric scanner.

biometrics

Establishing personal identity using technology to measure a physical or behavioral trait – for example, a fingerprint.

cipher lock

A lock that is opened by pressing its buttons in a specific sequence. It differs from a **coded lock** in that it typically has only 4-5 buttons, and each button can only be pressed once. The cipher lock, with metal buttons, was the mechanical precursor of today's electronic coded lock with a telephone-like keypad.

coded lock

A lock that is opened by typing a code on a keypad.

contact smart card

A **smart ca**rd that must make contact with the reader. Compare with **contactless smart card**.

contactless smart card

A **smart card** that uses **RFID** technology to enable its use without physical contact with the reader. Maximum distance from the reader is either the **proximity** range (10 cm. / 4 inches) or the **vicinity** range (one meter / 3 feet) depending upon which of two RFID standards is used.

depth of security

Concentric perimeters of security having different or increasingly stringent access methods. An inner area is protected both by its own access methods and by those of the areas that enclose it and must therefore be entered first.

facial geometry

One of the physical traits that can be measured by biometric technology – the relative position of eyes, nose, and mouth on the face.

false acceptance

In biometric identification, the erroneous result of identifying someone who isn't in the database of known people. It is one of two ways biometric identification can fail; the other is **false rejection**.

false rejection

In biometric identification, the erroneous result of failure to recognize a known person. It is one of two ways biometric identification can fail; the other is **false acceptance**.

FAR

False Acceptance Rate. For a biometric device, the percentage of readings that are a **false acceptance.**

FRR

False Rejection Rate. For a biometric device, the percentage of readings that are a **false rejection**.

hand scan

A technique for biometric identification that measures three-dimensional hand geometry – the shape of the fingers and the thickness of the hand.

iButton®

A microchip similar to those used in a **smart cards** but housed in a round stainless steel button about a half-inch in diameter, which can be attached to a key fob or jewelry. iButtons are extremely rugged, but (as of May 2004) are not available with **RFID** technology for contactless use.

IFPO

International Foundation for Protection Officers. A non-profit organization founded for the purpose of standardized training and certification of protection officers. Its *Security Supervisor Training Manual* is a reference guide for protection officers and their employers.

infrared shadow card

A type of **access control** card that has a bar code sandwiched between two layers of plastic. The reader passes infrared light through the card, and the shadow of the bar code is read by sensors on the other side.

iris scan

A technique for biometric identification that maps the pattern of colors in the iris of the eye.

levels of security

The range of security protection, low to high, provided at concentric perimeters – the least secure at the outermost perimeter (such as entry to the building) and the most secure at the innermost perimeter (such as access to a rack).

magnetic stripe card
magstripe card

A type of **access control** card that uses a magnetic strip to store information; read by swiping through a reader.

manageable

Able to be monitored and controlled remotely. Manageable **access control** devices can communicate with a remote management system for *monitoring* (who's coming and going and when), *control* (configuring the device to allow access to certain people at certain times), and *alarm* (notification of repeated unsuccessful access attempts or device failure).

management

Automated communication with remote devices for monitoring, control, and alarm. Traditionally called "building automation" or "household automation," the new term *management* refers to network-based communication with all elements of a data center, including both the IT equipment itself (servers, storage devices, telecommunications, and network devices) and the physical infrastructure (power, cooling, fire protection, and security).

mantrap

An airlock-style arrangement having secured doors for entry and exit, with room for only one person between the doors. It is a solution to the security loophole called **piggybacking** or **tailgating**, in which an unauthorized person freely passes a security checkpoint by following an authorized person through an open door.

NCPI

Network-Critical Physical Infrastructure. Elements of a data center's *physical* infrastructure (as distinguished from IT infrastructure such as routers and storage managers) that contribute directly to **availability** by ensuring uninterrupted operation. NCPI includes power, cooling, fire suppression, and **physical security**.

need to know

A very high level of security, with access restricted to people who have a specific, immediate need to be in the secured area (for access to particular data, for example), with access only allowed for the time period during which that need exists.

Network-Critical Physical Infrastructure – see NCPI

PAC

Personal Access Code. Another name for PIN (Personal Identification Number) – a code or password that identifies a user at an **access point**.

physical security

Protecting physical facilities from accidents or sabotage caused by the presence of unauthorized or ill-intentioned people. A physical security system always includes **access control** devices for automated screening at entry points, plus a sensor-based alarm system. Additional protection may include camera surveillance and security guards. (*Physical security* is sometimes used in a more general way to refer to protection from all kinds of physical damage including weather, earthquakes, and bombing. In this paper it refers only to protection from trouble caused by unauthorized *people* inside the facility.)

piggybacking

The security breach that occurs when an authorized person, having unlocked a door using legitimate credentials, holds the door open for an unauthorized person to follow through the checkpoint with no credentials. (A similar breach is **tailgating**, where the unauthorized user slips through undetected behind the authorized user.)

prox card
proximity card

An **access control** card that has an onboard **RFID** transmitter/receiver, allowing it to communicate with a reader from a distance of up to one meter (3 feet).

proximity smart card

A **smart card** that has **RFID** technology in its chip, so that it can communicate with the reader from a distance of up to 10 cm. (4 inches). Also called a **contactless smart card**.

retinal scan

A technique for biometric identification that maps the pattern of blood vessels in the retina of the eye.

RFID

Radio frequency identification. Communication between card and reader without physical contact. RFID technology is what makes **proximity cards**, **vicinity cards**, and **contactless smart cards** work. The RFID chip is powered by an electromagnetic field from the reader, and so does not need a battery.

smart card

A type of **access control** card that stores information in a microchip. The chip not only stores data, but can perform computation and exchange data with the reader. It is read by touching the card to the reader so that the electrical the contacts line up. See also **contactless smart card**.

smart media

Small objects of any shape that contain the same type of chip used in a **smart card**. Smart media are typically small objects (**tokens**) that can be attached to a key ring or worn as jewelry.

social engineering

The use of ordinary guile and deceit to con people into relaxing security procedures – for example, such as revealing passwords, lending keys, or opening doors.

tailgating

The security breach that occurs when an unauthorized person slips past a checkpoint undetected, by following an authorized user through an open door. (A similar breach is **piggybacking**, where the authorized user is complicit and holds the door open.)

template

In **biometrics**, a computed transformation of a scan – still unique to the individual but taking up much less storage. It is the template, not the raw scan, that is stored in a database of users or on the chip of a **smart card**, for comparison to a live scan taken at an **access point**.

threshold

In **biometrics**, the user-adjustable parameter that can be used to adjust the two failure rates (**false acceptance**

and **false rejection**). Since it represents "How close is close enough?" decreasing one of the failure rates automatically increases the other.

token

A small object with a microchip that carries your personal identifying information. The token is touched to a reader, or simply brought within range if it includes **RFID** capability.

vicinity card

An **access control** card that has an onboard **RFID** transmitter/receiver, allowing it to communicate with a reader from a distance of up to one meter (3 feet).

voice print

In **biometrics**, a digital representation of a user's voice used for comparison with the user's live speech at an **access point**.

Weigand card

A type of **access control** card that uses specially treated and magnetized imbedded wires to hold information; read by swiping through a reader.

what you have

In **access control**, any method of identification based on an object in your possession, such as a card or **token**. It is the least secure category of identification because there is no guarantee that the object is being used by the intended person.

what you know

In **access control**, any method of identification based on something that you know, such as a numeric code or a password. It is more secure than **what you have**, but can still be told to someone else, or written down and discovered.

who you are

In **access control**, any method of identification based on a biological or behavioral trait unique to you. It is the most secure category of identification because it very difficult to forge such a trait, but it is not 100% reliable because of the risk of errors in reading or interpretation. Another name for this type of identification is **biometrics**.