Software Engineering Institute

Insider Threat Security Reference Architecture

Joji Montelibano
Andrew Moore

April 2012

CERT® Program

http://www.sei.cmu.edu

Carnegie Mellon

Table of Contents

List of Figures

List of Tables

Abstract

The Insider Threat Security Reference Architecture (ITSRA) provides an enterprise-wide solution to insider threat. The architecture consists of four security layers: Business, Information, Data, and Application. Organizations should deploy and enforce controls at each layer to address insider attacks. None of the layers function in isolation or independently of other layers. Rather, the correlation of indicators and application of controls across all four layers form the crux of this approach. Empirical data consisting of more than 700 cases of insider crimes show that insider attacks proved successful in inflicting damage when an organization failed to implement adequate controls in any of three security principles: authorized access, acceptable use, and continuous monitoring. The ITSRA draws from existing best practices and standards as well as from analysis of these cases to provide actionable guidance for organizations to improve their posture against the insider threat.

1 Introduction

From the time an insider decides to attack to the point at which damage is done, there exist opportunities for the prevention, detection, and response to the attack. Ideally, the organization will be able to prevent the attack altogether. Failing this, the organization should have adequate controls in place to detect the malicious activity. Finally, the organization should have a proper incident response plan to mitigate the damages resulting from the insider's actions. The areas above and below the timeline in Figure 1 denote the data the organization should collect. The top portion represents nontechnical data, such as human resources (HR) records and physical security logs, while the bottom portion represents technical data, such as database logs and remote access logs.
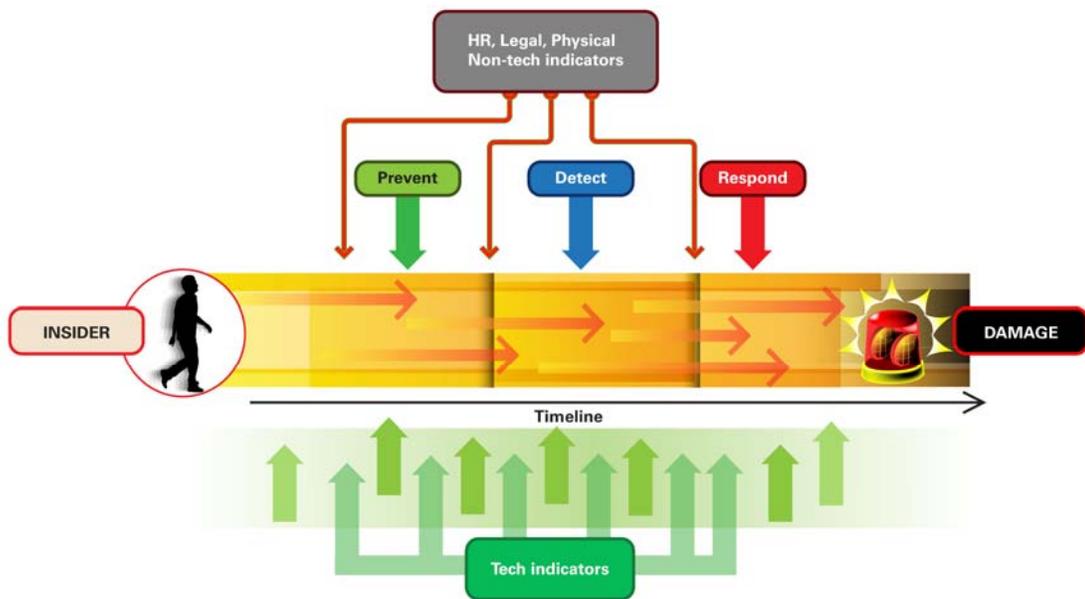
Figure 1:   Opportunities for Prevention, Detection, and Response for an Insider Attack

Correlation of data is the key. Such data will come from disparate sources across the enterprise, and the challenge is the correlation of such data to inform security staff without overwhelming them. The Insider Threat Security Reference Architecture (ITSRA) is designed to address this challenge.

2   The Components of the ITSRA

Figure 2 shows the four layers of the ITSRA. The Business Security layer contains high-level business requirements, such as an organization's mission. This layer involves the creation of policies, procedures, and other guidance that determines the level of security to be implemented in other layers. The Information Security layer describes the organization's underlying information infrastructure. This includes the information network and the components necessary to operate the organization's information services, such as routers, switches, and servers. This layer also contains the operating systems and software required to manage the infrastructure. The Data Security layer involves information assets considered to be proprietary to the organization. Such data can take the form of documents, spreadsheets, or databases. Finally, the Application Security layer addresses both internal development of software, as well as the deployment of non-operating-system applications used to fulfill a particular business mission. A Content Management System (CMS) and a Customer Relationship Management (CRM) system are examples of such applications. The Application Security layer ensures that these programs adhere to the security requirements defined at the Business Security layer.
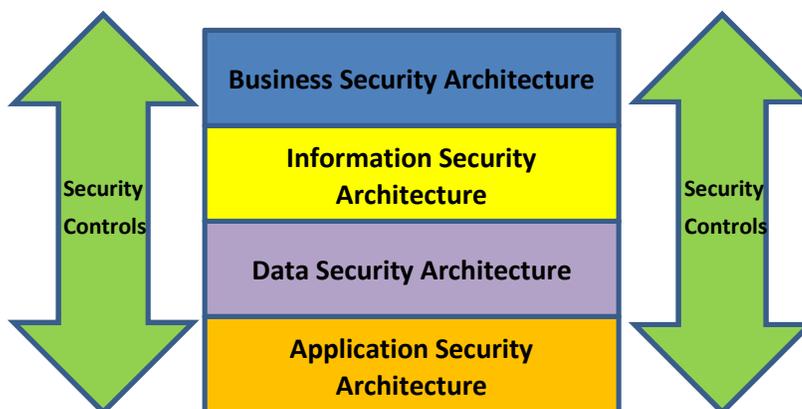
Figure 2:   Insider Threat Security Reference Architecture

Security is the common thread running through all levels of a sound enterprise architecture. The two arrows on each side of the four layers indicate this cross-cutting role of security. There exists a wide body of research and products to help organizations implement security measures at each layer. For instance, secure business processes secure the business architecture layer, data protection mechanisms secure the data architecture layer, and so on. What is missing is a cohesive instrument that integrates disparate security controls into a single, comprehensive strategy. The ITSRA seeks to cover this gap by offering a structured approach to help organizations improve their level of preparedness to address the insider threat.

3   Empirical Foundations and Standards

Our research on insider threat is empirically based, drawn from rigorous analysis of over 700 actual cases of malicious insider activity [Cappelli 2009, Hanley 2011]. One such case illustrates how the ITSRA can be put to use. In this case, a foreign currency trader was able to cover up nearly $700 million in losses over a five-year period. He accomplished this by modifying the source code for his organization's trading system. The trader violated a number of Human Resources (HR) policies, including improper treatment of colleagues. However, because of his status as a "star performer" within the firm, he did not incur the organization's standard disciplinary actions. In this case, conventional standalone detection mechanisms, such as intrusion detection systems and configuration management (CM) systems, did not prove adequate to detect, let alone prevent, the crime. Rather, had principles of the ITSRA been applied, correlation of HR data would have triggered increased scrutiny and monitoring of this individual's online activities. This, when combined with an alert raised by the CM application revealing his changes to source code, would have uncovered the insider's illicit activities and, we believe, would have led to his earlier arrest and conviction. Needless to say, this may have even saved the organization itself from the financial loss and damaging publicity that followed.

The NIST Enterprise Architecture Model (EAM) [EOPUS 2007, NIST 2009] and the Federal Enterprise Architecture (FEA) [CIOC 2001, EOPUS 2007] form the foundations of the ITSRA. The ITSRA uses these enterprise-level models as a basis because insider threat cannot be fully addressed by a single department within an organization. That is, insider threat is an enterprise-wide problem and must be confronted with an enterprise-wide solution. The ITSRA is such a solution.

Figure 3 captures the approach that is used to create the ITSRA. The arrows indicate the cross-dissemination of security data gathered from different sources in the enterprise. This information sharing best informs and prepares the organization to prevent, detect, and respond to insider threats.

NIST Enterprise Architecture Model

Federal Enterprise Architecture

Controls include those that are administrative, technical, and physical as well as preventive, detective, compensating, and corrective.

Insider Threat Security Reference Architecture
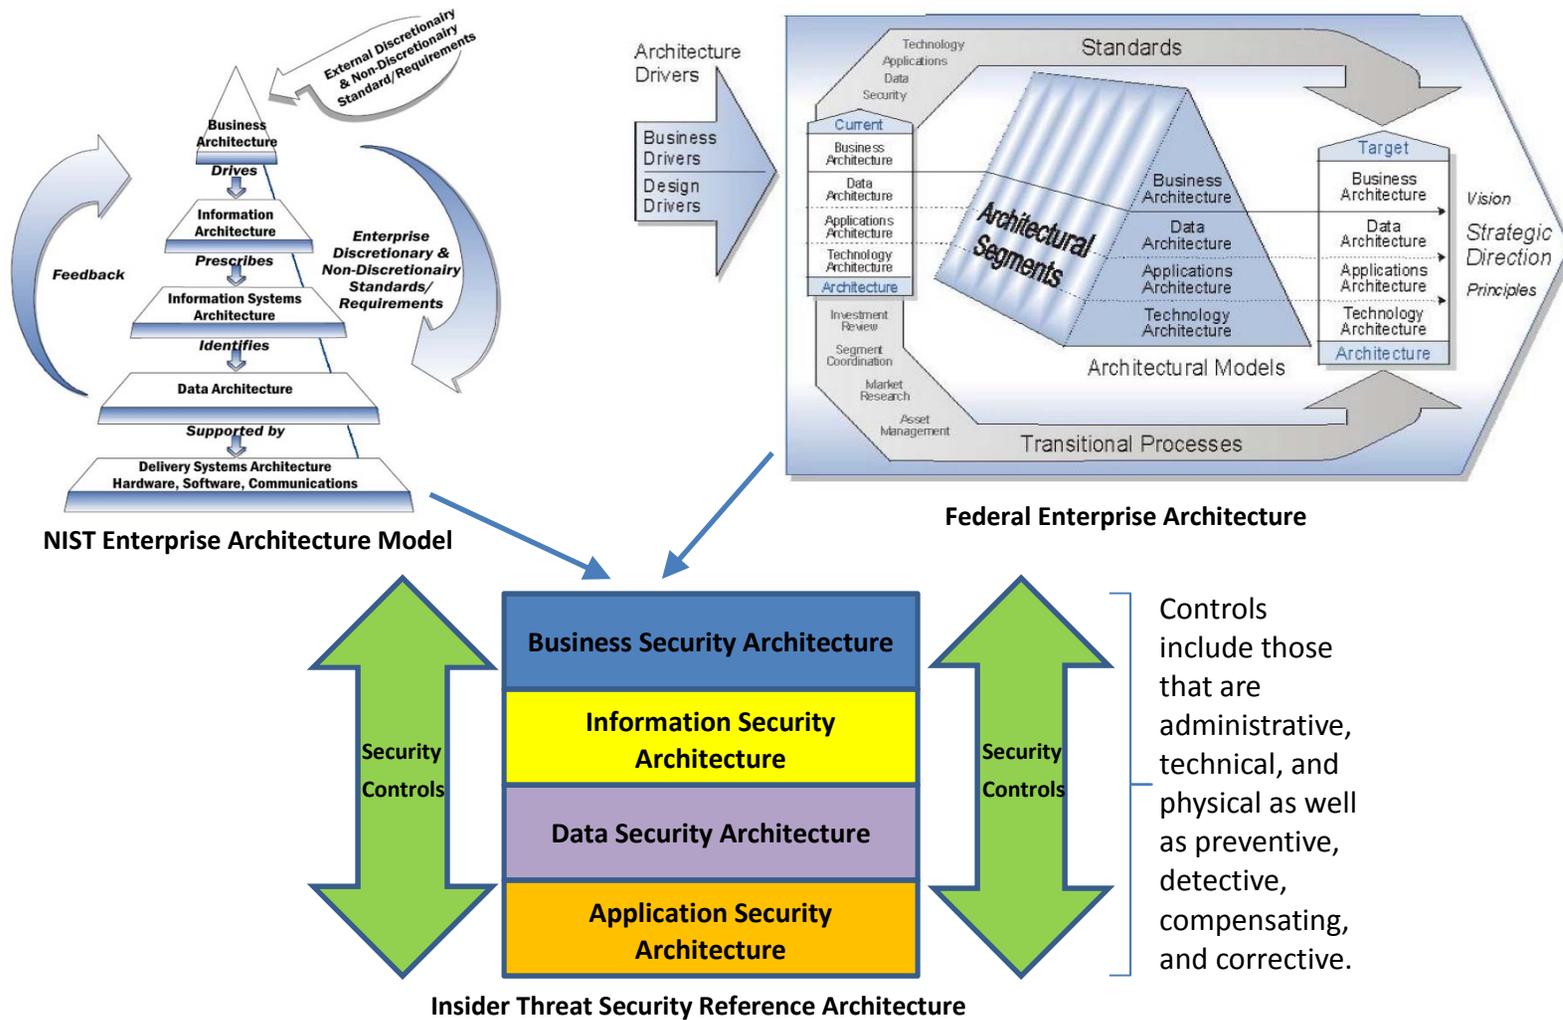
Figure 3: *The Insider Threat Security Reference Architecture Is Derived from the NIST Enterprise Architecture Model [EOPUS 2007, NIST 2009] and the Federal Enterprise Architecture [CIOC 2001, EOPUS 2007]*

4 Application of the ITSRA

The process of applying the ITSRA should involve refinement and customization at each layer to make the corresponding controls applicable to the organization in question. To this end, we created and analyzed an insider threat database to develop models of insider attack. Analysis of the more than 700 cases reveals that each case can be categorized into one of the following:

- IT sabotage
- theft of intellectual property (IP)
- fraud [Cappelli 2009]

IT sabotage describes an insider's use of information technology (IT) infrastructure to direct specific harm at an organization or individual. Theft of IP involves the use of IT to steal IP from an organization. This includes industrial espionage involving insiders. Finally, fraud includes all cases involving insiders who used IT for the unauthorized modification, addition, or deletion of data for personal gain or theft.

Based on these three categories, we have created different models or patterns that represent the sequence of events in a given attack vector [Hanley 2011, Moore 2009]. An organization wishing to apply the ITSRA can choose the relevant crime patterns that present the most visible threat to its operations. The ITSRA then offers granular recommendations at each security layer to address that particular attack vector. Figure 4 below describes how the ITSRA transitions from a high-level reference architecture to an instantiated enterprise architecture customized to fit a specific organization's requirements.

High-level ITSRA

Insider threat attack pattern library
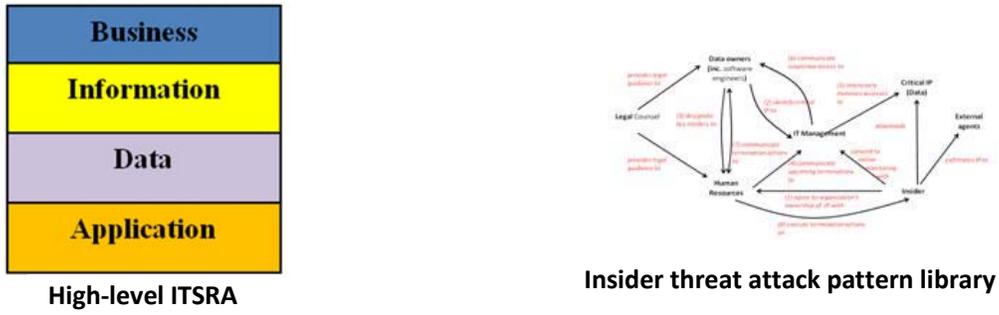
Organization chooses insider threat pattern

Detailed enterprise security architecture

Figure 4: ITSRA Combines with Attack Pattern Library to Form a Customized Enterprise Security Architecture

The customized enterprise security architecture shown above in Figure 4 takes the form of an ITSRA matrix, which we will describe below in Section 4.1.

The ITSRA offers granular recommendations by combining both operational- and policy-based guidance from the insider threat research and existing best practices to provide recommended controls such as

- business process guidelines
- policy formulation
- legal controls
- switch configuration
- security information and event management (SIEM) rules
- intrusion prevention system signatures

- HR procedures
- physical security practices

Most importantly, it combines all of these measures into a comprehensive and holistic framework that will better prepare an organization to prevent, detect, and respond to malicious insider activity. Table 1 below gives a sample listing of best practices available to security practitioners today. The ITSRA describes how these best practices can be integrated to form a truly enterprise-oriented framework.

Table 1:   Sample Security Architectures

ITSRA Layer	Security Architecture
Business	• Sherwood Applied Business Security Architecture (SABSA) [SABSA 2011] • NIST SP 800-37 [NIST 2010] • Zachman Framework [SABSA 2011] • Six Sigma
Information	• Open Security Architecture • Cisco SAFE [Chung 2010] • NIST SP 800-53 [NIST 2009]
Data	• Common Data Security (CDSA) [Blackwell 2009] • Oracle Database Security [Oracle 2011]
Application	• OWASP • CERT® Secure Coding Standards • Microsoft Application Security [Microsoft 2010]

4.1 The ITSRA Matrix

To develop the ITSRA, the authors superimposed the analysis of the insider threat database onto the best practices gleaned from the security architectures listed in Table 1 above. Specifically, in reviewing each architecture, we derived commonalities between different approaches and determined how these practices could have been applied to prevent or detect a specific attack. This approach revealed that security architectures are crafted to enforce the following most fundamental principles:

- authorized access
- acceptable use
- continuous monitoring

Applying these three principles to the cases of insider crimes does indeed affirm that each criminal act can be attributed to an organization's failure to implement one or more of the three security principles above.

Consider the case of the currency trader described in Section 3 above. Although his access to source code was authorized by the organization, his use of this privilege constituted unacceptable use. Although policies were in place restricting what was deemed acceptable, clearly in this case, there were insufficient means of enforcing such policies. The organization had separation-of-duties controls such that the back office verified every trade entered into the system. However, this particular insider social engineered the back office into skipping verification of his trades since he was "the star." In addition, when back office personnel questioned some of his illegal

---

® CERT is a registered mark owned by Carnegie Mellon University.

trades, he bullied them into overlooking their suspicions. So the business process controls were there, but not enforced. And because there was no correlation of issues between layers, management did not put the pieces of the puzzle together. In other words, although the organization may have instituted appropriate controls in the Business Security layer, it failed to extend such controls into the Information and Application layers, which allowed the trader to commit his crime. This reinforces the need to have a cross-cutting architecture that spans the breadth of all security layers. From another perspective, the organization granted authorized access to the trader and defined acceptable use, but it failed to apply continuous monitoring strategies to ensure that the trader adhered to these restrictions.

Table 2 shows the ITSRA Matrix, which gives a high-level summary of controls recommended by various security architectures, categorized by layer and which security principle (authorized access, acceptable use, or continuous monitoring) it best addresses.

Table 2: The ITSRA Matrix – Sample Subset of Controls per ITSRA Layer

	Authorized Access	Acceptable Use	Continuous Monitoring
Business	• legal guidance • physical security • separation of duties • need-to-know	• legal guidance • acceptable use policy • change management	• legal guidance • audits • assessments • asset prioritization
Information	• account management • host authentication (e.g., MAC address authentication) • authentication, authorization, and accounting (AAA) • multifactor authentication	• firewalls • proxies • IDS/IPS • file read/write restrictions	• SIEM rules • log correlation • intrusion detection • automated alerts • incident response • antivirus
Data	• account management • role-based access	• data classification • data tagging • least privilege	• data loss prevention (DLP) • intrusion detection • database alerts
Application	• account management • separation of duties	• code review • quality assurance • email filters • HTTP/HTTPS proxies	• audits • peer review • configuration and change management

5   Correlation

Insider cases in our database involved the exploitation of any one or more of the vulnerabilities shown in the ITSRA matrix. Current security architectures generally emphasize the need for controls to be in place to address the vulnerability in that particular layer [Blackwell 2009, Jabbour 2009]. What is needed is a formalized process to ensure that any countermeasures, whether they provide prevention, detection, or response controls, cross vertically across all security layers to provide the best protection possible.

Figure 5 below shows a snapshot of the "Authorized Access" column of the ITSRA matrix, focusing on select controls per security layer.
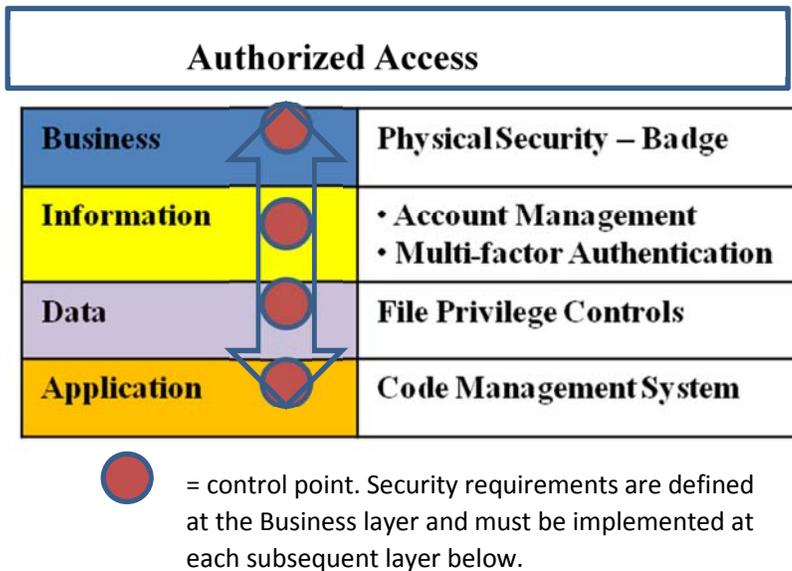
Figure 5:   Authorized Access Controls Span All Layers of the ITSRA

Any controls meant to enforce physical access to a closed area should also be extended into the logical realms of information, data, and application controls. Information controls may involve the existence of a dedicated account for that particular individual inside the closed area, along with multifactor authentication to confirm that individual's identity. Moving down to the data layer, file access controls for read-write privileges must be in place to restrict the individual's need-to-know access. For instance, an employee dedicated to biological research should not have any read access to internal salary records. Finally, if the closed area contains any sensitive applications, those applications must have appropriate safeguards in place to ensure authorized access. Using the investment trader as an example once again, if portions of the code should not have been accessible to him, a code management system such as CVS should have been in place to restrict his access.

Just as in Figure 2, the arrows in Figure 5 go both directions across the ITSRA stack. While it is true that high-level security requirements should be defined at the Business layer, any data collected from controls in any of the bottom layers should likewise inform the top layers. This has

major implications especially with regards to targeted monitoring, which will be discussed in Section 5.1 below. For instance, once controls have been implemented in a way suited to the ITSRA model, that in and of itself is not the desired end state. Rather, the continuous monitoring component is still in place. So if an insider were to violate a control in the Information, Data, or Application layers, this information should be communicated up to senior leadership so that they can take action and implement high-level business decisions, such as sanctioning the employee or even terminating employment.

5.1 Incident Response and Targeted Monitoring

The preceding discussion describes the preventive and detective elements of the ITSRA. That is, defining requirements at the highest level and implementing relevant controls through the other ITSRA layers will position an organization well in countering any insider threats. There is, however, an additional dimension to the ITSRA, and that is incident response. Specifically, what mechanisms does the ITSRA have in place to adequately respond to an insider who has violated any security controls? For any clear violations or even criminal acts, the organization should define clear, unequivocal repercussions at the Business level in the form of policies. These policies should include specific actions to be taken, most commonly employee sanctions or termination, in the event of malicious insider activity.

There may be cases, however, where the organization may have due cause to monitor a particular employee's activities after a policy violation takes place and even if sanctions are enforced. For example, many of the insiders in our database did indeed receive sanctions from leadership, but these did not deter them from committing their crimes. Rather, in some cases, insiders were further provoked and angered by these sanctions and were emboldened to carry out their attacks. In these cases, the organization should have a targeted monitoring policy in place that allows it to selectively monitor both online and at-work activities of any individual if warranted. Of course, such a policy will have to be crafted and approved by legal counsel to ensure that it conforms to local laws and that it does not violate privacy rights of individuals.

Performing targeted monitoring does not require any additional investment in infrastructure. Rather, with the tools in place that already conform to the ITSRA, the organization will have the ability to fine tune its security devices to observe any person's activities. Empirical data shows that employing such targeted monitoring might have prevented many insiders from causing damage to their respective companies [Cappelli 2009, Hanley 2011, Moore 2009]. In one case, an insider was the subject of many complaints from fellow workers, who reported inappropriate behavior including workplace intoxication and sexual harassment. Although the HR department did issue a formal reprimand to the employee, they did not inform the Information Security (IS) department of this individual's actions. Since the IS department had no cause to monitor this user's online activities, they failed to detect his planting a logic bomb in the company's infrastructure, which threatened to destroy several years' worth of critical data. The insider also had access to backup media, whose sabotage could have prevented the organization from restoring its normal operations in a timely manner.

In several other cases, insiders had previous criminal records for the very crimes they went on to commit again [Cappelli 2009]. In these cases, background checks failed to reveal the insiders' criminal histories, and even when they did, personnel management did not inform the appropriate

security departments to keep a closer eye on each person's activities. What we discovered in such cases was the existence of a figurative barrier of communication between respective departments, and an inordinate reluctance especially on the part of HR or personnel management to share any negative information about employees with any other party. The ITSRA seeks to break down these barriers where appropriate and necessary, and to enable a free flow of communication across all departments within an organization.

6 Sample Instantiation of ITSRA: Theft of Intellectual Property

As we mention in Section 3, the insider threat database contains more than 700 cases of insider crimes [Cappelli 2009, Hanley 2011]. Of that 700, roughly 90 of these cases deal with the theft of intellectual property (IP). To illustrate the utility of the ITSRA in addressing theft of IP, we consider the problem of an organization attempting to mitigate the risk of loss of its proprietary data [Moore 2009].

Our data on theft of IP shows that well over 50% of the insiders who stole information did so within 30 days of resignation [Cappelli 2009, Moore 2009]. Current case trends suggest that organizations regularly fail to detect theft of IP by insiders, and even when theft is detected, organizations find it difficult to attribute the crime to any specific individual. Monitoring employee behavior for suspicious actions worthy of further investigation can be expensive. These costs must be balanced against the risk of losing the organization's IP. Organizations need to ensure they do not violate employees' privacy rights and have a valid ownership claim to the IP they wish to protect. Given these challenges facing an organization, we propose a solution described in Section 6.1 below.

6.1 Solution

The solution is described in the context of Figure 6 below. Relationships are indicated as labeled arrows between distinct groups (e.g., HR) or bodies of information (e.g., critical IP). Relationships that exist as part of a sequential ordering are numbered accordingly. The relevant layer of the ITSRA is superimposed upon this diagram in Figure 7 to illustrate how the ITSRA addresses each component of the solution. As this figure illustrates, the principle of acceptable use is the most appropriate for application to theft of IP situations. This is because insiders who committed theft of IP most often had authorized access to the data they removed from the organization.
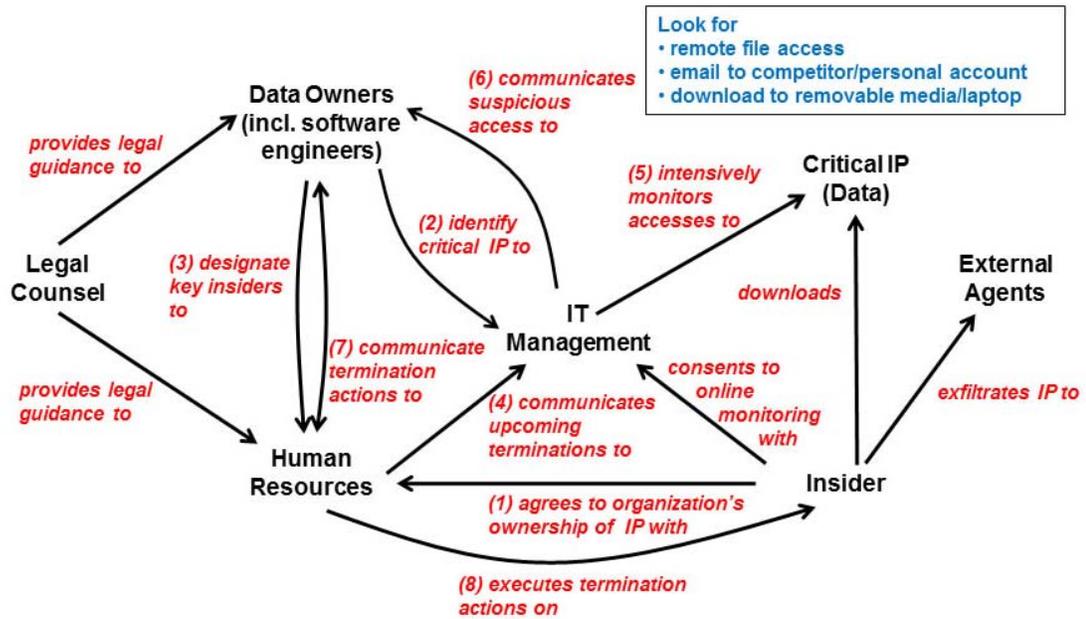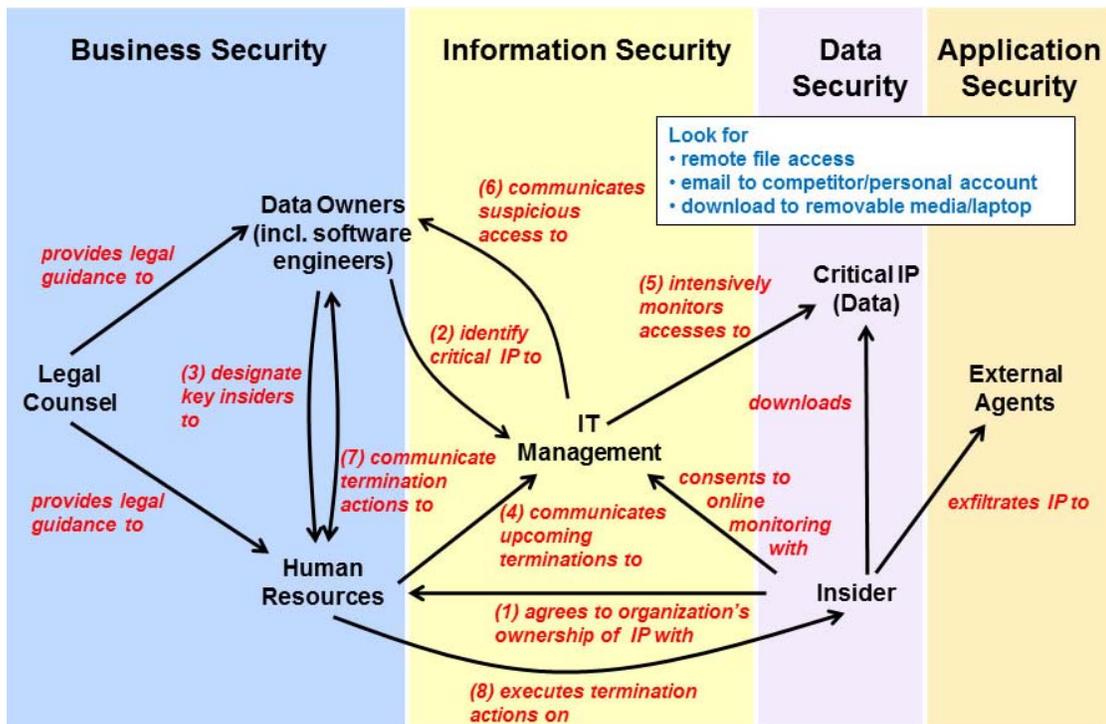
Figure 6:   Theft of IP Pattern

Figure 7:   Theft of IP Pattern with ITSRA Superimposed

At the Business Security layer, an organization needs to make sure its employees agree, as a condition of employment, that the organization owns the critical IP (see Relationship 1 in Figure 6). The employee's clear and formal acceptance of the organization's IP ownership helps ensure that the organization's right to ownership will stand up in court. Consulting with the organization's legal counsel will help ensure the organization is on firm legal ground. The organization can convey ownership to employees through devices such as nondisclosure agreements, IP ownership policies, and references to IP ownership in a network-acceptable-use policy.

At the Data Security layer, data owners need to identify and properly label their IP. They need to communicate the existence and sensitivity of the IP to IT Management (Relationship 2). They also need to communicate to HR the key insiders with access to critical IP (Relationship 3). Our data shows that scientists, engineers, programmers, and salespeople are especially likely to steal IP.

At this point, the principles of targeted monitoring described in Section 5.1 above come into play. HR needs to track insiders who have access to the IP so that, when the insider resigns, HR can notify IT Management to monitor that insider's online behavior for signs of suspicious exfiltration of IP (Relationship 4).

IT Management needs to take action concordant to the controls implemented in the Information Security (monitoring mechanisms, such as a SIEM) and Application Security layers of the ITSRA—in this case, email is the application to be monitored. In particular, IS staff should closely monitor the insider's access to critical IP during the 30-day window around termination (Relationship 5) because many IP thieves stole information within this window [Hanley 2011]. Although the organization may decide to monitor beyond the 30-day window, restricting monitoring to this period may allow the organization to balance the monitoring costs with the risks of losing the IP. No matter what level of monitoring is used, organizations must ensure that insiders are treated consistently and fairly. Typically, insiders need to consent to monitoring of their online actions as a condition of using the organization's systems, consistent with business and legal requirements previously defined at the Business Security layer.

Investigation and response activities may be necessary if IT Management discovers suspicious activity by the insider. During the 30-day window, several items may warrant a detailed investigation:

- **Download of a large volume of critical IP to removable media/laptop or via remote file access:** Large-volume downloads close to insider termination may indicate that the insider is preparing to exfiltrate data. Case information suggests that users who exfiltrate a large amount of information via email or other means first move that data over the network to their workstation. Movement of data within enclaves or across enclaves that exceeds normal traffic patterns may signal this type of event.

- **Email to the organization's competitors or the insider's personal account:** Most insiders who steal information through networked systems do so by either emailing information off the network through a corporate account or through webmail. Corporate email accounts can be configured to alert the organization to suspicious events from mail transaction logs. For example, if an organization enumerates (but does not blacklist) suspicious transactions, such as data transfers to competitors, then it can be alerted to any mail traffic generated to/from

the outgoing employee, particularly if these messages appear to have attachments or have relatively large byte counts. Further, many insiders email IP to their personal webmail accounts and then forward it to an outside collaborator. For example, a user can simply open a browser, log into a personal Gmail account, attach documents, and send them off the network. Organizations need to consider monitoring for uploads to known webmail domains to mitigate these behaviors [Hanley 2011].

Data owners need to be informed of any suspicious access to critical IP and be included in the response decision-making (Relationship 6). The organization needs to be able to either block exfiltration or detect it and confront the employee. If the suspicious activity occurs prior to termination, HR and the data owners need to formulate an appropriate response as part of the termination actions (Relationship 7). The organization can then confront the employee with that response during the exit (termination) interview. If the insider has violated an agreement regarding IP, the organization may wish to pursue legal remediation, with advice from its legal counsel.

Figure 8 below isolates the "Acceptable Use" column of the ITSRA matrix. As with the "Authorized Access" column shown in Figure 5, security controls for acceptable use should span all layers of the ITSRA. High-level policies and requirements should begin at the Business Security layer, and subsequent controls should be implemented at all other layers below.
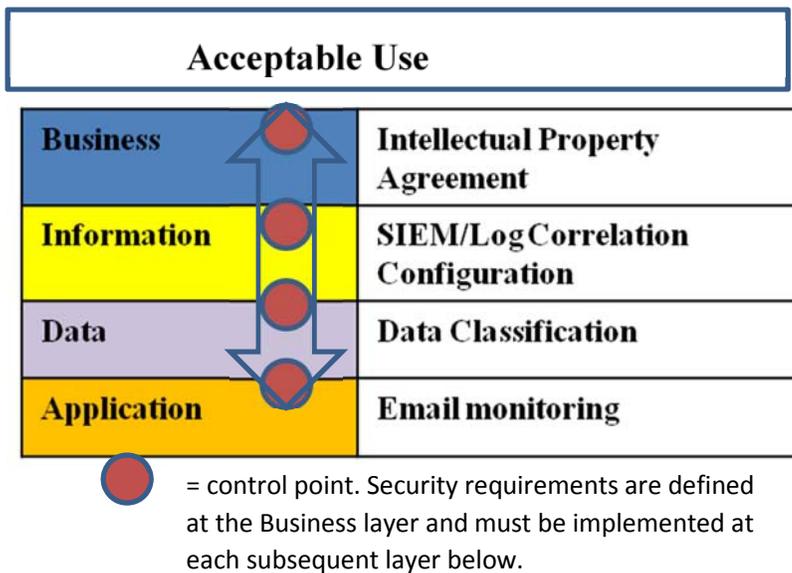
Figure 8: ITSRA Acceptable Use Controls for Theft of IP

Expanding on these forces, monitoring employee behavior to identify suspicious actions worthy of further investigation can be an expensive proposition. These costs must be balanced against the risk of losing the organization's intellectual property. Organizations need to ensure they do not violate employee's rights of privacy and have a valid ownership claim to the IP that they wish to protect.

7 Conclusion

The immediate goal of the ITSRA is to provide an end-to-end architecture that provides actionable guidance to minimize the threat of malicious insider actions. It is intended to support

- senior leadership
- system implementers
- network architects
- Security Operations Center (SOC) operators
- HR managers
- security guards
- enterprise architects

The ITSRA is a dynamic model that will adapt to the continuously changing climate of business processes, information technologies, and security practices.

Insider threat is a serious problem, and our research has demonstrated the urgent need for a documented, comprehensive, standard Insider Threat Security Reference Architecture. Therefore, we have designed a reference architecture based on current state of hardware, operating systems, and networking infrastructures. However, the ITSRA is intended to serve another key function as well: though the ITSRA will be designed to solve a current problem using current technologies, it will also form the foundation for the next generation of operating systems, applications, and information infrastructure. It is critical that developers of next-generation technologies consider insider threat mitigation requirements at the initial design stages. Next-generation hardware, operating systems, and networking infrastructures should have the ITSRA designed into the individual components. Otherwise, we will find ourselves in the same position down the road as we are in now: attempting to solve the insider threat problem with a collection of policies, processes, and applications bolted onto available components rather than built into the infrastructure. The ITSRA will provide a common foundation for those next-generation engineers to provide a more effective solution in the future.

References

URLs are valid as of the publication date of this document.

[Blackwell 2009]
Blackwell, C. "A Security Architecture to Protect against the Insider Threat from Damage, Fraud, and Theft." *Proceedings of the 5th Annual Workshop on Cyber Security and Information Intelligence Research Workshop: Cyber Security and Information Challenges and Strategies* (CSIIRW '09). Knoxville, TN, USA, April 2009. ACM, 2009.

[Cappelli 2009]
Cappelli, D.; Moore, A.; Trzeciak, R.; & Shimeall, T. *Common Sense Guide to Prevention and Detection of Insider Threats, 3rd Edition – Version 3.1.* Software Engineering Institute, Carnegie Mellon University, 2009.

[CIOC 2001]
Chief Information Officer Council. *A Practical Guide to Federal Enterprise Architecture.* Federal Architecture Working Group (FAWG), 2001.

[Chung 2010]
Chung, J.; Pueblas, M.; Nadimi, A.; Hamilton, D.; & Farrington, S. *Cisco SAFE Reference Guide: Cisco Validated Design.* Cisco Systems, Inc., 2010.

[EOPUS 2007]
Executive Office of the President of the United States. *FEA Consolidated Reference Model Document Version 2.3.* Office of Management and Budget, 2007.

[Hanley 2011]
Hanley, M. *Deriving Technical Controls and Indicators of Insider Attack from Socio-Technical Models and Data* (CMU/SEI-2011-TN-003). Software Engineering Institute, Carnegie Mellon University, 2011. http://www.sei.cmu.edu/library/abstracts/reports/11tn003.cfm

[Jabbour 2009]
Jabbour, G. & Menascé, D. "The Insider Threat Security Architecture: A Framework for an Integrated, Inseparable, and Uninterrupted Self-Protection Mechanism." *Proceedings, 12th IEEE International Conference on Computational Science and Engineering* (CSE 2009). Vancouver, Canada, August 2009. IEEE, 2009.

[Microsoft 2010]
Microsoft Corporation. *Planning and Architecture for SharePoint Server 2010.* Microsoft, May 12, 2010. http://technet.microsoft.com/en-us/library/cc261834.aspx

[Moore 2009]
Moore, A.; Cappelli, D.; Caron, T.; Shaw, E.; & Trzeciak, R., "Insider Theft of Intellectual Property for Business Advantage: A Preliminary Model," 1-1. *Proceedings of the 1st*

International Workshop on Managing Insider Security Threats (MIST 2009). Purdue University, West Lafayette, IN, June 2009. Purdue University, 2009.

[NIST 2009]

National Institute of Standards and Technology. *Special Publication (SP) 800-53, Revision 3, Recommended Security Controls for Federal Information Systems and Organizations*. NIST, 2009.

[NIST 2010]

National Institute of Standards and Technology. *Special Publication (SP) 800-37, Revision 1, Guide for Applying the Risk Management Framework to Federal Information Systems: A Security Life Cycle Approach*. NIST, 2010.

[Oracle 2011]

Oracle Corporation. *Oracle Database Security Overview*. Oracle, 2011. http://www.oracle.com/technetwork/database/maximum-security-architecture-094265.html

[SABSA 2011]

The SABSA Institute. *Sherwood Applied Business Security Architecture*. SABSA, 2011. http://www.sabsa-institute.org

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE April 2012	3. REPORT TYPE AND DATES COVERED Final

4. TITLE AND SUBTITLE	5. FUNDING NUMBERS
Insider Threat Security Reference Architecture	FA8721-05-C-0003

6. AUTHOR(S)

Joji Montelibano, Andrew Moore

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Software Engineering Institute Carnegie Mellon University Pittsburgh, PA 15213	8. PERFORMING ORGANIZATION REPORT NUMBER CMU/SEI-2012-TR-007

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) HQ ESC/XPK 5 Eglin Street Hanscom AFB, MA 01731-2116	10. SPONSORING/MONITORING AGENCY REPORT NUMBER ESC-TR-2012-007

11. SUPPLEMENTARY NOTES

12A DISTRIBUTION/AVAILABILITY STATEMENT Unclassified/Unlimited, DTIC, NTIS	12B DISTRIBUTION CODE

13. ABSTRACT (MAXIMUM 200 WORDS)

The Insider Threat Security Reference Architecture (ITSRA) provides an enterprise-wide solution to insider threat. The architecture consists of four security layers: Business, Information, Data, and Application. Organizations should deploy and enforce controls at each layer to address insider attacks. None of the layers function in isolation or independently of other layers. Rather, the correlation of indicators and application of controls across all four layers form the crux of this approach. Empirical data consisting of more than 700 cases of insider crimes show that insider attacks proved successful in inflicting damage when an organization failed to implement adequate controls in any of three security principles: authorized access, acceptable use, and continuous monitoring. The ITSRA draws from existing best practices and standards as well as from analysis of these cases to provide actionable guidance for organizations to improve their posture against the insider threat.

14. SUBJECT TERMS insider threat, reference architecture, espionage, fraud, theft	15. NUMBER OF PAGES 29

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL