Wetware’s Intrusion Prevention Systems: Defending Against Social Engineering

Gartner IT Security Summit 2005

Ant Allan

6–8 June 2005
Marriott Wardman Park Hotel
Washington, District of Columbia

Millennia of Social Engineering

12th Century BCE

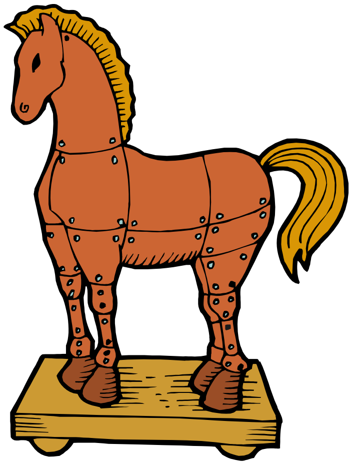

2005 CE

~~f~~ ph

Despite all of our technical advances, it’s far easier to fool a fool than circumvent some security defenses. Phishing, viruses (but not worms) and scams are rampant on the Internet, and for every new technology defense there’s a new trick to get information or cash out of even intelligent individuals.

The latest rounds of social engineering attacks on the Internet are increasingly sophisticated as the cybercriminals improve their technical skills. In the past year we’ve seen more — and more sophisticated — phishing attacks, which started nearly a decade ago on AOL as a technique for stealing usernames and passwords over instant messenger. But even “simple” classic attacks — like the Nigeria 419 scam that started through snail-mail even before faxes — are still widespread.

The basic concepts of human manipulation never really change, but in this evolving electronic world, thieves and tricksters take advantage of our unfamiliarity with this new global, environment. With millions of e-mail addresses at their disposal, then need fool only a fraction of a percent for a profitable venture.

Social engineering can lead to many damaging security compromises, and the threat shows no sign of lessening. As enhanced security infrastructures make technical attacks against the network less likely to succeed and less damaging, the focus of attack will shift to the organization's employees. By doing this, the attacker will aim to:

- Gain access to restricted areas or equipment.
- Gather logon information and passwords, or otherwise gain access to corporate systems .
- Discover confidential, proprietary, or personal information.
- Trick victims into opening e-mail or downloading attachments that, in turn, admit malicious code.

Client Issues

1. What is social engineering and how does it work?
2. How can a "wetware intrusion prevention system" defend against social engineering?
3. How can you build a "wetware Intrusion prevention system"?

Client Issue: What is social engineering and how does it work?

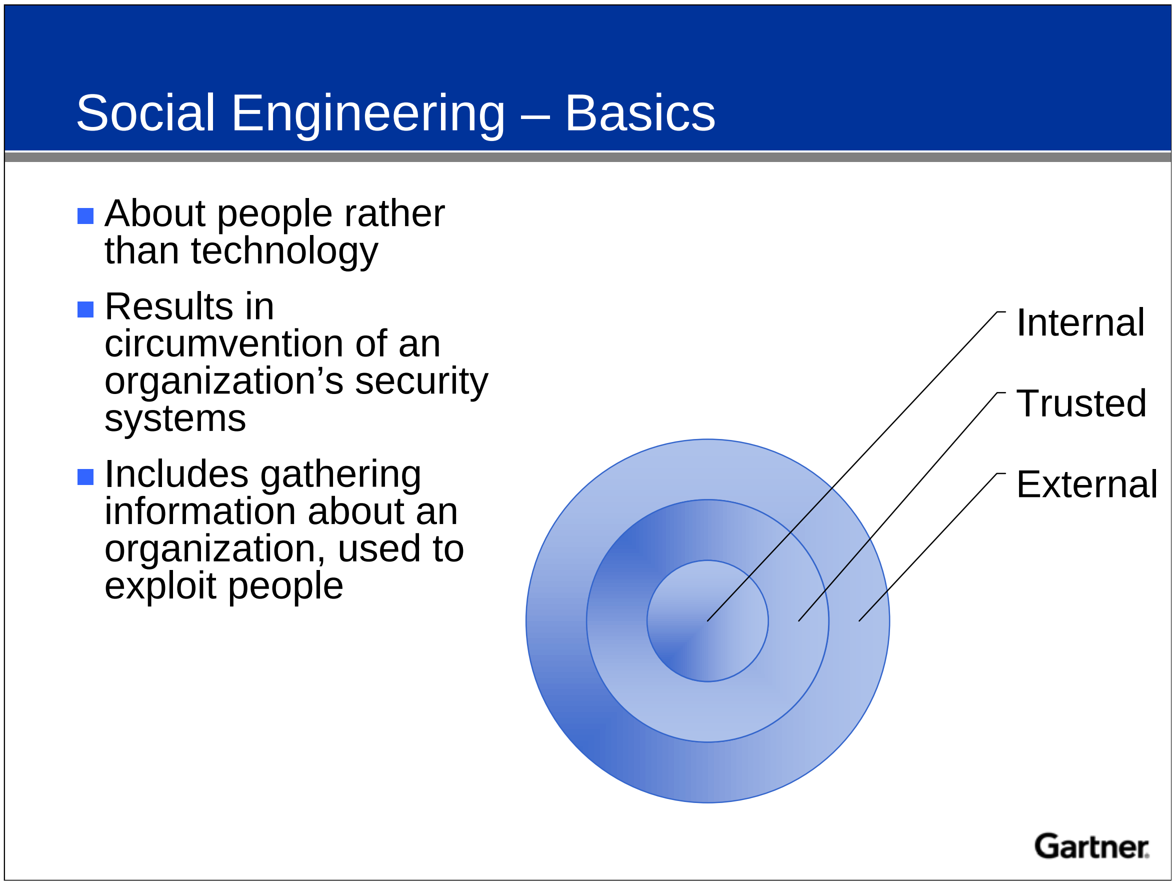

Social Engineering – Basics

- About people rather than technology
- Results in circumvention of an organization's security systems
- Includes gathering information about an organization, used to exploit people

Social engineering tends to involve manipulating people rather than technology to breach an organization's security.

Successful social engineering partially or completely circumvents an organization's security systems. The best firewall is useless if the person behind it reveals the access codes or information it is installed to protect.

Social engineering is the single greatest security risk in the decade ahead. Many of the most damaging security penetrations are due to social engineering. Most of the media-hyped "hacking" attacks are actually social engineering attacks.

Social engineering threats can be classified in three zones:

External Threats: Competitors, curious hackers, or individuals with criminal intent may create a sufficient level of short-term trust for the attacker to succeed.

Trusted Threats: Partner organizations, consultants and contractors deal with your sensitive information regularly, but may not be incorporated in your security plans.

Internal Threats: Employees may be dissatisfied, have innate criminal tendencies, or even be compromised by external sources. For high-value targets, it is not uncommon for an attacker to seek a temporary, "none-core" position to gain access to information. There may also be threats from former employees with remaining internal contacts who are willing to provide information that can compromise systems.

Client Issue: What is social engineering and how does it work?

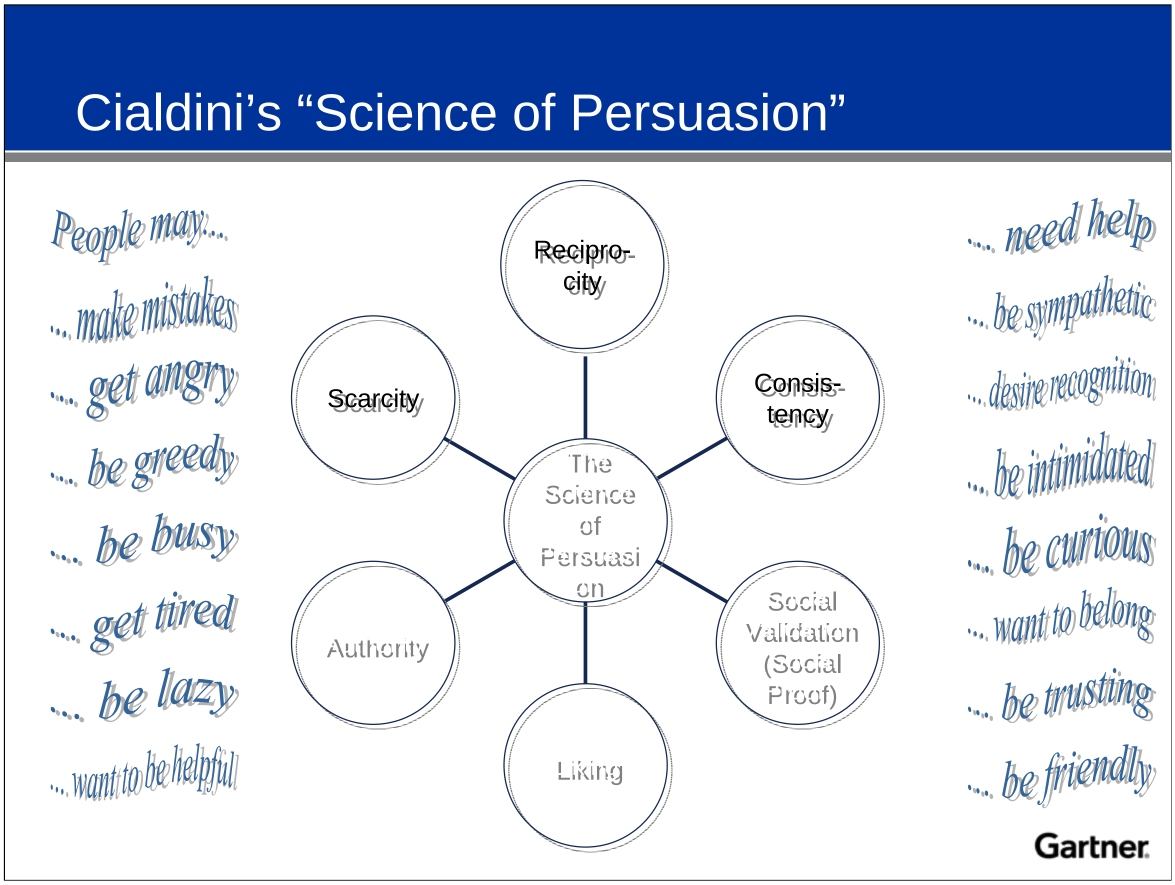

Cialdini's "Science of Persuasion"

People may...

... make mistakes

... get angry

... be greedy

... be busy

... get tired

... be lazy

... want to be helpful

... need help

... be sympathetic

... desire recognition

... be intimidated

... be curious

... want to belong

... be trusting

... be friendly

Social engineering attackers study human behavior closely, exploiting other people through careful manipulation to persuade them to release information or perform actions. Dr. Robert Cialdini ("The Science of Persuasion," Scientific American, February 2001), outlined six basic techniques:

Reciprocity: If an attacker helps the target, the target is more likely to help the attacker in return.

Consistency: Initially, the attacker may seek to get the target to commit to helping him in the future. When the attacker calls in the favor, he or she is more likely to help.

Social validation: If the target fears being left out of something, they will be keen to help an attacker.

Liking: People will perform favors for those they like, so the attacker must appear similar to the target or project an attractive personality.

Authority: Under pressure from someone who is ostensibly "in charge", people are likely to do things that they otherwise would not.

Scarcity: If the target fears missing out on an opportunity, they may be more willing to help..

These ploys are closely related to the confidence tricks that drive the plots of many films. However, social engineering is not so easy to detect in real life.

Client Issue: What is social engineering and how does it work?

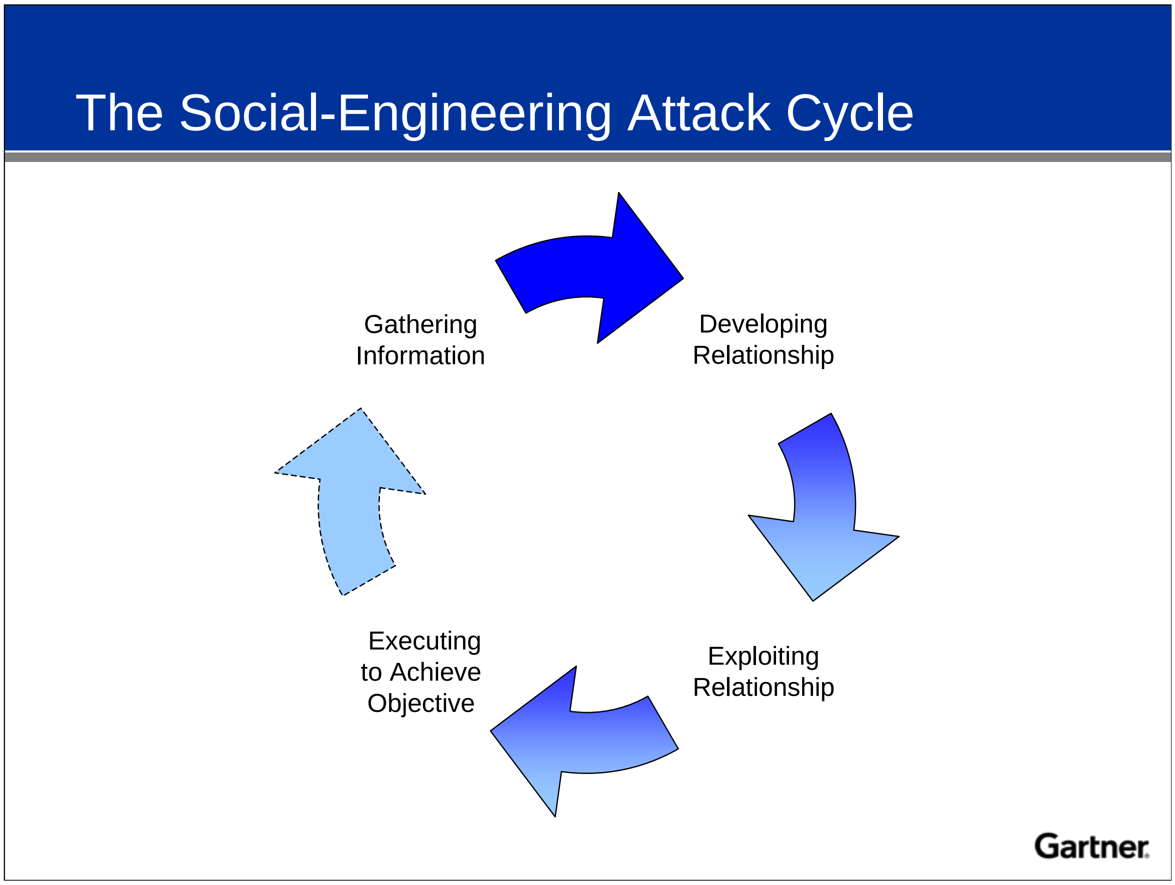

The Social-Engineering Attack Cycle

Gathering Information

Developing Relationship

Exploiting Relationship

Executing to Achieve Objective

Social engineering attacks , though varied, exhibit a common pattern that is often recognizable and preventable.

Information Gathering: The attacker uses a variety of information sources —a phone list, social security numbers, dates of birth, mothers’ maiden names, system architectures or organizational structures/procedures. This information serves as a basis for a relationship with someone connected to the eventual target.

Developing Relationship: The social engineer develops a rapport with the target, either briefly or over several weeks, creating a trust relationship that can then be exploited.

Exploiting Relationship: The attacker manipulates the target into revealing information (such as passwords, credit card numbers, vacation schedules) or performing an action (for example, creating an account, reversing charges) that would not normally occur. This information or action is the end objective, or it can be used to stage the next attack/phase of attack.

Executing to Achieve Objective: The result of Phase 3 is used to achieve the end objective, or it can be used to fuel the next stage of the attack. Often, an attack can include a number of these cycles — combined with traditional cracking methods and some physical information gathering — to achieve the end objective.

Client Issue: What is social engineering and how does it work?

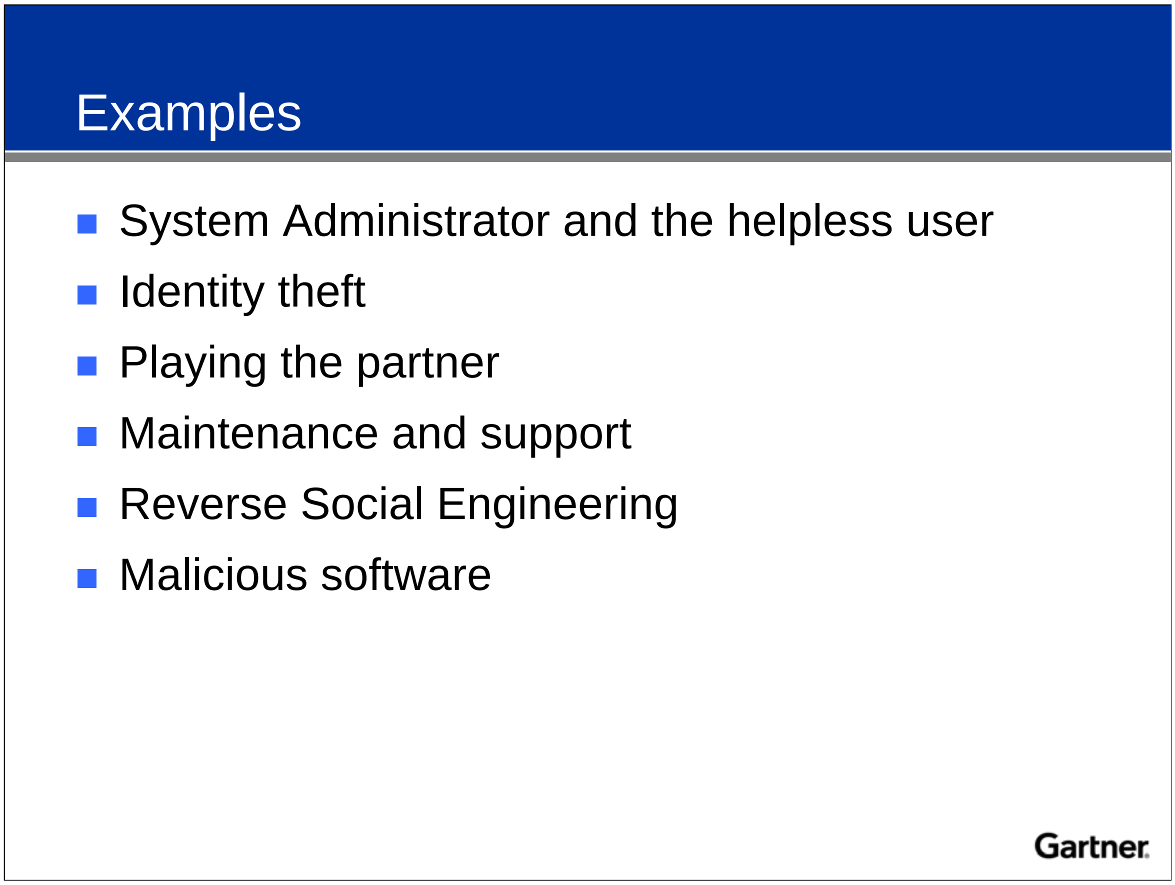

Examples

- System Administrator and the helpless user
- Identity theft
- Playing the partner
- Maintenance and support
- Reverse Social Engineering
- Malicious software

Think of any good scam, con, or “grift” and there is probably a social engineering equivalent.

Of increasing concern is identity theft. A busboy in New York City used social engineering and traditional cracking skills to assume the identity of a number of famous people and gain access to their financial information to place illegal orders.

Attackers will often gain access to an organization by taking a temporary job with it, or serving as part of the support staff. These workers tend to be somewhat invisible and go through less scrutinized hiring procedures, yet have complete access to the physical premises.

Reverse social engineering attacks tend to be more complicated, often involving the creation of false materials or altering existing materials to draw a target toward the attacker. The attacker then waits for the target to initiate a call, having developed a relationship and level of trust that can be exploited.

Many of the most prolific viruses are actually social engineering attacks, such as Melissa or ILOVEYOU. These viruses only work if the user executes them on his or her system, accomplished by creating compelling content or subject lines, or due to the origin.

Client Issue: What is social engineering and how does it work?

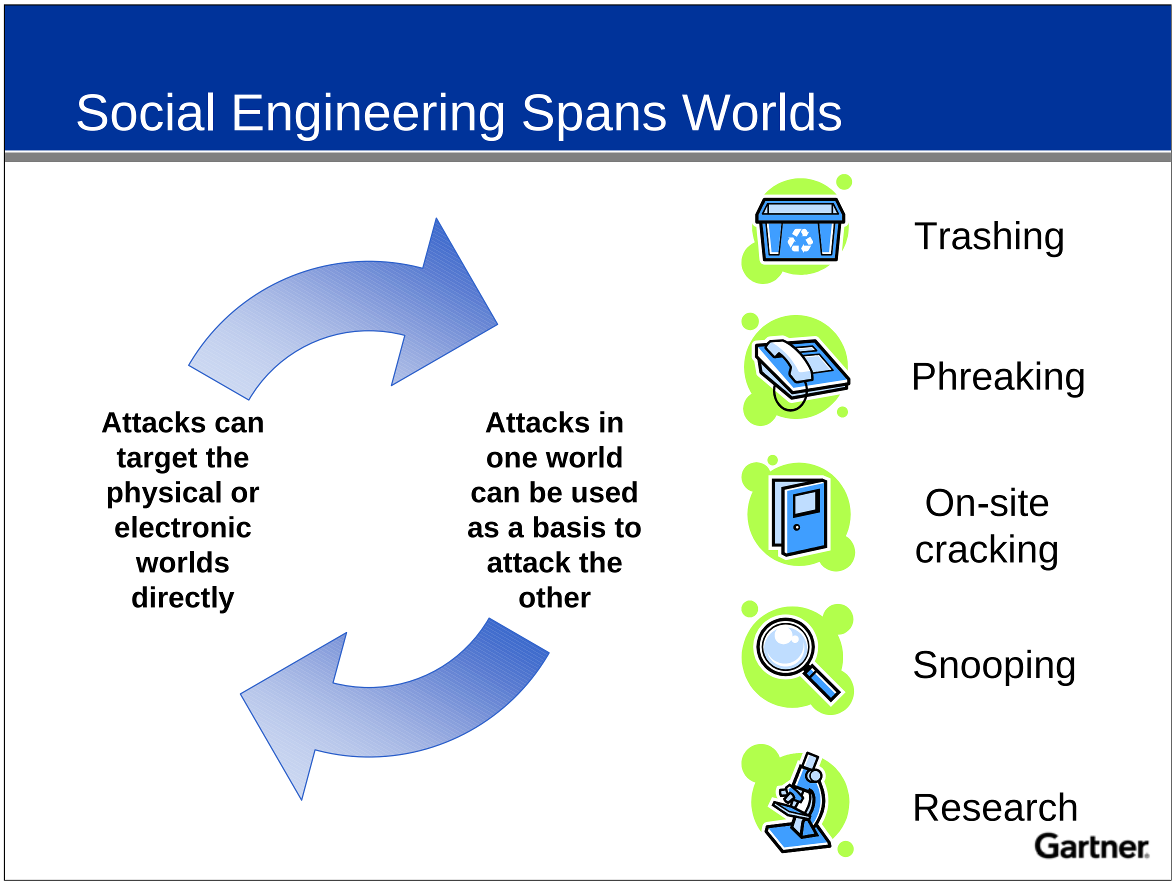

Social Engineering Spans Worlds

Attacks can target the physical or electronic worlds directly

Attacks in one world can be used as a basis to attack the other

Trashing

Phreaking

On-site cracking

Snooping

Research

Gartner

Attackers might exploit any information source. Old invoices or phone lists may not seem to pose a threat, the attacker can glean information from them to forge a relationship. Attackers may collect information remotely, over the Internet, or simply look over someone's shoulder. Corporate and personnel data is readily available and accessible on the Internet.

Every scrap of information is valuable to an attacker. A phone list or organization chart can reveal information most employees assume only their peers know. That information can then be used to delve deeper, until attackers finally convince their targets to release the information they need to compromise the organization’s security.

Other ways an attacker can gather information include:

Trashing — Examining garbage for sensitive information. Few organizations secure their trash, or train employees on what to throw away and what to shred. Source codes, passwords, schematics or internal memos often end up in unsecured dumpsters.

Phreaking — Cracking phone/voicemail systems to listen to recorded messages or leave false messages.

On-site cracking — Visiting a physical location to gather information.

Snooping — Examining paper mail, e-mail, “sniffing” networks or simply listening to conversations.

Research — Incredible amounts of information are readily available on nearly any person or organization.

Client Issue: How can a “wetware intrusion prevention system” defend against social engineering?

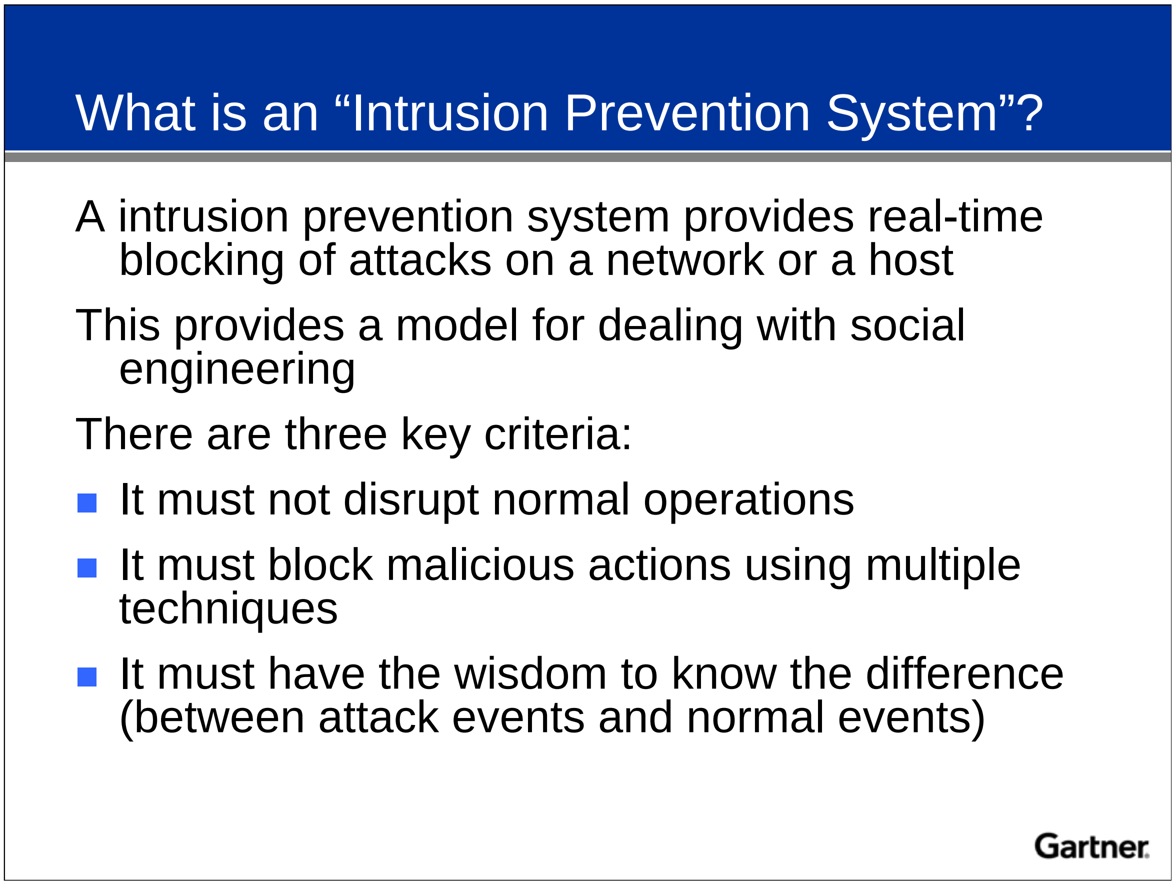

What is an “Intrusion Prevention System”?

A intrusion prevention system provides real-time blocking of attacks on a network or a host

This provides a model for dealing with social engineering

There are three key criteria:

- It must not disrupt normal operations
- It must block malicious actions using multiple techniques
- It must have the wisdom to know the difference (between attack events and normal events)

To defend against social engineering attacks, you need a system of people and processes that protects the organization against intrusion.

Such a system parallels the function of an intrusion prevention system on the network, for which there are three key criteria:

- It must not disrupt normal operations
- It must block malicious actions using multiple techniques
- It must have the wisdom to know the difference (between attack events and normal events)

Client Issue: How can a "wetware intrusion prevention system" defend against social engineering?

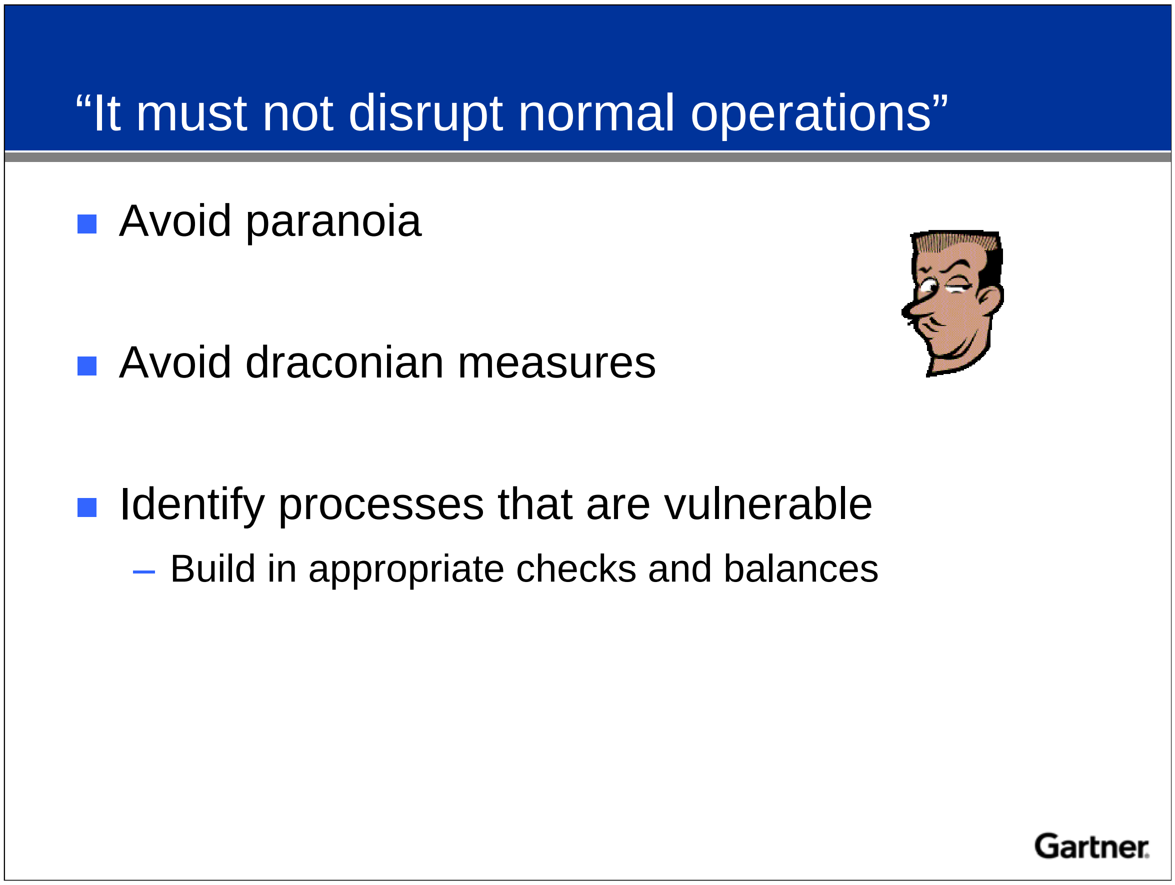

"It must not disrupt normal operations"

- Avoid paranoia
- Avoid draconian measures
- Identify processes that are vulnerable
  - Build in appropriate checks and balances

Paranoia and draconian measures against social engineering risk disrupting the company's normal operations. Appropriate checks and responses must be built into all business processes that might be vulnerable to social engineering, but these must avoid generating "false positives."

While anything clearly hostile should be blocked, suspicious requests may be passed and should be reported to management or security personnel.

Client Issue: How can a “wetware intrusion prevention system” defend against social engineering?

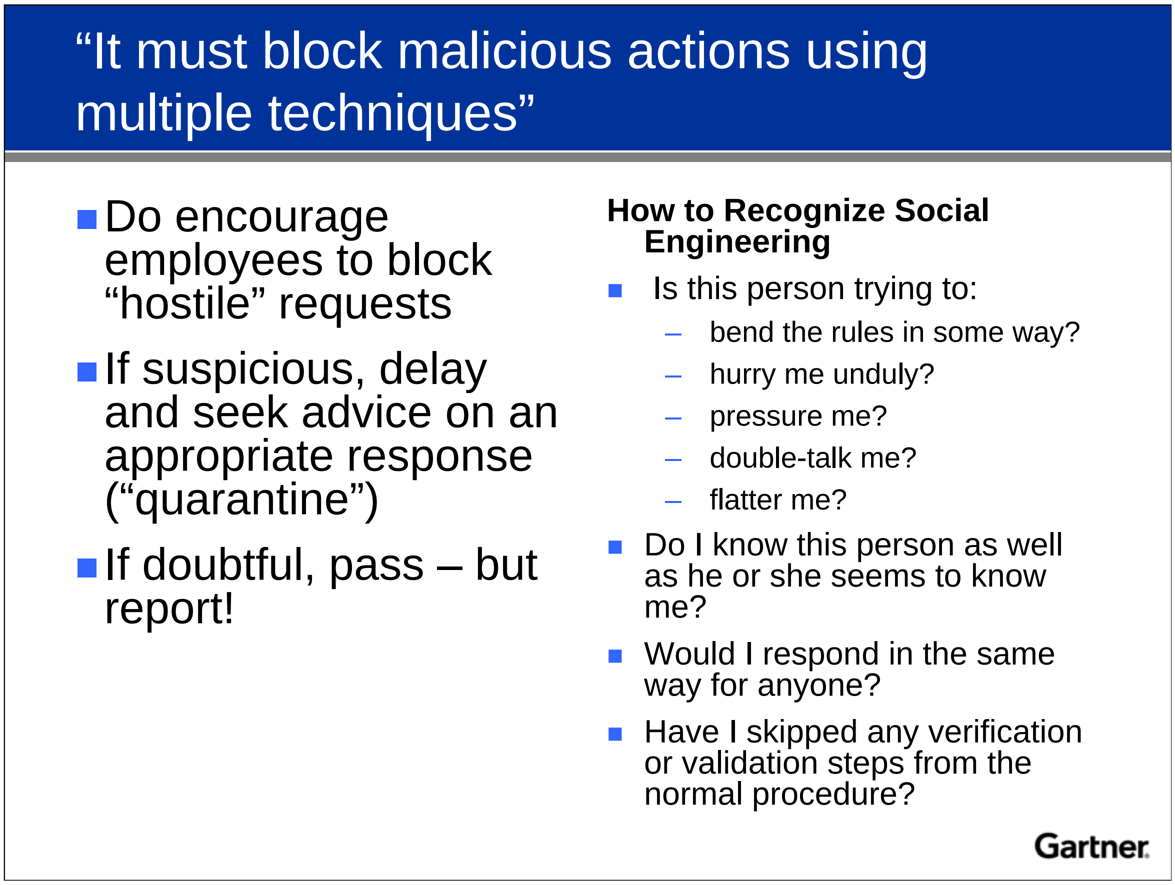

“It must block malicious actions using multiple techniques”

- Do encourage employees to block “hostile” requests
- If suspicious, delay and seek advice on an appropriate response (“quarantine”)
- If doubtful, pass – but report!

How to Recognize Social Engineering

- Is this person trying to:
  - bend the rules in some way?
  - hurry me unduly?
  - pressure me?
  - double-talk me?
  - flatter me?
- Do I know this person as well as he or she seems to know me?
- Would I respond in the same way for anyone?
- Have I skipped any verification or validation steps from the normal procedure?

Gartner

The best outcome is to block an attack in real time, and this should be made clear in the organization's security policies. Employees should be encouraged to refuse requests that are clearly hostile or, if the requests only seem suspicious, delay until they can discuss an appropriate response with line management or security personnel.

People must be trained to recognize the ploys used in social engineering and these lessons should be reinforced in documented procedures.

Client Issue: How can a “wetware intrusion prevention system” defend against social engineering?

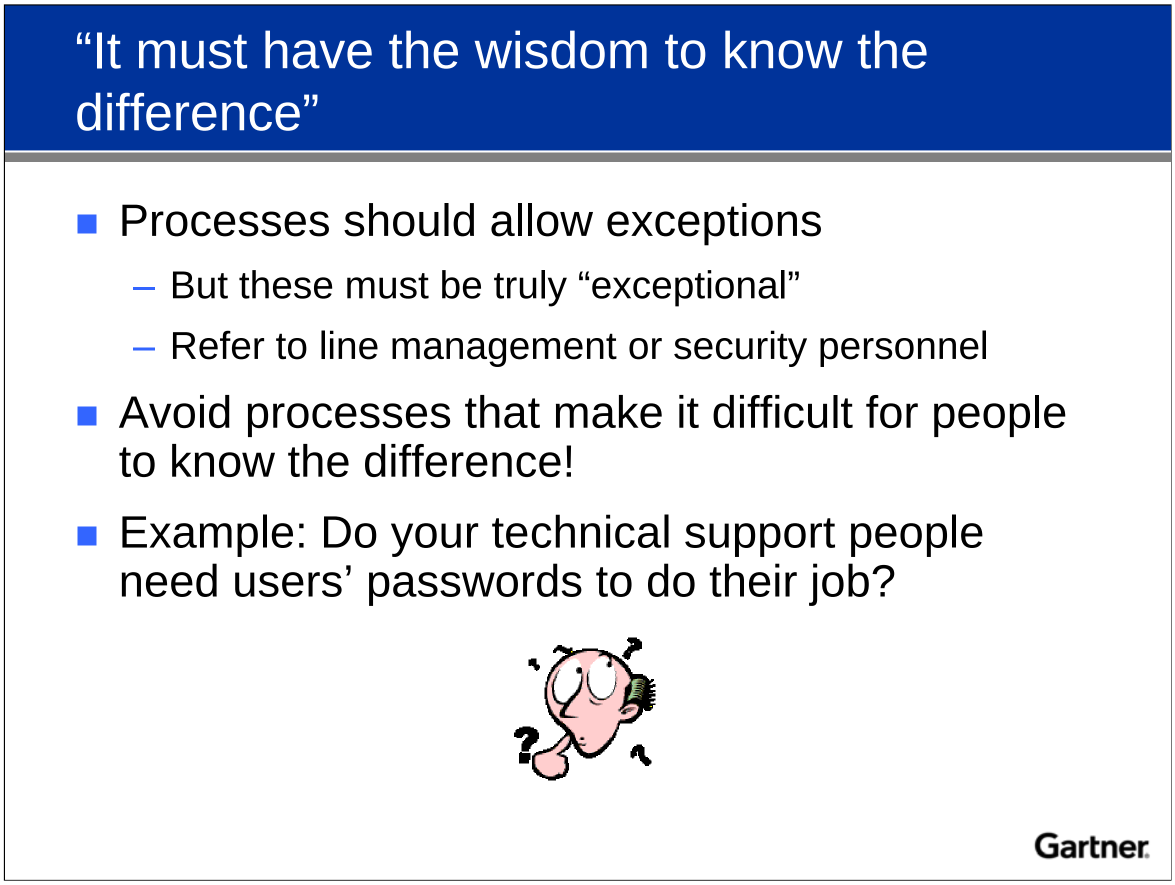

“It must have the wisdom to know the difference”

- Processes should allow exceptions
  - But these must be truly “exceptional”
  - Refer to line management or security personnel
- Avoid processes that make it difficult for people to know the difference!
- Example: Do your technical support people need users’ passwords to do their job?

No policy should be so rigid that it does not allow for exceptions. On the other hand, lax procedures around information security will likely make it difficult for people to distinguish between attack events and normal events. Any exceptions that blur this difference can create new vulnerabilities.

Exceptions should be decided by referral to line management or security personnel.

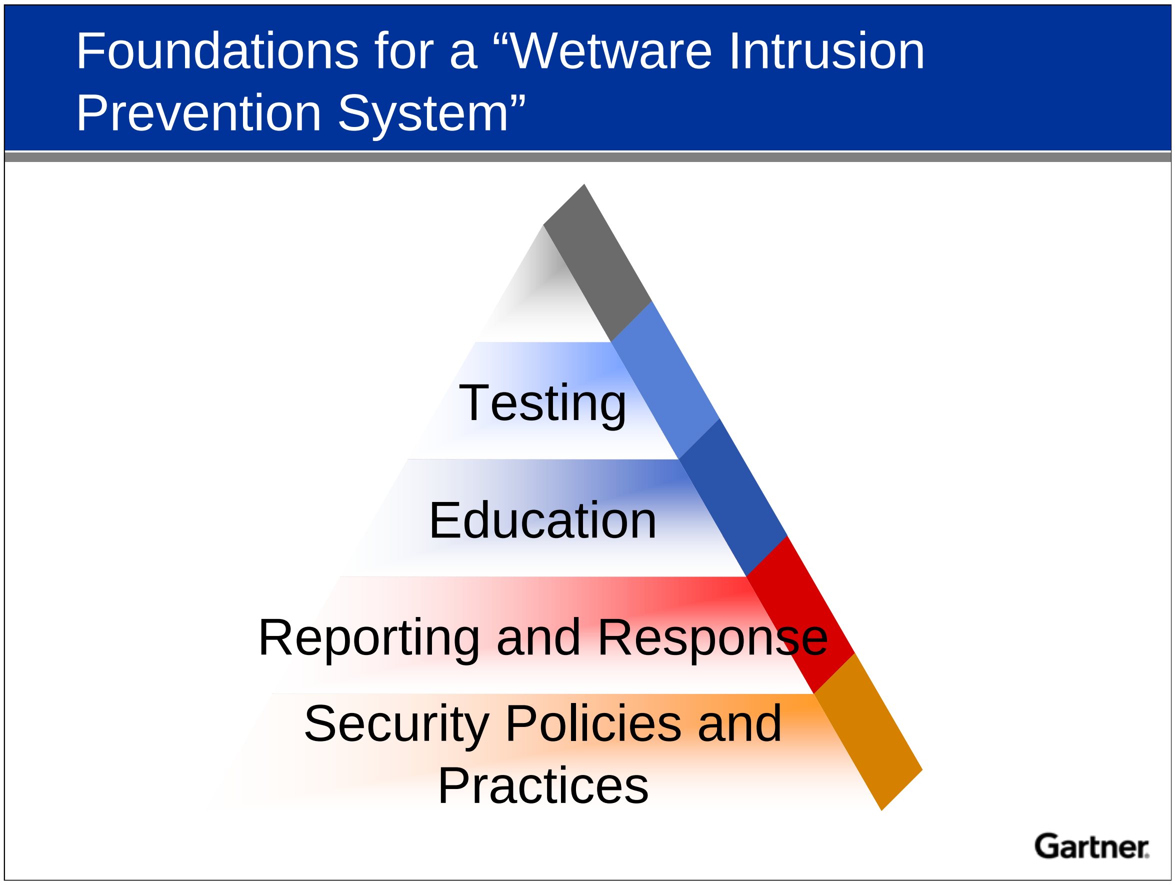

Foundations for a “Wetware Intrusion Prevention System”

Testing

Education

Reporting and Response

Security Policies and Practices

Gartner

Defense against social engineering relies on:

Clear, well understood and consistent security policies and practices

Clear channels for reporting suspicious contacts and active processes for responding to reports

Educational programs at all levels of the organization, so that everyone understands the nature of the attacks and how best to foil them

Inclusion in the security tests that the organization periodically carries out

Policies, processes, and testing are issues of governance and require more executive attention. Education requires executive buy-in, but will be be managed more day-to-day by others in the organization.

Client Issue: How can you build a “wetware intrusion prevention system?”

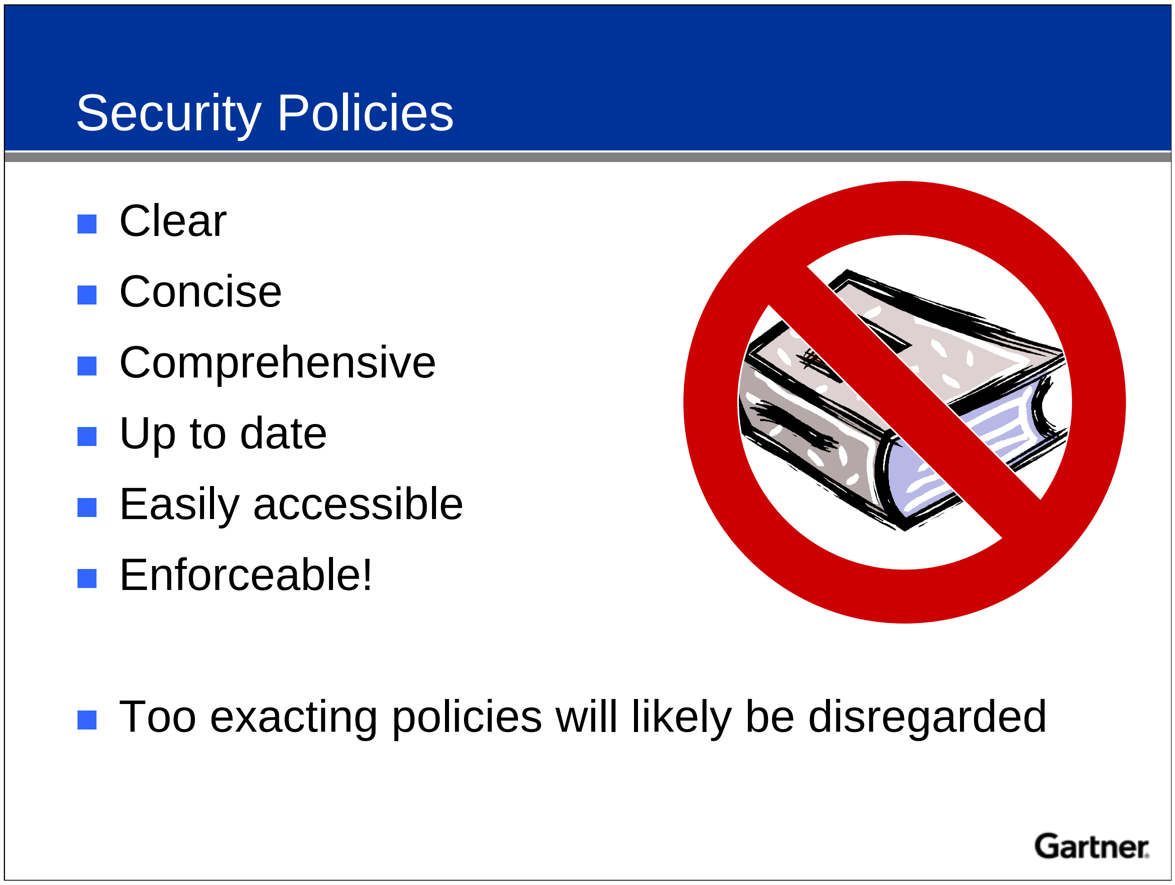

Security Policies

- Clear
- Concise
- Comprehensive
- Up to date
- Easily accessible
- Enforceable!

- Too exacting policies will likely be disregarded

To be effective, a security policy must be clear, concise, comprehensive, up to date, easily accessible and — most importantly — enforceable. When a security policy is draconian or exacting, compliance will be lax and overwhelmed users will circumvent or disregard it.

Because social engineering is primarily a human issue, it needs to be recognized especially in your personnel security policy.

It’s important to ensure that members of your organization are screened for criminal behavior (background checks) or prior work issues. Ancillary workers (maintenance, housekeeping) also need to be screened and should be given only limited physical access.

Terminated employees can leave with valuable information and retain access to internal systems; organizations must develop and enforce a comprehensive lockout procedure when employees are terminated.

Remaining employees need to be notified of the termination (preferably in a positive tone) so they don’t continue to feed information to the terminated individual.

The best way to hinder disgruntled employees and internal theft is to develop a positive working environment where employees feel valued and care about the organization’s success. This probably won’t keep everyone happy, but it will seriously limit anger against the organization and create a positive environment where employees are concerned about the company and security threats against it.

Client Issue: How can you build a "wetware intrusion prevention system?"

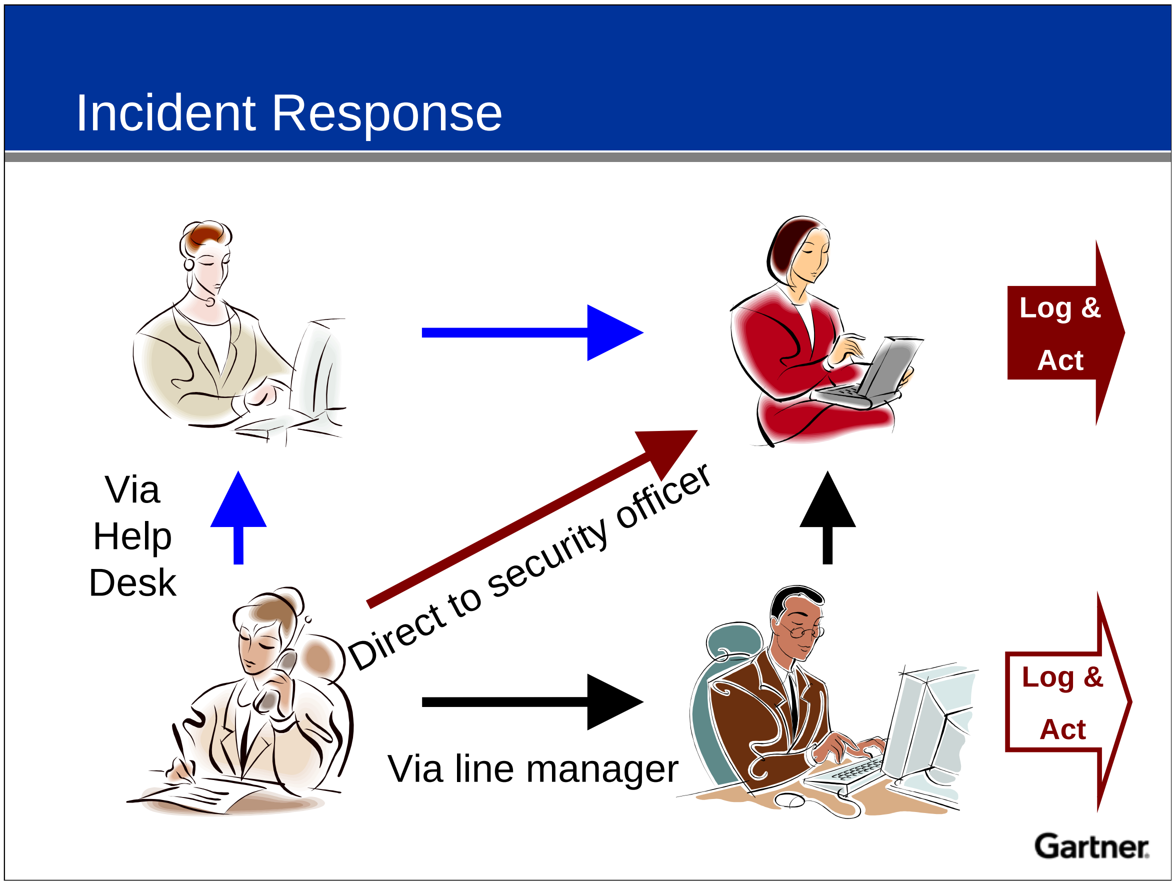

Incident Response

Via Help Desk

Direct to security officer

Via line manager

Log & Act

Log & Act

Gartner

Employees should know when and to whom to report suspicious contacts. Management and security personnel must, in turn, log and act on each report. Otherwise, employees will not bother to report their suspicions. Openness and transparency should be encouraged.

Give security its own place within the organizational structure. Those responsible for keeping systems running don't generally have sufficient resources or training to manage security. Plus, managing security can conflict with keeping systems accessible and efficient.

Many security breaches are overlooked because employees do not have clear channels to report suspicious behavior or incidents. Provide one point of contact and respond positively to every report. People may not always understand the difference between IT and physical security and, thus, will need that single point of contact.

There should be clear channels of communication between those responsible for physical and digital security. The loss of a laptop computer is the loss of a physical asset used to interact in the digital world, and both digital and physical security experts have roles in the investigation and management of the loss.

Action Item: Establish clear communications channels between the cyber-incident response team (CIRT) and physical security teams, including cross-training and joint command structures.

Client Issue: How can you build a "wetware intrusion prevention system?"

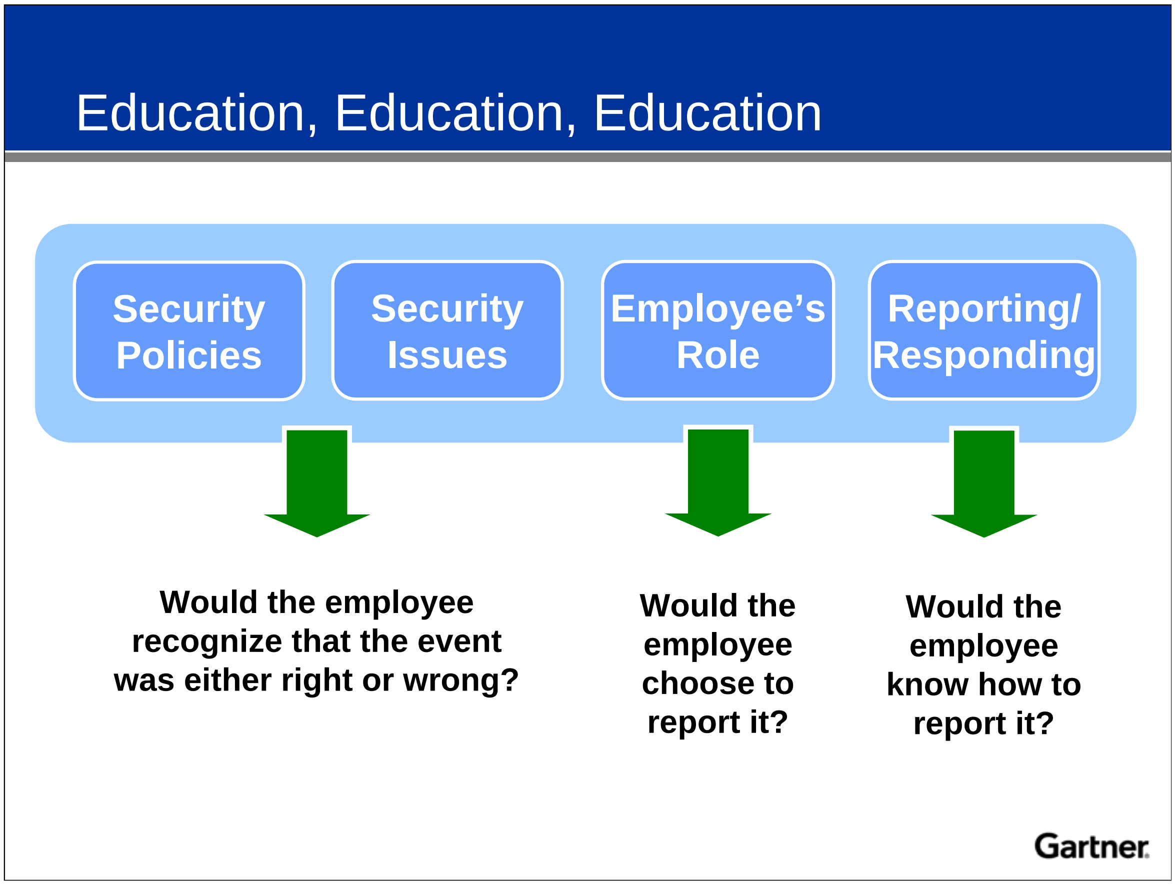

Education, Education, Education

Security Policies

Security Issues

Employee's Role

Reporting/ Responding

Would the employee recognize that the event was either right or wrong?

Would the employee choose to report it?

Would the employee know how to report it?

Gartner

The single strongest defense against social engineering is a base of educated employees, management, security personnel and administrators. Everyone must be trained on the security policies and on general security issues. When people understand the issues, they are more likely to comply with policy and procedures.

Employees need to be trained in the following areas:

Security policies: Employees need to understand policies to both limit their personal violations and allow them to recognize when others violate policies.

Security issues: What is a virus anyway? Employees need training on a variety of security issues, from physical access to information misuse to e-mail safety. Ongoing training should include new security issues as they arise, as well as signs of an impending incident before it causes damage.

Impact on the organization/employee: People tend to pay less attention to issues that don't directly affect them. Employees are more likely to be aware and proactive if they understand the negative consequences, both to the organization and themselves.

How to report/respond: Obviously, not everyone needs to be trained to put out a fire, but they do need to know how to hit the fire alarm, dial 911 and safely evacuate the building.

However… if management doesn't set an example and the culture is to "look the other way," then security becomes unenforceable. A security-conscious culture can have the greatest overall, positive impact on an organization.

Client Issue: How can you build a “wetware intrusion prevention system?”

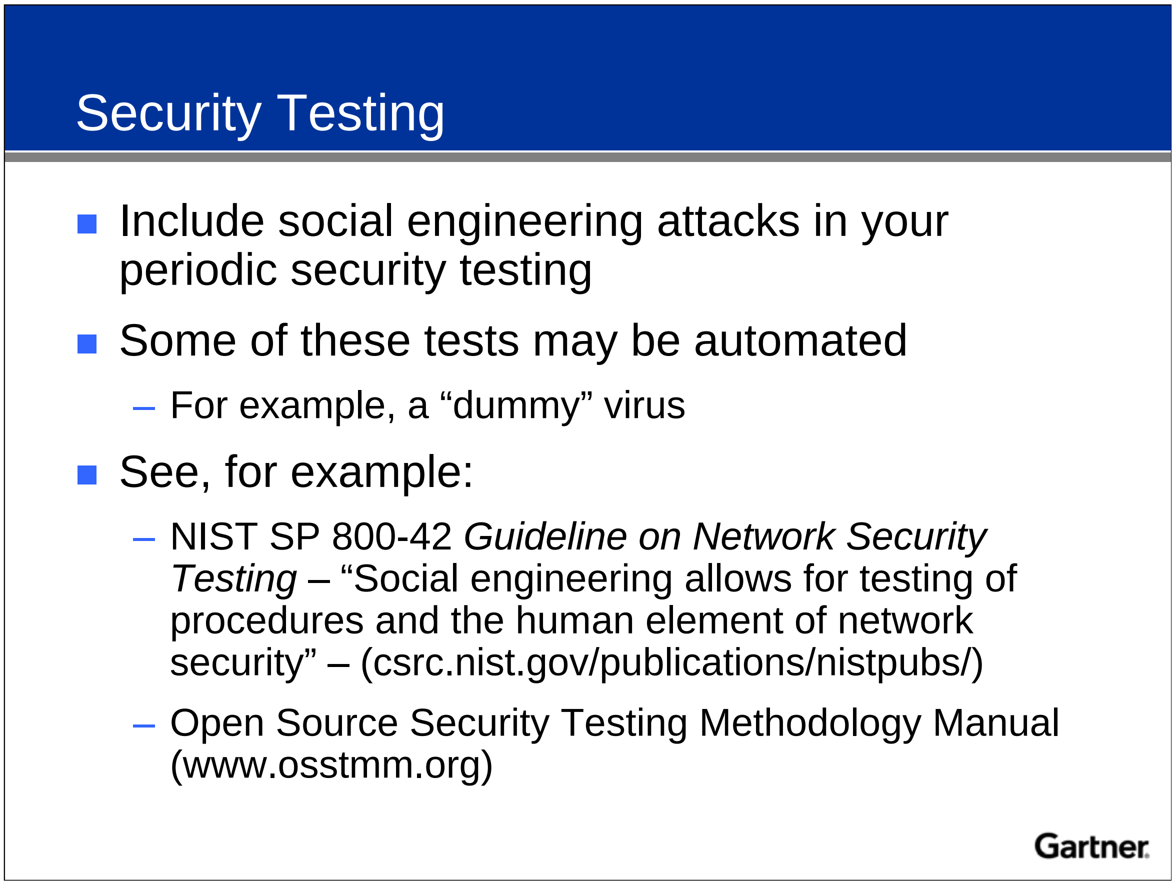

Security Testing

- Include social engineering attacks in your periodic security testing
- Some of these tests may be automated
  - For example, a “dummy” virus
- See, for example:
  - NIST SP 800-42 *Guideline on Network Security Testing* – “Social engineering allows for testing of procedures and the human element of network security” – (csrc.nist.gov/publications/nistpubs/)
  - Open Source Security Testing Methodology Manual (www.osstmm.org)

Gartner

All organizations should conduct periodic testing of their information security status. Such a test should include social engineering attacks to assess the effectiveness of policies, education and incident response processes.

Good social engineering skills are an asset for a security professional, allowing you to probe for vulnerabilities in your “procedures and the human element of network security”, just as a tool such as Nessus allows you to probe for technical vulnerabilities. Social engineering is an important part of any penetration test.

Some social engineering tests can be automated. Remember that many viruses (but not worms) rely on people’s curiosity (or baser instincts) to execute them. How can you be sure that your users respond appropriately and do not double-click on attachments in suspicious emails? One Gartner client wrote a dummy virus that was benign and sent it at random to users in a deliberately “spammy” email: double-clicking it displayed a reminder about policy for the user and sent an alert to the security officer. The idea here is not to censure users, but to reinforce education and provide some measure of effectiveness of that education for the security officer.

Client Issue: How can you build a "wetware intrusion prevention system?"

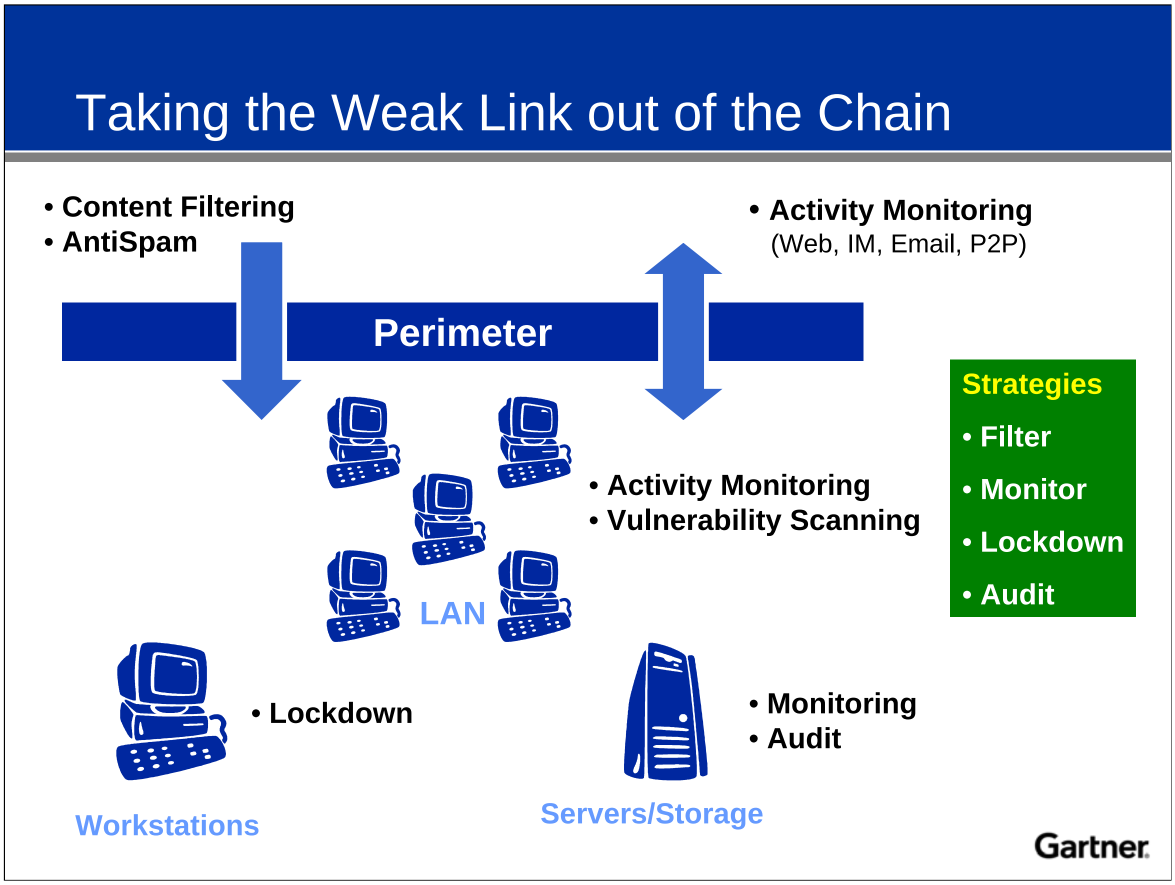

Taking the Weak Link out of the Chain

- Content Filtering
- AntiSpam

- Activity Monitoring
(Web, IM, Email, P2P)

Perimeter

Strategies
- Filter
- Monitor
- Lockdown
- Audit

- Activity Monitoring
- Vulnerability Scanning

LAN

- Lockdown

Workstations

- Monitoring
- Audit

Servers/Storage

Gartner

The best defense against social engineering is to modify user behavior… a kind of positive social engineering! But there are some recent advances in security technology that enhance our ability to protect our users and organization. Gartner recommends a four-layer strategy to defend against social engineering across the technology infrastructure:

Filter harmful material before it reaches users. This includes URL filtering, and email, IM, and P2P filtering for inappropriate sites, spam, spim, and unapproved content types.

Monitor inbound and outbound employee activity for intellectual property leaking out, inappropriate activity, and unacceptable use.

Lockdown workstations to prevent Trojan horses, other malicious software, or employees from installing potentially harmful applications with legitimate purposes.

Audit information repositories and applications for unusual access or behavior. Internal threats can be discovered by monitoring usage patterns in this way. Tools are available to analyze logs for these patterns, and observant IT staff may discover problems early in the process.

- **Use effective checks and balances in processes that are vulnerable to social engineering.**
- **Review what corporate information is publicly available.**
- **Use technical controls – filter, monitor, lockdown, audit – but understand their limits.**
- **Have an effective security policy.**
- **Have clear incident reporting paths.**
- **Educate users – security-aware employees are the most effective defense.**
- **Test!**

Because social engineering is more about people than machines, effective defenses must involve an organization's greatest resource — its personnel. Effective checks and balances and security-aware employees are the most effective methods to defend against social engineering attacks. Those that fail to educate their employees risk falling victim to malicious con artists.

This is the end of this presentation. Click any where to continue.