21st National Information Systems Security Conference

PRACTICAL STEPS TOWARD PROVIDING A COST-EFFECTIVE SECURITY ARCHITECTURE BASED ON THE COMMON CRITERIA

Bernard Ramsey, Jr.
Federal Aviation Administration

Abstract: In developing security for a complex distributed system in the Federal Aviation Administration, the Common Criteria were used as the basic guidance document for writing a Protection Profile. The Neutral Security Architecture Model has been developed to produce a cost-effective and tailored information security solution for the system and to ensure that the requirements in the Protection Profile are complete and comprehensive. A second model, the Risk Mitigation Factor Model, describes steps to rank and evaluate risk mitigation alternatives.

INTRODUCTION

A federal government agency developing an INFOSEC program is subject to general federal direction and guidance[1], as well as agency-specific regulations. This material, while helpful as a checklist of what services should be provided, offers little in the way of concrete information on how to implement the functional requirements. Another difficulty has been the lack of analytical tools to present the results of security analysis to management in a way that provides a cogent business rationale for the security program.

Two things are needed:

- The first is a set of generic procedures that security professionals can follow to identify threats, perform assessments, evaluate risk mitigation techniques, and ensure that all relevant requirements are included in the Protection Profile.
- The second is a management tool to quantify the costs of residual risk and alternative mitigation techniques.

This paper proposes two models to address these needs. The techniques have been developed in support of a Federal Aviation Administration (FAA) program and will be validated in practice during calendar year 1998. It is expected that the results of the experience will be reported in subsequent reports to the INFOSEC community.

BACKGROUND

The National Airspace System Infrastructure Management System. The National Airspace System (NAS) Infrastructure Management System (NIMS) comprises the general infrastructure

[1] These include the Computer Security Act of 1987, which provides a general mandate for security; Office of Management and Budget (OMB) Circular A-130, Appendix III, "Security of Federal Information," which requires federal agencies to prepare security programs; and the Federal Information Resources Management Regulation (FIRMR), which regulates the use of computer resources in federal agencies.

support systems for the NAS—including local area networks, wide area networks, personal computers, servers, networks, and various information technologies—connected in a nationwide network that monitors and controls the health of air traffic control equipment. The NAS infrastructure includes over 38,000 individually-managed systems located throughout the United States, communicating via an infrastructure not completely under FAA control.

Building a Protection Profile for NIMS. The protection profile for NIMS was the first document of its kind developed in the FAA. The collaborative process that led to its completion is expected to stand as a benchmark for further security efforts in the agency.

The profile is structured to track closely the Common Criteria, v.1, and it is compliant with applicable agency and federal regulations and laws. The profile is written in language that makes it accessible to the layman and it has been used as a vehicle to generate stakeholder buy-in for the security program. The document underwent several reviews by headquarters and field personnel, INFOSEC professionals, FAA requirements experts, and NIMS Product Team members. As a result, the NIMS protection profile enjoys broad acceptance by the community it serves.

Implementing the Protection Profile. Perhaps more than any other aspect of federal acquisition, INFOSEC is perceived as expensive and operationally complex. To compound the problem, the art of estimating the costs of security solutions is still in its infancy. As a result, security often operates in a budget “cloud” of programmatic and budgetary uncertainty.

When developing NIMS security, it was decided that one way to ensure that security was integrated into the systems under development in a cost-effective way was to develop a “neutral” security architecture.

NEUTRAL SECURITY ARCHITECTURE

A neutral security architecture is a set of generic security services that must be employed to manage risk to a system and that are not tied to a specific technology, vendor, or cost assumption. A neutral security architecture changes as the overall system architecture evolves.

The security architecture moves toward the system architecture *over time*:

- Imposing a security solution at the beginning of system development will tend to “drive” the system architecture, narrowing the options for innovations, and potentially ruling out many COTS products.
- Developing security solutions after the system architecture is fully mature can be very expensive *and may also prove impossible of practical integration so that some security requirements may not be implemented, leaving the system exposed to residual risks*.

The solution is to take security development one step at a time (see figure 1). Starting with a neutral security architecture, the security solutions are tailored to the system architecture as it develops.

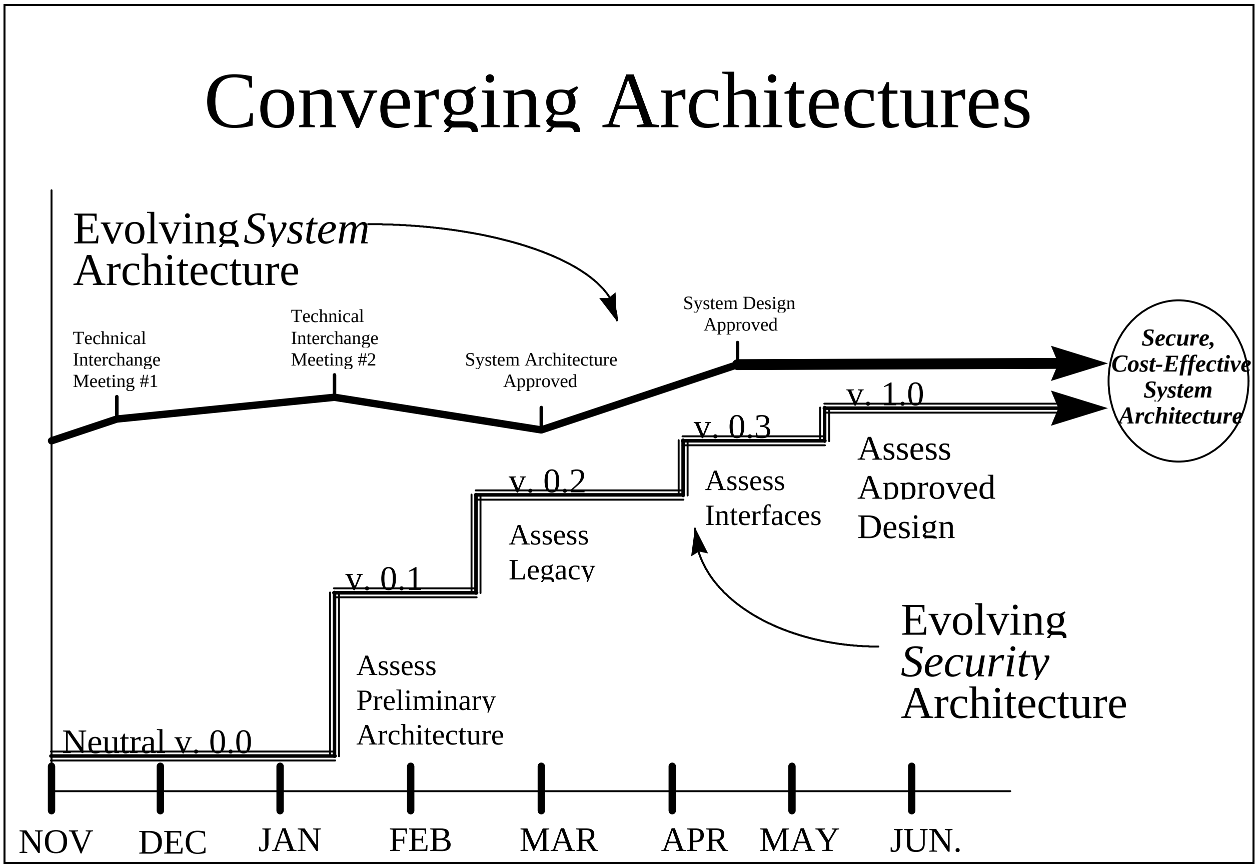
Converging Architectures
Evolving System Architecture
Technical Interchange Meeting #1
Technical Interchange Meeting #2
System Architecture Approved
System Design Approved
Secure, Cost-Effective System Architecture
Neutral v. 0.0
v. 0.1
v. 0.2
v. 0.3
v. 1.0
Assess Preliminary Architecture
Assess Legacy
Assess Interfaces
Assess Approved Design
Evolving Security Architecture
NOV
DEC
JAN
FEB
MAR
APR
MAY
JUN.

Figure 1. Converging Architectures

The architecture and the lower-level requirements in the protection profile are modified in a series of incremental assessments. These assessments are planned to take place at significant milestones in the developing system architecture and/or when new risk information becomes available. Unresolved issues are carried over to the next cycle as better information becomes available. These “cycles” of assessments use the classical methodology of risk management:

- Assess threats
- Assess vulnerabilities
- Validate assessments
- Develop countermeasures
- Complete risk mitigation with market studies, cost benefit analysis, and recommended COTS solutions.

If a requirement is identified during the assessment process that does not appear in the protection profile, it is added to the profile. The new set of requirements is then used in succeeding assessments. (It should be noted that assessments will not necessarily identify new requirements.)

The object is to upgrade the neutral security architecture and the protection profile incrementally until the architecture is system-specific. Fewer labor hours are required to perform incremental assessment cycles (to save money) and the work is performed parallel to the developing system architecture (to save time).

In this way, security becomes a flexible, but integral part of system planning. Security follows the lead of the overall system architecture, without driving the system design. Over the course of development of the system, the neutral security architecture becomes more and more closely tailored to the unique features of the system design, producing cost and schedule benefits.

THE RISK MITIGATION FACTOR MODEL

It is often difficult to "sell" security to management, because it is hard to quantify the value of their investment. The model below offers a way to rank and evaluate mitigation alternatives by combining quantitative standards with professional judgment and experience.

FAA Order 1600.66 (Telecommunications and Information System Security Policy, Appendix 2, "Threats to Information") lists specific threats that must be addressed in systems and networks throughout their life cycles. Four categories of threats are listed:

- common threat agents,
- inappropriate disclosure threats (confidentiality violations),
- fault-and-error threats (integrity violations), and
- loss of service threats.

For this system architecture, there are more than one hundred individual threat agents in the four threat categories. The number of threats and threat categories will vary for other systems based on operational needs.

Vulnerability Ratings. In the Risk Mitigation Factor Model, each threat category is assigned a rating based on two elements: the potential impact of the security violation on functional operations (severity of the hazard) and the probability that the violation will occur. These are the two elements specified in FAA guidance for assessing hazard risk for sensitive application certification[2].

[2] In accordance with FAA Order 1600.54B, Chapter 12.

The *severity of the hazard* is ranked on a four-point scale:

	Hazard Severity	
Value	*Severity*	*Characteristics*
4	Critical I	Extensive system loss and service interruption
3	Critical II	Severe system loss and service interruption
2	Marginal	Minor system loss and interruption
1	Negligible	Less than severe system loss or service interruption

A *probability* ranking is assigned on a five-point scale:

	Hazard Probability	
Value	*Description*	*Characteristics*
4	Frequent	Likely to occur frequently
3	Probable	Will occur several times during life cycle
2	Occasional	Likely to occur sometime during life cycle
1	Remote	Unlikely but possible to occur in life cycle
0	Improbable	So unlikely it can be assumed that occurrences may not be experienced

For each threat, the current algorithm asks the question: does this threat element apply to a particular vulnerability of the security target and, if so, how does it affect services? After values for the probability and severity of the threat are identified, the algorithm calculates the potential impact on services. The services included in the current algorithm are data integrity, data criticality, impact of loss, and availability of vulnerability information. Additional service elements will be added to the algorithm as the model moves toward maturity. The output of the algorithm is a *risk factor* for each of the threats.

Values are derived as follows. Each service is assigned a rating obtained from user surveys. When we perform the vulnerability assessments, users will be asked to complete a questionnaire to rate each of the elements (data integrity, data criticality, impact of loss, and availability of vulnerability information) on a 1-5 scale to describe the potential negative impact on the component or system. The arithmetic mean of responses to the questionnaire will be used in the model. The algorithm will be the product of the threat elements (probability x severity) multiplied by the sum of four service elements, divided by the total risk (sum of the maximum value for the service elements).

$$\text{Risk Factor} = (H_{sev})(H_{prob})(D_{integ} + D_{crit} + I + A)\ (1/R_{sum})$$

Note that if the probability of the occurrence of the threat is zero, the risk factor is also zero.

A example of how the algorithm works is shown below. A hypothetical vulnerability has been assumed. Values were assigned for illustration purposes only and do not relate to any actual vulnerability.

High-level Example for a Hypothetical System Vulnerability

Hazard risk rating[3]:

Common threat agents category
- Severity (H_{sev}): 4
- Probability (H_{prob}): 2

Impact on services[4]

Data integrity (D_{integ}): 4
- No impact (1)
- Minimal impact (2)
- Moderate impact (3)
- Substantial impact (4)
- Critical impact (5)

Data criticality (D_{crit}): 2
- Low; local importance (1)
- Moderate; regional importance (2)
- High; national importance (3)
- Critical; national NAS and defense importance (4)

Impact of loss (I): (3)
- No impact (1)
- Minimal impact due to redundancy or workarounds (2)
- Moderate impact; delays in service delivery (3)
- Substantial impact; major delays (4)
- Critical impact; denial of service (5)

Availability of information on vulnerabilities (A): 1
- Not available outside FAA (1)
- Restricted availability (2)
- Available upon request from FAA or vendor (3)
- Available in widely distributed written materials (4)

[3] In an actual assessment, these values would be chosen by the security manager, based on FAA guidelines.

[4] In an actual assessment, these values would be derived from a user survey.

Freely available on Internet or in public materials (5)

Common threat agent risk factor[5] = (4)(2)(4 + 2 + 3 + 1)/19= 4.21

Following the same process, (hypothetical) *risk factors* are calculated for each threat category. The *total risk* for all threat categories in this example is 13.63 (the sum of the risk factors):

- common threat agents: 4.21
- inappropriate disclosure threats (confidentiality violations): 0.42
- fault-and-error threats (integrity violations): 4.74
- loss of service threats: 4.26.

Dividing the threat category risk factor by the total risk gives us the *risk mitigation factor[6]* for that threat category, which is expressed as a percentage of total risk. For instance, the fault-and-error threat category has a risk mitigation factor of .3477 and represents 34.77 percent of the total risk to the identified vulnerability.

The risk mitigation factor tells us how much of the risk would be eliminated by "removing" the particular threat or threat category, in this case, by applying a countermeasure. In other words, if the "loss of service" threat agents could be completely mitigated by a given security solution, 31.25 percent of the risk to the system would be removed.

The risk mitigation factors and the total risk are unique to the system architecture; other systems will assess threats to their system's vulnerabilities quite differently. Each organization will need to generate its own security threat analysis.

<u>Risk and Cost</u>. Calculating the risk mitigation factors enables us to identify what proportion of the risk is mitigated by a given countermeasure. The results of such an analysis can be graphed against the cost of the countermeasure to show the cost of eliminating that risk. The following graph is a hypothetical example.

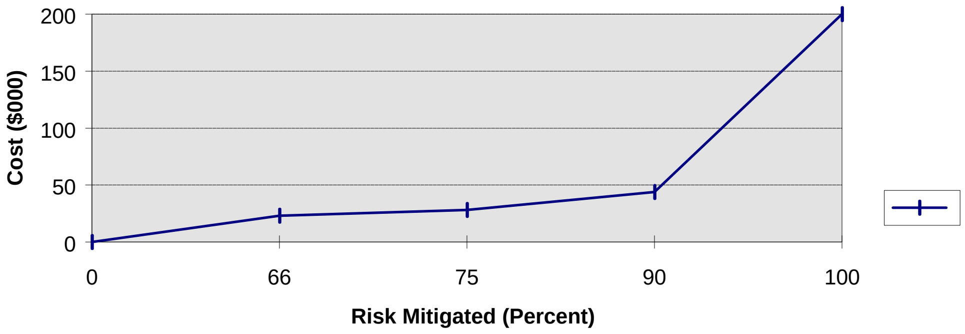

[5] The factor "19" is the sum of the maximum values for each of the service elements (5+5+4+5).
[6] The sum of the threats is 1.00 (may not add to total due to rounding).

To compare alternative solutions, we can show the information in a table. Analyzing what risks are mitigated will enable management to make an informed judgment about the cost of security and about the acceptability of the residual risk.

Product	Percentage of Risk Mitigated	Cost
A	53.9%	$ 10,000
B	32%	$ 7,500
C	28.3%	$ 55,000
D	40%	$ 50,000
E	55.5%	$200,000

Cost information must be comprehensive, including life cycle costs for the initial acquisition, plus training, maintenance, installation, and other costs.

More Detailed Analysis. The same method could be applied to individual threats within the threat categories to produce a ranking of threats and a risk mitigation factor for each. However, using only the two hazard criteria in use here (severity and probability) would not allow sufficient precision to avoid many ties. One approach would be to rank the threats individually based on their severity rating. For example, all "5's" could be ranked as 5.1, 5.2, and so on. A higher number would indicate increased severity. More investigation into the technique is warranted. Improving the precision of the model by addressing individual threats may not produce appreciably more information that would be significant to upper-level management.

In addition to the straightforward ranking strategy described in this paper, we are also investigating the use of the Borda method, a positional voting method for ranking alternatives. This ordinal voting method has been used as the framework for a study of maintenance drivers in a recent U.S. Air Force project and it may hold promise for NIMS information security.[7]

CONCLUSION

We have established two ways to analyze our security direction and to provide support to management in security planning:

- The neutral security architecture model provides a flexible and cost-effective method of tailoring security solutions to an evolving system architecture, without "driving" the system architecture or delaying its implementation; and
- The risk mitigation factor model provides quantitative data to enable upper-level managers to know how much risk is being mitigated by the countermeasure and its cost. The cost includes all life cycle costs from acquisition through disposition. The model also provides a graphical representation of security solutions and costs.

[7] Zachary F. Lansdowne and Beverly S. Woodward, "Applying The Borda Ranking Method," *Air Force Journal of Logistics*, vol. 10, no. 2, pp. 27-29.

Practical Steps Toward Providing a Cost-Effective Security Architecture Based on the Common Criteria

Bernard Ramsey, Jr.
Federal Aviation Administration

21st National Information Systems

Security Conference

October 5-8, 1998

Overview

- National Airspace System (NAS) Infrastructure Management System (NIMS)
- NIMS Protection Profile
- Security and development of a system architecture
- Neutral security architecture
- Risk mitigation model

National Airspace System Infrastructure Management System (NIMS)

- Nationwide network to monitor and control the status of air traffic control equipment
  - 38,000 pieces of equipment
  - COTS system, to the extent possible
- Tiered architecture
  - One National Operations Control Center (NOCC)
  - Three Operations Control Centers
  - 26 Service Operations Centers
  - 300+ Work Centers (WC)

NIMS Operational Architecture

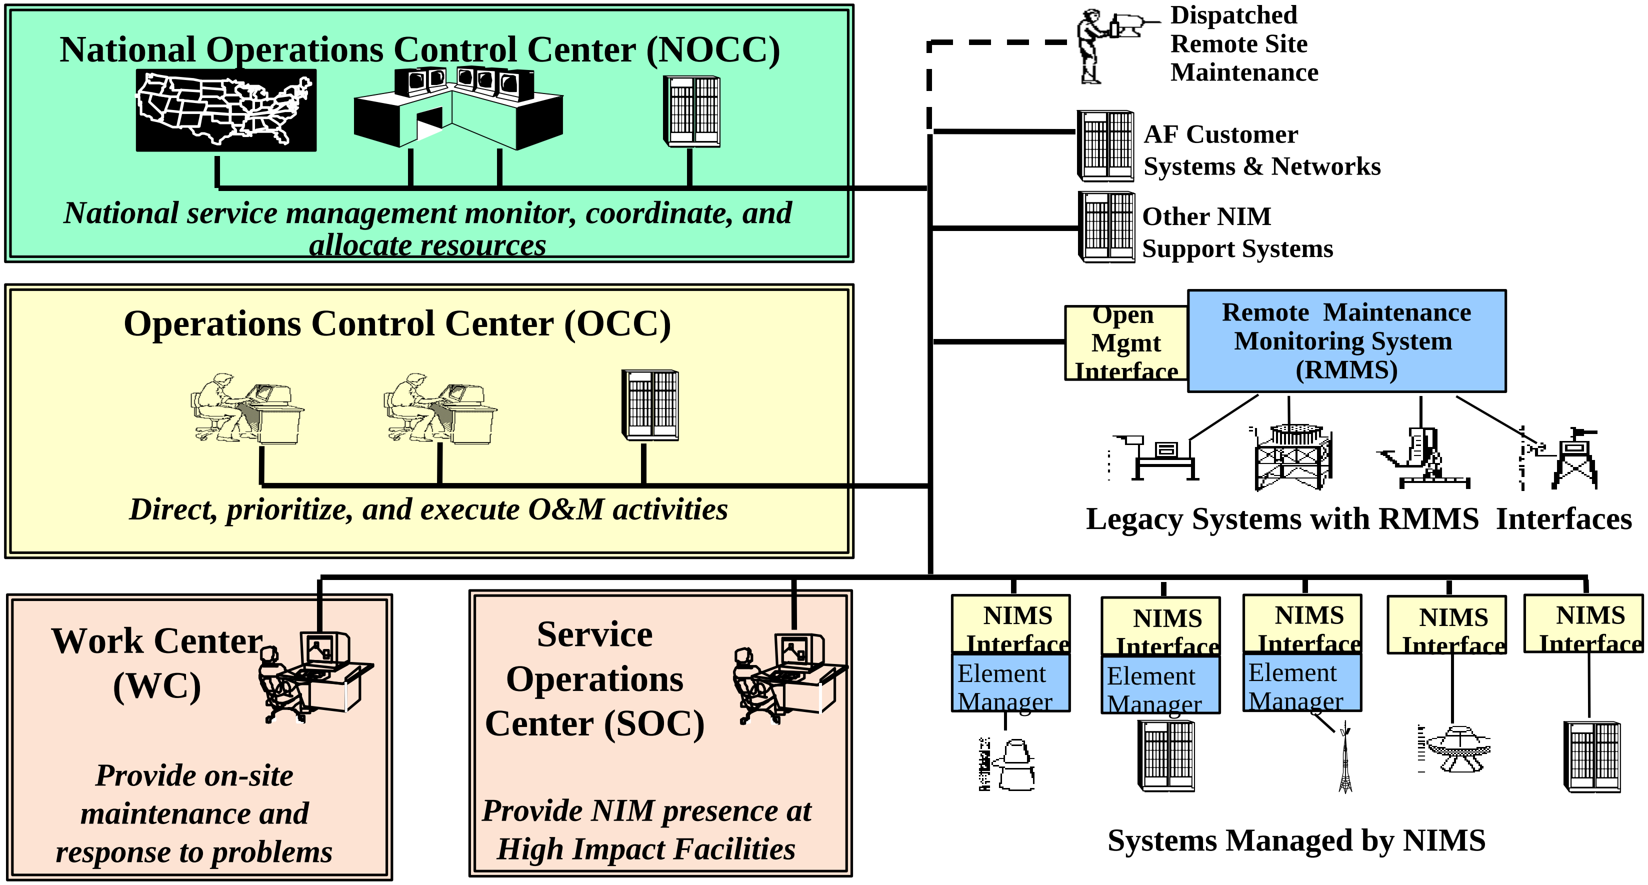

National Operations Control Center (NOCC)

National service management monitor, coordinate, and allocate resources

Operations Control Center (OCC)

Direct, prioritize, and execute O&M activities

Work Center (WC)

Provide on-site maintenance and response to problems

Service Operations Center (SOC)

Provide NIM presence at High Impact Facilities

Dispatched Remote Site Maintenance

AF Customer Systems & Networks

Other NIM Support Systems

Open Mgmt Interface

Remote Maintenance Monitoring System (RMMS)

Legacy Systems with RMMS Interfaces

NIMS Interface

Element Manager

NIMS Interface

Element Manager

NIMS Interface

Element Manager

NIMS Interface

NIMS Interface

Systems Managed by NIMS

NIMS will improve support to overall NAS operations by integrating and standardizing information at centralized locations

Major Types of Connections

- Telco lines
- Leased services
  - Administrative Data Transmission Network (ADTN)-2000
  - Federal Telephone System (FTS)
  - Leased Interfacility NAS Communications System (LINCS)
- Dial-up modem to modem
- Limited distance modem (FAA private line)
- Radio Frequency (RF) modem
- Data Multiplexing Network (DMN) (FAA owned)
- FAASAT (Satellite Network)
- Alaska NAS Interfacility Communications System (ANICS)
- National Airspace Data Interchange Network (NADIN) I&II

NIMS Protection Profile

- First of its kind in the FAA
- Follows the Common Criteria, v.1
- Developed in close coordination with:
  - Headquarters and field personnel
  - INFOSEC professionals
  - FAA requirements experts
- Enjoys broad acceptance
- Reviewed by National Information Assurance Program

Protection Profile Structure

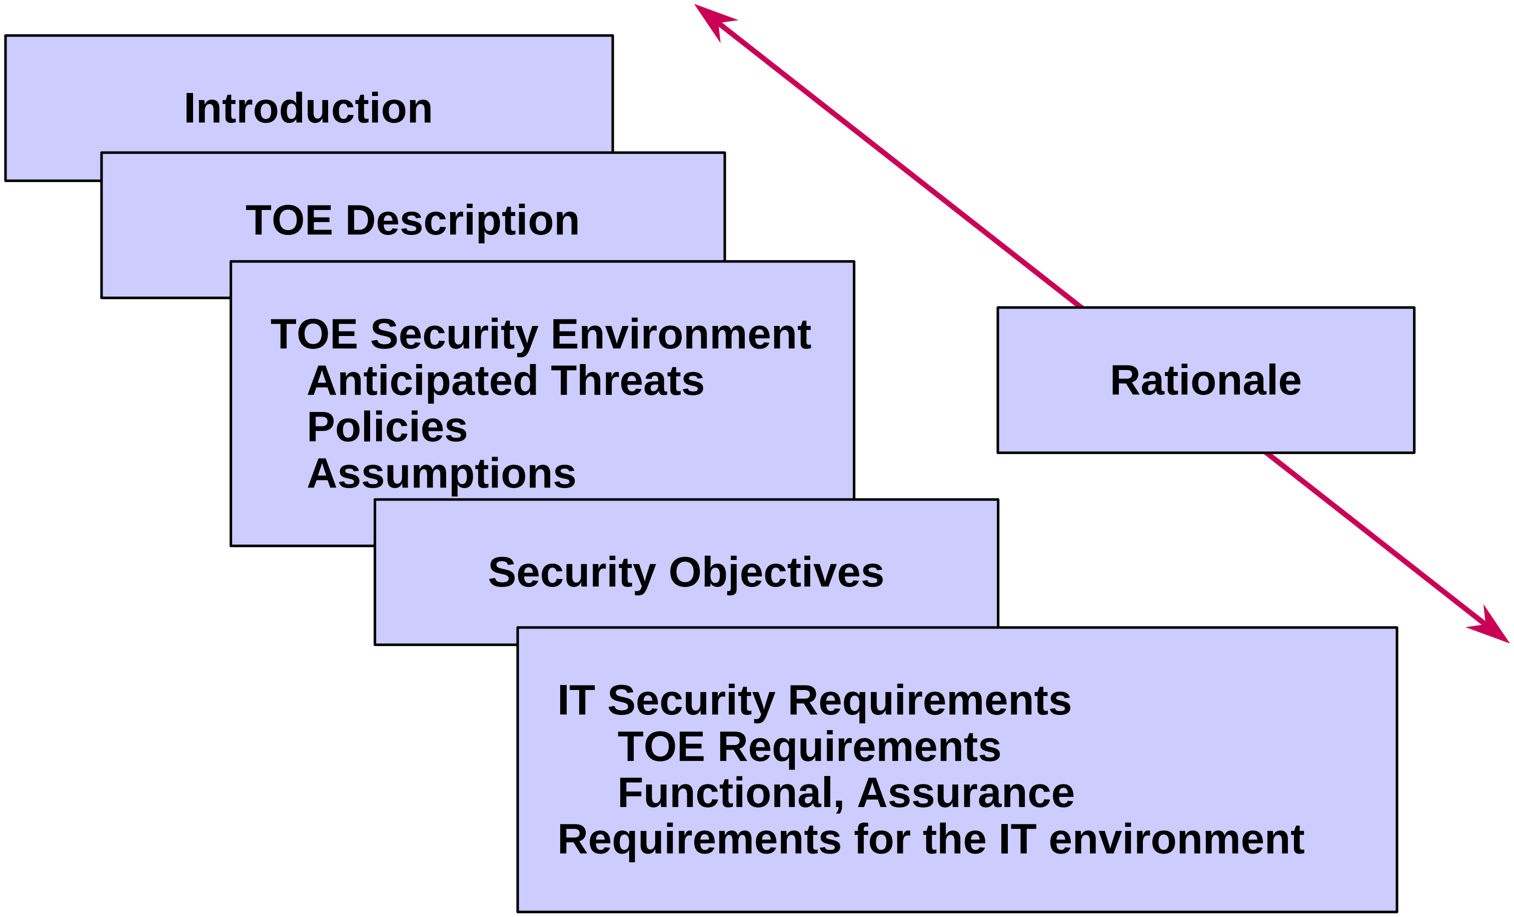

Introduction

TOE Description

TOE Security Environment
- Anticipated Threats
- Policies
- Assumptions

Security Objectives

IT Security Requirements
- TOE Requirements
- Functional, Assurance

Requirements for the IT environment

Rationale

Background on the NIAP Review

- The National Information Assurance Program (NIAP)
  - NSA / NIST alliance provides technology to civil agencies
  - Brings security technology to NVLAP program
- NIMS security management sought NIAP review
  - CC as vehicle for moving FAA to the forefront
  - Improve NIMS security requirements
  - Validate approach
  - Benefit to rest of FAA, other civilian agencies
- NIAP accepted the challenge
  - Entrée into the civil sector
  - Validate Common Evaluation Methodology

Security and the Development of a System Architecture

- *Need to integrate security development into system development at the earliest stages*

Security Development

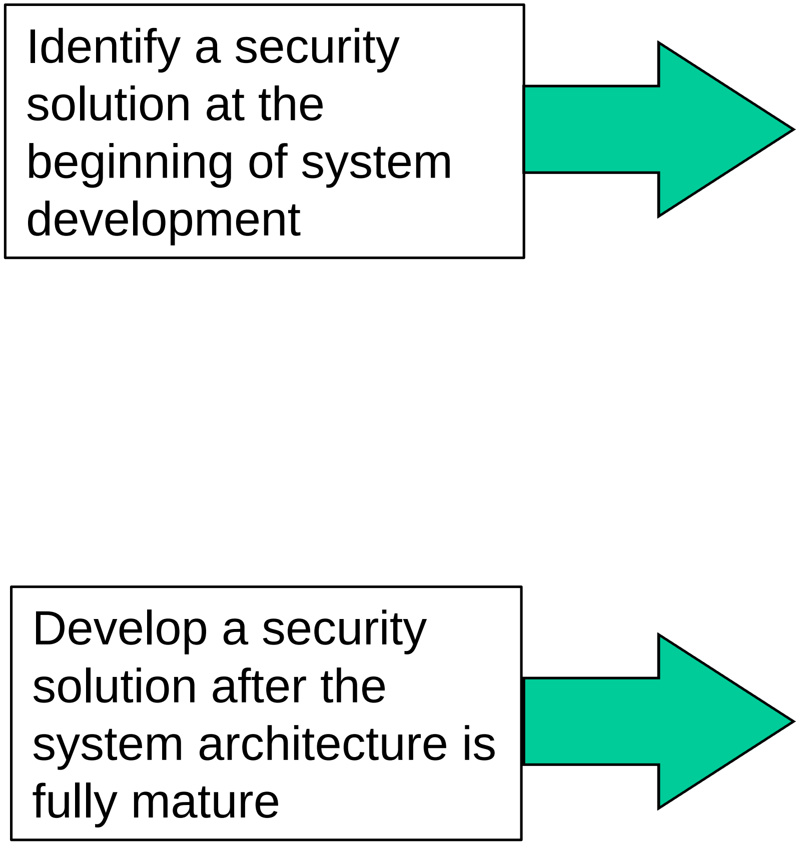

Identify a security solution at the beginning of system development

- May drive system architecture
- Narrows options for innovation
- May rule out some COTS products

Develop a security solution after the system architecture is fully mature

- Potential high cost
- Integration issues

Neutral Security Architecture

- Provides a starting point for development of the security architecture
- Security solution is tailored to the system architecture as it develops
- Generic security services
- Not tied to a specific technology, vendor, or cost assumption

Neutral Security Architecture

- Incremental assessments
  - At significant milestones
  - When new risk information becomes available
- Unresolved issues are carried over to the next cycle
- Upgrade the neutral security architecture and the protection profile incrementally until the architecture is system-specific

Assessment Cycle

- Assess threats
- Assess vulnerabilities
- Validate assessments
- Develop countermeasures
- Complete risk mitigation with market studies, cost-benefit analyses, and recommended COTS solutions
- If new requirements are identified, add them to the protection profile
- ***Repeat cycle of assessments***

Converging Architectures

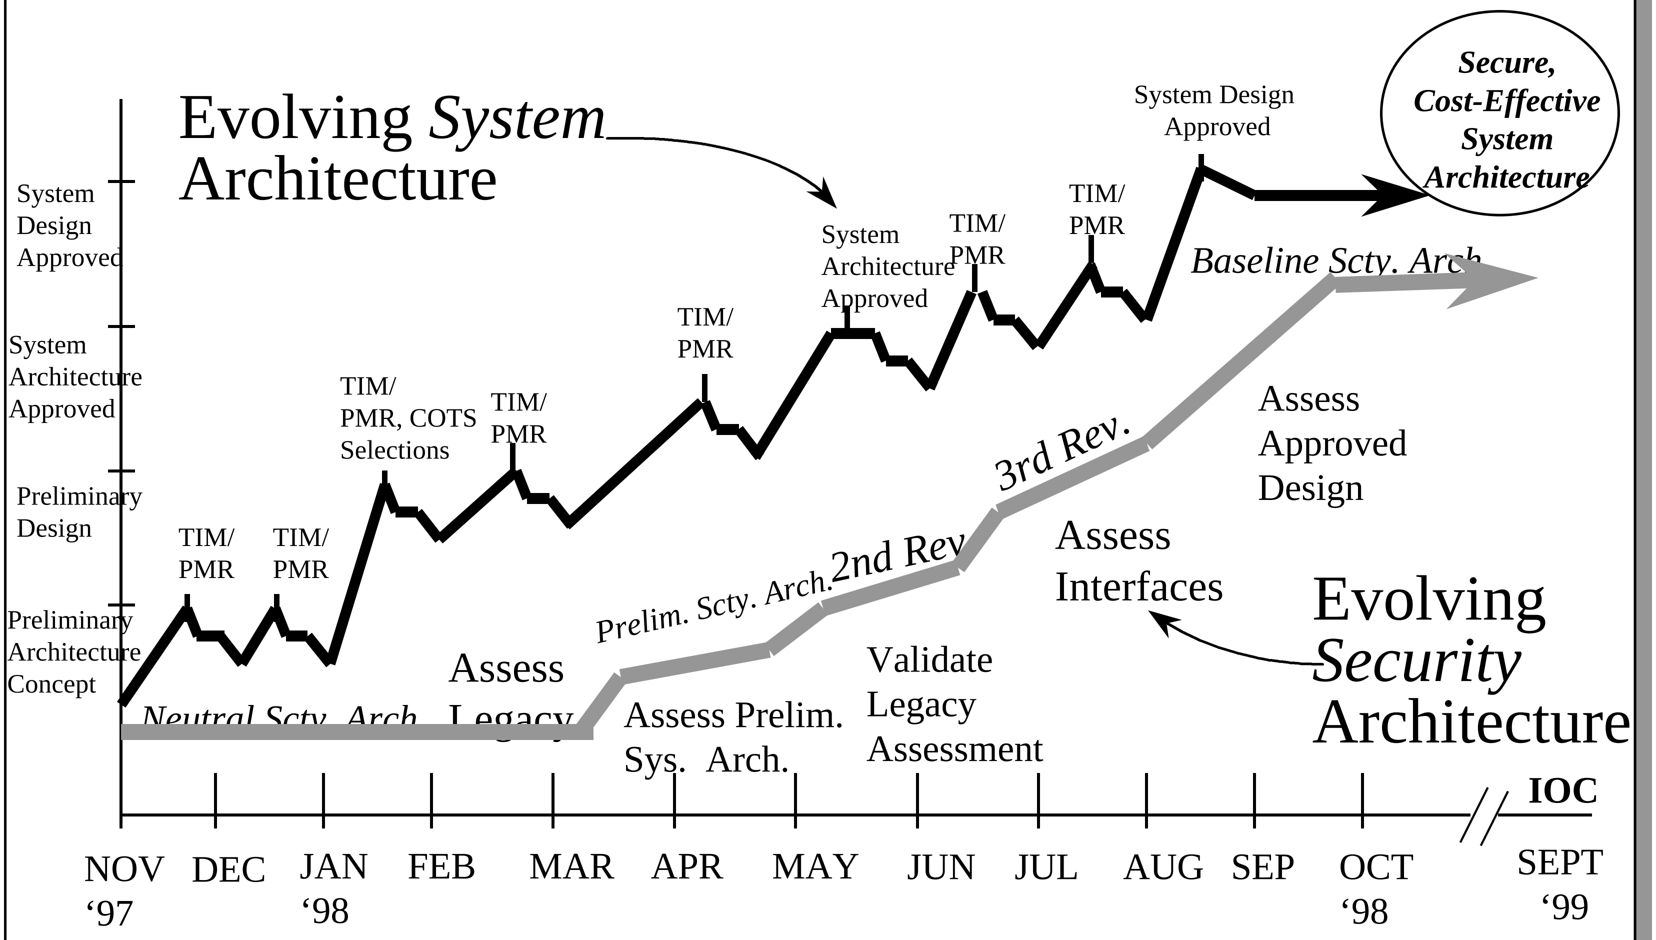

Risk Mitigation Factor Model

It is often difficult to "sell" security to management because it is hard to quantify the value of the investment.

The model offers a way to rank and evaluate mitigation alternatives by combining quantitative standards with professional judgment and experience.

Risk Mitigation Factor Model

- Identify threat categories
- Assign each threat category a vulnerability rating based on:
  - Hazard severity
  - Hazard probability
- User surveys
  - What is the potential impact on services?

Factors

- H_{sev} = hazard severity
- H_{prob} = hazard probability
- D_{integ} = data integrity*
- D_{crit} = data criticality*
- I = impact of loss*
- A = availability of information on vulnerabilities*

* Ratings assigned from information gathered in the user survey

Algorithm

$$\text{Risk Factor} = \frac{(H_{sev})(H_{prob})(D_{integ} + D_{crit} + I + A)}{R_{sum}}$$

Risk Mitigation Factor

- The risk mitigation factor tells us how much of the risk would be eliminated by applying a countermeasure to the particular threat or threat category.
- The risk mitigation factors and the total risk are unique to the system architecture.

Risk and Cost

- Calculating the risk mitigation factors identifies what proportion of the risk is mitigated by a given countermeasure.
- Cost of reducing/eliminating some percentage of risk = cost of the countermeasure.

Risk and Cost: Example

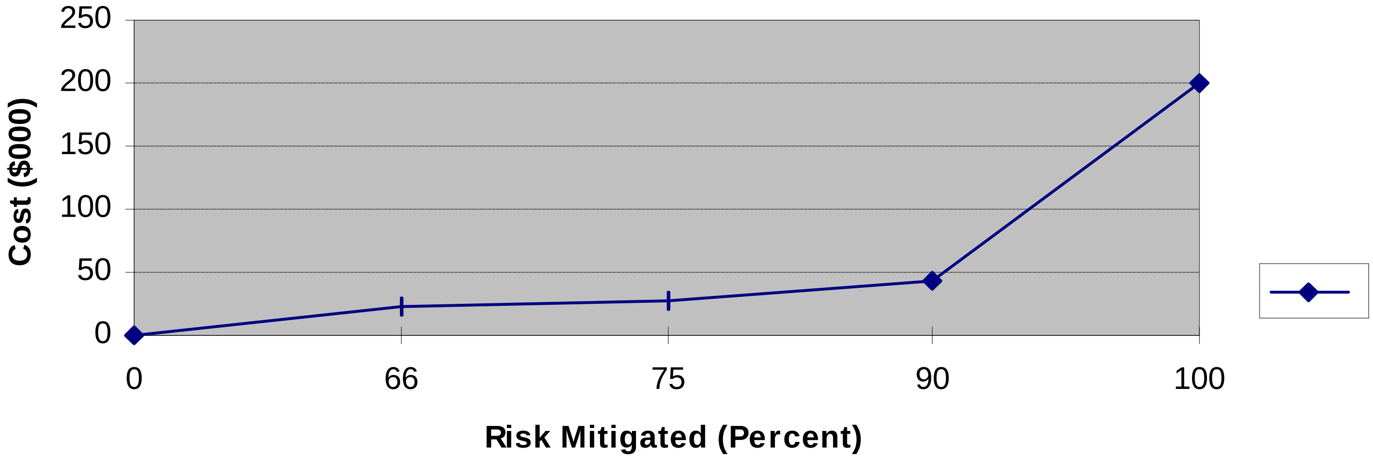

Product	Percentage of Risk Mitigated	Cost
A	53.9%	$ 10,000
B	32%	$ 7,500
C	28.3%	$ 55,000
D	40%	$ 50,000
E	55.5%	$200,000

Conclusion

- The neutral security architecture model provides a flexible and cost-effective method of tailoring security solutions to an evolving system architecture, without “driving” the system architecture or delaying its implementation
- The risk mitigation factor model provides quantitative data to enable managers to know how much risk is mitigated by a given countermeasure and its cost. The cost includes all life-cycle costs from acquisition through disposition. The model also provides a graphical representation of security solutions and costs.