SAND2015-8995C

TE Framework: A Framework for Securing COTs Applications

Melissa Tucker
Sandia National Laboratories
Albuquerque, NM
matucke@sandia.gov

ABSTRACT
The Trust Enhancement (TE) Project at Sandia National Laboratories (SNL) sought to provide a solution for the risk of commercial-off-the shelf (COTS) and government-off-the shelf software (GOTS). As part of the TE Project, a TE Framework was developed to aid organizations in the integration of semi-trusted software. This TE Framework is comprised of two areas: TE Analysis (design assurance, red-teaming, risk assessment, threat analysis and malicious actor definition) and TE Architecture (secure system design, security enclave and container integration and monitoring software). This paper is a follow-on to a white paper written by Subject Matter Experts on the TE Project.

Categories and Subject Descriptors
K.6.5 [**MANAGEMENT OF COMPUTING AND INFORMATION SYSTEMS**]: Security and Protection – *authentication, invasive software, physical security, unauthorized access.*

General Terms
Management, Documentation, Performance, Design, Reliability, Experimentation, Security, Human Factors, Standardization, Verification.

Keywords
Cyber-security, Red-teaming, Commercial-off-the-shelf software, Government-off-the shelf software

1. INTRODUCTION
With the increased utilization of commercial-off-the-shelf (COTS) and government-off-the-shelf (GOTS) softwares and the expanded use of virtual machines and cloud, a new cyber security issue has arisen. Where once organizations developed almost all software in-house, it is now essential to integrate outside developed software, even when the software may have inherent security issues and vulnerabilities. This software may also threaten the security posture of the internal network and offer a possible attack vector.

With the rapid growth of cybercrime, information security has become a major area of concern. However, the reality for most organizations is that securing information in the modern age has never been more challenging as Abayomi Oloko suggests in an article titled, *Information Security in the Enterprise and Modern Challenges* [1]. "The 21st century has been bombarded with technological innovations aimed at a tech savvy youthful market and communication is getting more open in a world that is truly getting smaller by the day; we are living in a global village." "In the past", writes Oloko, "system security has been seen as a government or a political function but the trend of modern technology in the 21st century and beyond definitely shows that the traditional view of securing information, data, and assets within the enterprise needs a rethink."

There is no mistaking the fact that information security management has evolved beyond government and politics. In fact, information security management has become a critical function for all types of organizations; especially with the global economy, the Internet and the increased fielding of more complex information technology (IT) infrastructures. Various factors have caused the discipline to mature and it has now become one of the core business operations. This shift means there is little room for error when dealing with critical and important data.

The challenge is to integrate COTS hardware and software applications in a way that lowers the risk of threats to confidentiality, integrity, and availability as well as theft or loss. It is important to keep in mind that an information assurance solution can never fully solve the ever evolving cyber-security problems. The TE Framework can be used to address the three core fundamental principles of all security systems: confidentiality, integrity, and availability risks.

2. TE Framework
The TE Framework was designed to assist engineers and developers in strengthening an organization's security posture when integrating COTS software applications. The goal of the TE Framework is to identify the flaws and risk associated with the COTS software and to address them. The TE Framework is broken into two areas: TE Analysis (design assurance, red-teaming, risk assessment, threat analysis and malicious actor definition) and TE Architecture (secure system design, security enclave and container integration and monitoring software). These processes, procedures, and an engineered security container and enclave architecture approach address the risks to confidentiality, integrity and availability of a system and its data.

The TE Framework is tailorable to fit an organization's needs and not all protections may be necessary for every organization. The TE Framework is not intended to be the end all/be all for the security of an organization. It offers a framework that can be used to better protect an organization and not all components are necessary in every situation. The implementation of the TE Framework does however focus on the entire lifecycle of a system: design to decommission.

2.1 TE Analysis
The TE Analysis area is broken into three sections: 1) Brainstorming Sessions 2) Risk Assessment & Design Assurance 3) Red-teaming

2.1.1 Brainstorming Sessions
There are many brainstorming sessions that an organization can use to better understand the scope and effort needed for their

particular situation. All brainstorming sessions may not needed to be utilized for each organization. It is important for these brainstorming sessions to ensure all the stakeholders and responsible parties are in attendance so that a 360-degree view can be achieved. System Administrators, Security Engineers, Hardware Engineers, and many others should be attendance.

2.1.1.1 Questions Brainstorming Session

During the Questions Session, it is important to start by answering some fundamental questions as a group. The types questions to be answered are as follows:

- What are we trying to secure?
- Which key approach should be the focus in how we address the security of this system: security depth or breadth?
- Even though we cannot solve all cyber-risks, are there things that we can design or implement within the system that will reduce the external risks?
- Should the security design apply the same approaches and techniques for all functional areas of the system, or should the strategies and techniques be tailored to the functional area?
- How many different security domains are necessary?

The questions above are not an exhaustive list, but provide a starting point for the brainstorming session. The answers to these types of questions will help an organization to better understand the software/environment they are trying to secure and begin thinking of the areas of concern. They also serve as an input into the TE Architecture.

2.1.1.2 Operational Challenges Brainstorming Session

In addition to the Questions Brainstorming Session, an organization should also conduct an Operational Challenges Brainstorming session. This session will help and organization to better evaluate the possible operational challenges that their system(s) face or may face. By defining these challenges, an organization can better tailor the TE Architecture suggestions for their particular situation.

The Operation Challenges Brainstorming session includes the following types of questions:

- Evaluation of an organization's current or proposed environment for key critical systems, including possibly numerous COTS and custom applications.
- Evaluate the current methods used for managing the systems. The management of these systems can be by using standalone, custom information technology systems, and/or processes that are optimized to support the organizations unique mission/goals.
- Evaluate the data that is stored and the enterprise-information architectures. An organization's enterprise-information architectures and information systems may collect, process, store, and transmit large amounts of data.
- Evaluate the combination of varying applications that may create a complex architecture with many interdependencies. As the number of interdependent applications increases, it becomes exponentially more difficult to secure the system because of the many inflows and outflows.
- Evaluate the amount of current data. An abundance of existing data adds to the complexity of securing the system.
- Evaluate and establish the method(s) used for physical access to hardware and networking devices, etc, as well as users and administrators local and remote access capabilities.

It is difficult to require, enforce, and maintain cyber-security protections for systems after the fact and creates a considerable security risk for the system. The outcome(s) from this TE Operational Challenges brainstorming session should be compiled and used as an input for the development of the TE Architecture.

2.1.2 Malicious Actors Brainstorming Session

Due to the nature of the information an organization wants to protect, many types of malicious actors seek different kinds of information; each with a different intent. Figure 1 below identifies several kinds of possible malicious actors that may seek information stored within an organization. Sometimes they may not be the individuals that carry out the exfiltration, but may be the end user of the information gained through subversion.

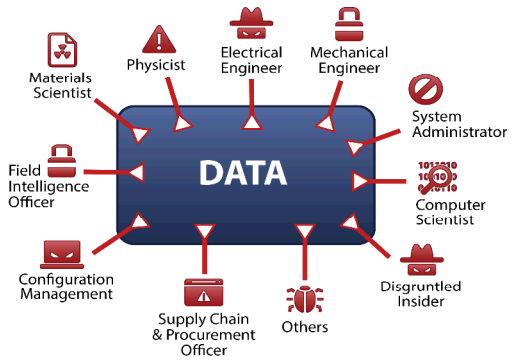

Figure 1. Types of Malicious Actors

It is important to fully understand the risk and possible malicious actors to an organization's system(s) and to mitigate and manage operational security concerns before implementing the TE Architecture.

The suggestions of the TE Architecture below and the information gathered from the other brainstorming sessions can help an organization to better make design decisions and better incorporate security as well as best engineering practices. Design decisions should be made so that an increase in system security posture would be realized. If the impact of an engineering decision reduces the system's security posture, the change should not be incorporated.

2.1.2.1 Defining Trust Brainstorming Session

Trust engineering is a relatively new term in information technology (IT) borne in part out of growing concern about the level-of-trust in modern information systems. Government agencies, who have significant interest in ensuring data confidentiality, integrity, and availability, were some of the first proponents of the approach now called trust engineering.

"In 2002, the National Security Agency's (NSA) Information Assurance Research Group coined the term trust engineering to describe a methodology for making use of software of uncertain

provenance in mission-critical systems."[2] The discipline described in the first paper, *Trust-engineering: An Assurance Strategy for Software based Systems*, no longer seems heretical today, writes Susan Alexander in a follow-on paper, *Trust Engineering – Rejecting the Tyranny of the Weakest Link.* [2]

"The problem with COTs products", says Alexander, "is that organizations have to grapple with untrusted components to get the functionality needed to make the business run smoothly. The outcome is that software has now become "the unwitting delivery mechanism for network attacks" due to a number of factors:

1. Computers have made it possible to do far more interesting things with information than merely communicate it, and the appetite for new functionality has become insatiable.
2. Protection of information is not the killer app for most customers of the newer functionalities, and information technology vendors.
3. A flatter world has produced a dynamic, global IT supply chain offering state-of-the-art functionality much more cheaply than it can be obtained from vetted providers."[2]

These days it is not just government agencies looking to perfect an information technology security strategy that moves away from the idea that the security chain is only as strong as its weakest link. Numerous private companies from banks to electronic lotteries are mobilizing under a variety of national-level directives to protect critical infrastructure and key resources against a broad spectrum of new threats. Given the complexity of modern hardware and software, coupled with user behavior and the possibility of insider threats, providing a more secure system environment has become a priority.

Utilizing trust principles can help an organization to better prepare themselves and protect their data and environments from compromise.

2.1.3 TE Baseline Assessment of Cyber Risk

The Brainstorming Sessions, explained above are essential in beginning to understand the current architecture and the possible areas of concerns. It is also important to complete a full risk assessment and a baseline assessment of cyber risk.

The importance of conducting a cyber-risk baseline assessment of information technologies before they are integrated cannot be emphasized enough. Many IT organizations are required to integrate COTS solutions into their systems with reasons of functionality, design cost, or maturity of product. Yet the question of introduced cyber-risk associated with the COTS solution is seldom asked. In creating a trusted system such as the enclave, the cyber-risk must be understood before the end product and integrated solution is turned on. If a baseline is not conducted, system owners may never be alerted to a successful malicious attack that came from within the COTS solution.

Another key reason to conduct a cyber-risk baseline assessment is to understand the extent of the potential issues and work towards specific mitigations to address them. While many best practice techniques of hardening a system such as applying DoD STIGS generally reduce cyber-risk, they are generally focused on securing the operational environment and may not mitigate malicious insertions into the system through other venues. Conducting a cyber-risk baseline assessment applies the Trust Engineering principles and characterizes the system or COTS solution so that the potential risks are fully understood and mitigations can be engineered to reduce the risk and increase the difficulty for an attacker.

This risk assessment can be completed in whichever way an organization wishes to conduct it. A few ways that this risk assessment can be completed are explained below.

2.1.3.1 TE Design Assurance

The definition and application of design assurance greatly varies between design teams, organizations, companies and even government entities. Regardless, the necessary end goal of design assurance is to follow a well-defined process to increase the security posture throughout the design phase through disposal of the product.

There are many things that design teams or their supporting staff can do to analyze there system: run COTS vulnerability analysis tools, conduct fault tree analysis, conduct code reviews looking for weak validation of function parameters, extensive testing using statistical-based test pattern design (e.g. orthogonal arrays or robust parameter design), and all the other reliable system design practices.

However, while most of these help discretely measure the improvements to system reliability and indirectly reduce the cyber-vulnerability surface, they may constrain a threat model to only the attacks that can be scripted and replayed in an autonomous manner. While this type of extensive testing is necessary, it is important to check and test for specific items that may directly eliminate an attacker's critical path.

It is important to collect data and better understand the inter-workings and security gaps of the existing system as part of the TE Framework. The security efforts also assist with "best practices" for desired protection level, such as which security guidelines to follow like NIST 800-53 or U.S. Department of Defense's (DoD's) Security Technical Implementation Guides (STIGs)[3]. In addition to the risk assessment, and design assurance, identification of malicious actors is critical to the success of the TE Framework. A simplified design assurance characterization and analysis process (iterative in nature) should contain the following: Planning, Data Collection, Characterize, Analyze, Report, and Engage that TE used. Red-teaming can be used to acquire an independent view of the areas of vulnerability and concern for an organization.

2.1.4 TE Security Engineering

Cyber-security elements of projects are crosscutting - the planned work integrates/overlaps with the normal systems engineering and general function use-cases. In order to achieve a better and more secure product, system designers and security engineers needed to accomplish all of the below major tasks:

- Least-privilege security model,
- System architecture assessment,
- Adversary threat model,
- System monitoring and intrusion detection technology assessment,
- Information protection technology assessment,
- Security system requirements definition and meta-model,
- Baseline cyber-risk assessment and design assurance red teaming,
- Trusted architecture design and system mitigations,

- Cyber-security acceptance test plan & tests, and
- Cyber-security maintenance plan.

2.1.4.1 Security Threat Model

A threat model primarily identifies issues that challenge the system's confidentiality, integrity, and availability. This includes determining appropriate threat capabilities from a generic threat matrix, nightmare consequences, attack graphs, strengths, weaknesses, and mitigation strategies.

The threat model will include a number of generic attacks that apply broadly to many networked information systems. Associated with these general system attacks are cyber-security best practices that will help mitigate some vulnerabilities and weaknesses.

However, any application can have many specific threats associated with it due to the sensitivity of the data and customization of the system. These identified specific threats, vulnerabilities, and weaknesses may be from previous studies on the existing system as well as new ones discovered during the baseline assessment and interim testing. The discovered vulnerabilities should be generalized and summarized so system developers can easily apply the appropriate defenses.

2.2 TE Architecture

The information gleaned during TE Analysis is a direct input into the TE Architecture. Within the TE Architecture, a system is broken and defined into a number of functional zones within the system. In each of these functional zones, cyber-protection technologies are implemented. Each functional zone is comprised of nine primary factors that provide a unique view on the security of the system or subsystem:

- System monitoring & analysis,
- Hardware and Network,
- Virtual environments (virtual machines, virtual networks, etc.),
- Software security (including operating systems, applications, middleware, etc.),
- Human Factors,
- Patches and upgrades,
- Configuration management, and
- Post-incident forensics.

These factors should all be considered in order to reduce the overall risks to the system. If one factor is not considered during the design, it leaves a major opportunity for an adversary. The TE Architecture focuses on the key areas of a containered/enclave approach to system security, defense-in-depth, layered defense, and least-privilege.

2.2.1 Containered/Enclave Approach

A pivotal part of the TE Architecture is the implementation of a security container/enclave that follows an object-oriented design methodology using the following design criteria:

- Containerize (physically isolate) COTS system (Hardware, Application, Middleware, Database), critical environments;
- Implement application firewall(s) with advanced monitoring and detection;
- Encapsulate all internal processes, operation, and monitoring from external view;
- Continual and ongoing internal testing of files and software to ensure system integrity;
- Implementation of a least-privilege model - limit container personnel access;
- Tightly control configuration management of entire system (Hardware, Software, and People); and
- Securely manage upgrades and patches to reduce possibility of introducing new vulnerabilities.

The container, as illustrated in Figure 2, "containerizes" core services and functions in a service domain specific in accordance with the ITIL model. This approach allows the container to protect the file integrity, utilize service/user profile monitoring, provide automated security defenses, and offers configuration management and monitoring. A key component to this containered approach and the TE Architecture is the use of defense-in-depth/layered security approach, least privilege, and application specific firewalls and monitoring.

2.2.2 Least Privilege

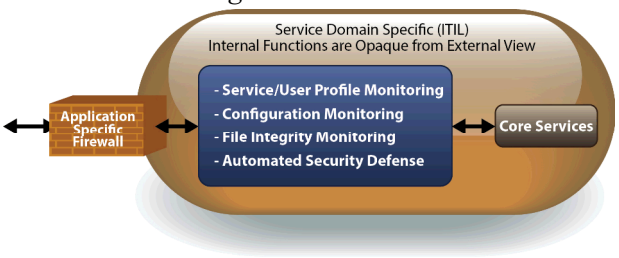

Figure 2. The TE Security Container

In the TE architecture, the least-privilege security model is one that starts at complete lock-down of a system, and incrementally adds access or communication capability to the desired level of functionality, but not more. This requires intricate knowledge of the existing software environment and data communication between services, normal user behavior, network traffic, etc. It also requires network traffic logs and potential instrumentation of the existing production system or client computer to capture the data. This is coupled with large dataset analysis to establish a technically-based model.

2.2.3 Defense-in-Depth

Defense-in-depth is essentially a multi-layered defense approach where the system does not rely solely upon a single cyber-defense mechanism. There should never be a single point of failure for a security system. This strategy is based on the military principle that it is more difficult for an enemy to defeat a complex and multi-layered defense system than to penetrate a single barrier. The TE Architecture implements a Defense-in-depth strategy.

According to a Defense-in-depth paper published by the Information Assurance Solutions Group at the National Security Agency (NSA) this strategy should have the following focus areas: 1) People 2) Technology and 3) Operations.

For the TE Framework, emphasis is placed on the people, technology, and operations areas. Figure 3, depicts these focus areas [4]. If not all areas are addressed, then due diligence has not been performed to ensure a proper strategy. The area of People should be addressed with the creation of policies and procedures, training, physical security, and personnel security.

Figure 3. Defense-in-Depth Strategy

The area of Operations should be addressed through a defined patching process, system security assessment (red-teaming), monitoring system, processes for response to attacks, and backup and recovery.

The TE Architecture is based on a Protect, Detect, and React model. This model is a replica of the Protect, Detect, and React paradigm referenced in the Defense-in-Depth paper published by the Information Assurance Solutions Group at NSA. The areas for defending are referenced in the below Figure 4 [4].

Figure 4. Defense-in-Depth Focus Areas

The Defense-in-Depth philosophy is necessary in order to fully protect the applications and data that reside in the TE enclave. Like traditional Defense-in-Depth models, the outermost layer of the defense philosophy is typically within the current networking layer implemented at an organization, but the TE Architecture takes it one step further. In addition to the traditional networking defenses, TE Architecture is implemented with flexibility and the COTS application in mind. The Defend the Network and Infrastructure and Defend the Enclave Boundary focus areas are detailed below.

2.2.4 Network Defenses

The TE Architecture suggests the use of external perimeter protections to provide application layer and network layer protections against internet-born attacks. These set of protections filter traffic on specific ports entering and leaving the TE Architecture. They also provide user authentication services, perform SSL interception and content filtering, and are a central tap point for an organization's Security Operations Center (SOC), where security analysts are able to monitor threats.

The first lines of defense in any layered defense philosophy is the network defenses and the TE Architecture is no different in this regard. Where the TE Architecture differs from other layered defense philosophies is that the architecture is created with flexibility and application visibility in mind. This is achieved by creating an architecture that contains two key elements:

1. Application traffic visibility in the form of SSL Intercept or Forward-Proxy, and
2. The creation of containers, in the form of zones using networking technologies.

SSL Intercept or Forward-Proxy is the use of a networking device – most commonly a load balancer or firewall – that sits between client-server SSL communications and decrypts that communication. The un-encrypted traffic can then be forwarded to any number of networking tools (application aware firewall, intrusion detection system, intrusion prevention system) where it can be analyzed. The TE Architecture then re-encrypts the traffic by using a second networking device. This is done to minimize the exposure of un-encrypted traffic on a network.

2.2.4.1 Network Enclaves

Now that the application visibility portion of the TE Architecture via SSL interception has been addressed the team can examine the zone-based philosophy behind the architecture. The TE Architecture suggests the creation of at least three network zones:

1. Tenant Zone
2. Infrastructure Zone
3. Monitoring Zone

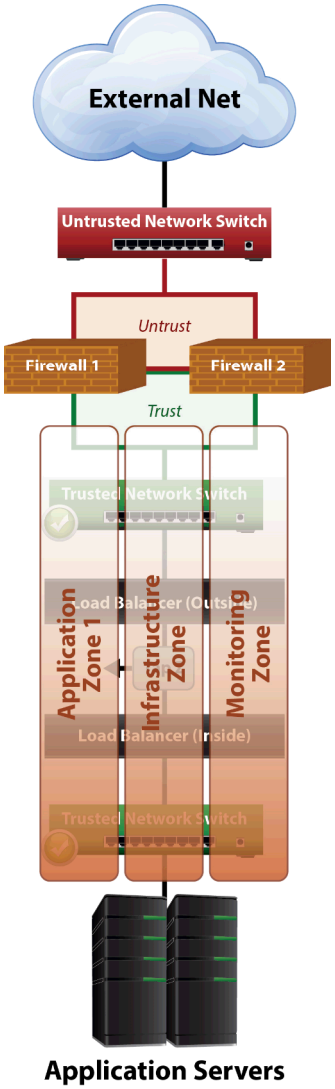

Figure 5. TE Network Architecture with Zones

Figure 5, above displays these zones in a possible scenario. The Tenant Zone should contain all of the servers/services that are required for the application being protected. The Infrastructure Zone should contain all of the shared elements that are required by all applications to run such as time services, shared databases, and authentication servers. Lastly, the Monitoring Zone should contain any of the tools that an organization may use for cyber-monitoring such as security information and event manager (SIEM), Log Aggregation, and IDS tools. The overall architecture provides the implementer flexibility in how the zones are constructed.

The three suggested zones can be implemented in any number of ways: VLANs on the Internal Trust Network Switch connected to the application servers, VLANs created on the Inside Load Balancer or, standard firewall zoning or virtual routing and forwarding (VRF) technologies.

The trend in computer networking is to move away from more traditional, purpose built hardware and toward virtualization of networking elements and functions. Network Functions Virtualization (NFV) is the term that is used to describe this movement and is essentially the transitioning of these network functions, defined in hardware, (firewalling, network address translation (NAT), IDS, load balancing, etc.) to software. In the first step towards NFV, many vendors have already created virtual appliances (load balancers, firewalls, and WAN optimization devices) that can be implemented in place of any network hardware appliance within the TE Architecture.

2.2.5 Monitoring

After the networking layer, there is a monitoring/application layer. This layer is comprised of many tools used to monitor the network traffic, encrypted and un-encrypted, in order to get a complete view of the activity within the TE Framework. Monitoring software was created as part of the TE Project at SNL that uses multiple data sources to draw an in-depth look and analysis at user and software behavior in order to form a profile. This monitoring software combined with hardened operating systems and zone limitations add to the defense-in-depth protection of the enclave. In addition to the monitoring software, the TE Framework can easily integrate with the tools currently used by an enterprise to monitor their network.

2.2.6 Operations

In addition to technical design decisions and monitoring, there are additional factors that must be taken into account when a system is actually put into operation. When a system reaches operational status, many things should be accounted for during operations in addition to an organization's current security policies and procedures. Two areas noted in the TE Framework are the addition of Policies and Procedures surrounding the security of the container and the implementation of a Trusted Software Process. These are not an exhaustive list of areas of improvement of operations, but merely a subset of what can and should be implemented. Both are explained in further detail below.

2.2.6.1 Procedures and Policies

Adding to the monitoring/application layer requires the use of operations policies put in place to further limit access to the enclave and better protect the data within. Operations are most definitely an important area that should be focused on as a possible avenue for a breach. The TE Architecture suggests the implementation of policies and procedures for limiting access to the machines within the defined enclaves.

This policy dictates that all access of an administrative nature to the enclave will be through one central point. This central access point shall be implemented in such a way that users have limited access to only the machines they need access to rather than provide open access. Additionally, different types of administrative users, such as application administrators and system administrators should be created. Each of these administrators should have unique rights and access depending on their role and need across the enclave.

2.2.6.2 Trusted Software Process

Another area of operations that should be accounted for is the process for the introduction of new files to the containers/enclaves. Typically one of the ways a user introduces malware to a system is through a download to their workstation that has not gone through a thorough scanning process. For the TE Architecture, the only point of transfer for files into the enclave should be through a trusted software process.

In the TE Architecture, tools should be used to scan new patches and necessary files for introduction in the enclave for malware. Approved and scanned files will then be transferred via sneaker net into the enclave and added to a central repository. Removing the ability of administrators of applications within the enclave to place harmful files into the enclave further protects the security posture of the enclave.

3. Why the TE Framework?

The TE Framework helps an organization to develop a new system or re-architect an existing system. It can be used to help an organization define the risks and areas of concern through the suggestions of the TE Analysis portions, which in turn helps the organization to define the necessary number of functional zones within a system where cyber-protection technologies should be implemented. The TE Framework is a Framework and is tailorable to an organization and has been utilized at SNL.

4. Conclusion

These days it is not just government agencies looking to develop an information technology security strategy focusing less with the concept you are only as strong as your weakest link. Numerous private companies from banks to electronic lotteries are mobilizing under a variety of national-level directives to protect critical infrastructure and key resources against a broad spectrum of new threats. Given the complexity of modern hardware and software, coupled with user behavior and the possibility of insider threats, providing a more secure system environment has become a priority. The TE Framework provides the building blocks for an organization to better protect themselves and their data from exposure and loss.

5. ACKNOWLEDGMENTS

6. REFERENCES

[1] Oloko, Abayomi, *Information Security in the Enterprise and Modern Challenges.* Dataversity, December 15, 2011

[2] Alexander, Susan D., *Rejecting the Tyranny of the Weakest Link*. U.S. Intelligence Advanced Research Projects Activity, ACSAC '12, December 3-7, 2012, Orlando, FL

[3] http://iase.disa.mil/stigs/Pages/index.aspx

[4] Author Unknown, Defense in Depth, A Practical Strategy for Achieving Information Assurance in Today's Highly

Networked Environments, National Security Agency, Ft. Meade, MD