INL/EXT-14-31534

ATAC Process Proof of Concept Final Report

Bri Rolston
Sarah Freeman

March 2014

The INL is a U.S. Department of Energy National Laboratory
operated by Battelle Energy Alliance

INL/EXT-14-31534

ATAC Process Proof of Concept
Final Report

Bri Rolston
Sarah Freeman

March 2014

Idaho National Laboratory
Idaho Falls, Idaho 83415

http://www.inl.gov

Prepared for the
U.S. Department of Energy
Office of Electricity Delivery and Energy Reliability
Under DOE Idaho Operations Office
Contract DE-AC07-05ID14517

Abstract

Researchers at INL with funding from the Department of Energy's Office of Electricity Delivery and Energy Reliability (DOE-OE) evaluated a novel approach for near real-time consumption of threat intelligence. Demonstration testing in an industry environment supported the development of this new process to assist the electric sector in securing their critical networks. This report provides the reader with an understanding of the methods used during this proof of concept project. The processes and templates were further advanced with an industry partner during an onsite assessment. This report concludes with lessons learned and a roadmap for final development of these materials for use by industry.

Executive Summary

The Attack Technology, Analysis, and Characterization (ATAC) process is a data organization and analysis framework that provides companies and practitioners with a methodical approach to facilitate the efficient aggregation and evaluation of cyber-risk related information. ATAC is a proof of concept Frontier project funded by the Office of Electricity Delivery and Energy Reliability in the Department of Energy (DOE-OE). The team at Idaho National Laboratory (INL) within the National and Homeland Security Critical Infrastructure Protection organization performed the work in partnership with an asset owner and refinement by DOE-OE experts.

Energy sector industries face an ever changing cyber threat environment with a constant flood of emerging threat information that is not specific to their implementation. Past threat sharing events with these asset owners in the energy sector has shown that there is no single approach which can be applied across the board. Mitigations for vulnerabilities and exploits along with incident response are static while the threat is dynamic due to the changing techniques of the attackers. Understanding and mitigating these threats with process improvement adds value to the business process of industry.

In the event of a security breach, ATAC can be employed to qualify the severity of the infiltration and to formulate the appropriate response. Alternatively, organizations can use ATAC in order to generate threat characterizations from open source research and technical cyber security information. Industry can then apply this threat characterization (*ATAC*) against their Response Analysis and Characterization Tool (ReACT) profile in order to develop a current perspective of their risk.

The ATAC process can be subdivided into four main areas of effort including threat identification, simple threat analysis, complex threat analysis, and the development of a strategy for process improvements after incident response. Products of the ATAC process include an Attack Surface Analysis (ASA) to understand the gaps in security of the targeted system; an Attack Timeline and Attack Path Model to understand the movements of an adversary within a system or network; Predictive Attack Path Analysis (PAPA) and an Attack Technology Forecast in order to ascertain an organization's risk; and, a ReACT Response Plan for threat incident response feeding into process improvements.

The industry partner's initial impression was positive to the ATAC proof of concept. The industry partner found value in the systematic process for threat assessment, which eliminated much of the subjective nature of analysis to which they were accustomed. Additionally, the industry partner praised the operational focus of the research which produced information that could be used to improve the organization's current defense strategies. During the assessment with the industry partner, INL collected valuable feedback for both processes which appears in the results and discussion section of this document. The ATAC process promotes a collaborative effort, not only with external subject matter experts (SMEs) and industry, but also among key groups within an organization. The most successful implementation of ATAC will involve further development of the process in partnership with industry.

Recommendations for future development were identified through asset owner partnerships. These recommendations include testing the process through additional case studies, developing additional training materials on how to best deploy ATAC within an organization, creating an attack path catalog and attack process encyclopedia, automating the process, and developing a framework to provide near real-time threat sharing.

Acknowledgement

DOE-OE funded this Frontier project based on a concept presented during the FY13 annual operating plan teleconference. The Attack Technology, Analysis and Characterization (ATAC) process was funded along with other synergistic Frontier concepts. Working with DOE-OE experts, INL staff developed a process which provides organizations with a systematic approach for the consumption and application of threat information. The ATAC process represents this cooperation between DOE-OE subject matter experts, asset owner partnerships and INL staff.

CONTENTS

Acronyms

ATAC	Attack Technology, Analysis and Characterization
ASA	Attack Surface Analysis
C&C	Command and Control
CIP	Critical Infrastructure Protection
CTA	Complex Threat Analysis
DOE-OE	Department of Energy Office of Electricity Delivery and Energy Reliability
EMS	Energy Management System
FSL	Functional Security Layers
FSM	Functional Security Matrix
NHS	National and Homeland Security
ICS-CERT	Industrial Control Systems –Cyber Emergency Response Team
IDRA	Impact Driven Risk Analysis
INL	Idaho National Laboratory
PAPA	Predictive Attack Path Analysis
PoE	Point of Entry
ReACT	Response Analysis and Characterization Tool
SCADA	Supervisory Control and Data Acquisition
SME	Subject Matter Expert
STA	Simple Threat Analysis
UR&R	User Roles and Responsibilities

1. Introduction

The Attack Technology, Analysis, and Characterization (ATAC) process provides organizations with a systematic approach for the consumption and application of threat information. A team of Idaho National Laboratory (INL) experts within the National and Homeland Security (NHS) Critical Infrastructure Protection (CIP) organization performed this work focused on cyber security risks to control systems. The process is designed to understand cyber security architecture and risks from the user's perspective, not the attacker's. This ATAC Proof of Concept Final Report fulfills the milestone *Final Operational Security Baseline Template and Final Report* (*ATAC Final Report)*, number 2.6.8. The ATAC process builds on the concepts developed in the Response Analysis and Characterization Tool (ReACT) Proof of Concept development project.

1.1 Background

Since Department of Energy Office of Electricity Delivery and Energy Reliability's (DOE-OE's) funding of the National Supervisory Control and Data Acquisition (SCADA) Test Bed in 2003, INL's work in CIP has focused on the cyber security of control systems supporting the energy sector. INL's direction continues to support work with vendors of control systems and asset owners who manage and operate controls for the energy sector. Based on these past relationships, INL understands the challenges associated with the consumption of threat information and its applicability to industry. Based on significant experience in evaluating cyber security vulnerabilities of control systems, exploits, and incident response, INL was able to identify the issue of addressing the static, one time, vulnerability in a dynamic threat environment defined by an attacker's changing tools, techniques, and procedures. The ATAC proof of concept project demonstrates a standard method for gathering and analyzing threat information. ATAC organizes this information in an accessible format, which facilitates threat intelligence consumption and the development of new defense strategies to lessen risk.

1.2 Scope

The scope of the ATAC proof of concept project includes developing a methodology for a consistent approach for understanding cyber threat information. In particular, this project sought to codify the common indicators of threat. Using these indicators, this project developed an analytical framework through which industry could assess and respond to cyber threats. This analytical framework has four main areas of effort including a) the identification of an emerging threat, b) simple threat analysis c) complex threat analysis and d) the identification of a response strategy to include process improvements. The development of ATAC was accomplished through collaboration with asset owners. This partnership allowed for a test of the proof of concept.

1.2.1 Basis for Development

Based on substantial threat-focused analytical experience, INL researchers examined the concepts and assumptions which dictate industry's current approach to threat identification, analysis, and management. Collaboration with industry partners identified the following issues:

1. Changes in the threat environment have fundamentally shifted the requirements of analysis. The current sophistication of cyber threats requires that industry take an active role in recognizing, interpreting, and responding to threat information.

2. Traditional approaches result in traditional results. That is, traditional perspectives on threat intelligence, developed in the national security sphere, fail to meet the needs of industry.[1] Since the most common threat information does not include how technical threat is relevant to the energy sector, each organization must develop their own methods to assess threat and its effect on sector-specific resources.
3. Organizations face challenges related to the effective consumption of threat information, not access. Today both public and nonpublic information is shared with industry. However, in many cases, this information is not presented in a format which is actionable for industry.
4. Industry security teams remain focused on short-term technology issues. Many lack the approaches necessary for long-term trend analysis.

Working in tandem with industry, INL researchers developed a process to provide the foundational analysis, information architecture, and relationship correlation necessary to address the obstacles previously identified.

1.3 Purpose

INL designed the ATAC process from the bottom-up with the intention of providing energy-sector members with a methodical and specialized approach for classification and mitigation of cyber threats. This approach differs from traditional perspectives of threat analysis by focusing on a threat's operational impact to industry. This information can then serve as the basis for developing a cyber security metrics program for critical infrastructure. Once cyber risk has been characterized, organizations can begin to address threat more proactively.

Industry, much like other members in the public and private sectors, face an ever changing cyber threat environment. However the security efforts of these organizations do not suffer from a dearth of threat data, but rather the constant flood of emerging information. Against this backdrop, industry partners need a consistent, organized approach for evaluating publically available threat information.

The challenge of threat intelligence consumption is compounded by variations in focus, methods, and goals, across sectors and industries. There is no single approach which can be applied across the board. However, personalized approaches will fail if these efforts ignore the business goals of an organization.

1.3.1 ATAC Proof of Concept Objective

INL sought to develop a methodology for a consistent approach for understanding cyber threat information. In particular, this project sought to codify the common indicators of threat and incorporate them within an analytical framework. Using these indicators, this project developed an analytical framework through which industry can assess and respond to cyber threats. Throughout the course of this work, INL researchers approached threat from a non-traditional perspective; that is, the ATAC process was designed from the ground up to focus on an operational perspective. For example, this work remains primarily concerned with how the energy sector is likely to be affected and what technical approaches are necessary for defense. These concepts make the threat assessment process more approachable for individuals who lack substantial analytical experience evaluating threat.

The evolution of the ATAC process, much like ReACT, expanded from the needs of the end users. Central to both projects was significant collaboration with industry, along with the realization of the

[1] Threat is a concept traditionally used in the national security, military, and federal government arenas to describe an individual actor or group of actors with the capabilities, opportunity, and intent to enact a range of negative actions.

projects through an onsite assessment with an industry partner. These efforts laid the foundations for this report and assisted with the identification of needs going forward.

1.3.2 Objective of ATAC

The ATAC process is concerned with the development of an organizational security posture based on adversarial capabilities. However, this methodology is not concerned with adversaries, but rather the likelihood that an attack will be successful. The ATAC process is intended to answer the question of how a change in attack technology affects an organization's cyber risk. This will allow security professionals to properly translate threat alerts into consumable intelligence for their individual business environments.

2. ATAC Methodology Summary

ATAC, much like ReACT, is a data organization and analysis framework that facilitates efficient aggregation and evaluation of cyber-risk related information. However, unlike the ReACT methodology, which approaches risk from an organizational perspective, ATAC forces analysts to consider the risk to an organization from an adversary's perspective. Essentially, ATAC views organizational security from the outside looking in. This provides organizations with a means to evaluate their own cyber risk. The ATAC process increases an analyst's situational awareness and leads to a greater understanding of the threat environment.

The ATAC process can be divided into four sections, threat identification, simple threat analysis, complex threat analysis, and the resulting ReACT Response Plan. In its simplest form, the ATAC process progresses from the identification of a potential threat, to information gathering *(Simple Threat Analysis)*, and finally to the development of improved security strategies.

1. ***Threat Identification*—**the initial identification of a cyber threat, originating from within the organization or externally (via threat reports, academic research, news articles, listservs, etc.);
2. ***Simple Threat Analysis (STA)*—**the collection and analysis of information based on an incident within or outside of the organization. STA is based on attack surface analysis (ASA), attack timeline development and attack path modeling;
3. ***Complex Threat Analysis (CTA)*—**based on an adversary's capabilities and historical preference, CTA identifies the most likely attack paths and technology that will be used to target an organization.
4. ***ReACT Response Plan*—** the creation of a response strategy to address issues identified during the ATAC process, namely based on identified likely attack paths.

The ATAC process alters the ways in which organizations face threat. Rather than focus on developing a myriad of specialized security plans, the ATAC method focuses on the *process* of an attack. The ATAC process is not concerned with the identification of nameless faceless adversaries, but with industrial operations protection. It provides security professionals with the means to secure their systems even with poor or incomplete information, and without any specific knowledge of the existence of an attack.

2.1 Moving from ReACT to ATAC

Through the course of industry collaboration, INL researchers became aware of the need to reorder the ATAC and ReACT methods. By moving from ReACT to ATAC rather than the alternative, organizations experience the following benefits:

- **Increases the relevancy of an analytical product:** Specifically, adoption of the ATAC process increases the likelihood that in-house analysis can be done in a timely manner.
- **Develops an organization's analytical foundation:** After the emergence of a potential cyber threat, in-house teams can immediately begin to process what risk this cyber threat poses to the operational capabilities of an organization. This allows groups to produce higher quality products even with limited capabilities.
- **Promotes predictive rather than reactive responses:** Similar to above, the ATAC process is most effective when implemented following a ReACT assessment. This ordering allows organizations to jump start incident response by amassing key information about their organizations before the introduction of an issue.

3. Assumptions

1. The ATAC development team worked with the DOE-OE experts on understanding the assumptions, and products prior to finalizing the methods developed. Asset owner partnerships validated these assumptions. Unlike in the national security sphere, industry lacks a common approach for sharing information and responding to technical cyber threats.
2. Current threat information is shared with industry in an unwieldy format. That is, threat analysts have not focused on what threat information is and what format would assist industry in their defensive efforts.
3. In order to calculate risk posed by a specific threat, groups must have an understanding of the potential operational impact of an event.
4. Working within the reality of limited time and resources, industry needs a means to prioritize threat information and defensive efforts, allowing justification within the business sphere.

4. Products of the ATAC Process

Standard methods and steps have been defined for the four subsections of the ATAC process identified above (Threat Identification, Simple Threat Analysis, Complex Threat Analysis, and ReACT Response plan). An overview of the critical workflow components is depicted in Figure 1. A detailed description of the ATAC process can be found in Appendix I.

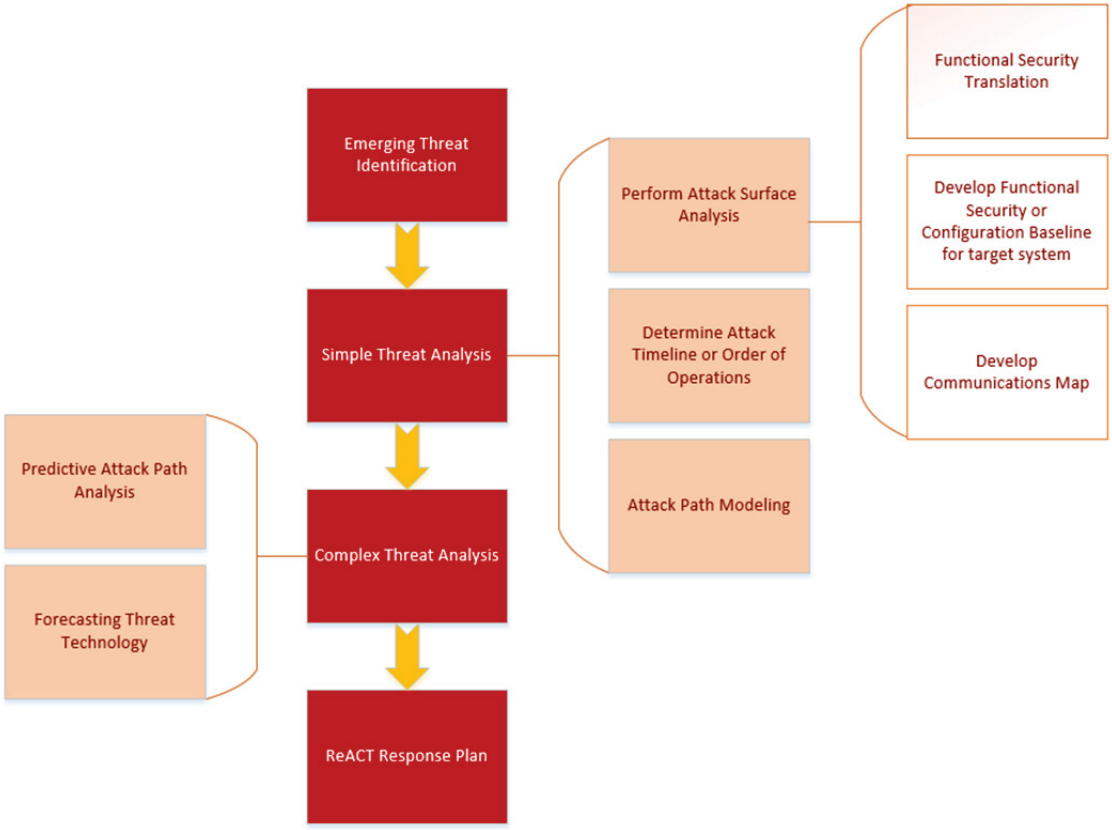

Figure 1: An overview of the ATAC process.

In general, the ATAC process collects information regarding cyber threats. However, how ATAC is deployed will dictate the type of data gathered and the products of the process. Typically, ATAC will be used either following an incident or preemptively. For clarity, the products have been separated in Figure 2, according to these two use cases:

ATAC following an incidence (STA)	ATAC used preemptively (CTA)
• Attack Surface Analysis • Attack Timeline • Attack Path Model • Predictive Attack Path Analysis *(optional)*	• Predictive Attack Path Analysis • Forecasting Attack Technology

Figure 2: The products of the ATAC Process. The left column includes products of Simple Threat Analysis (STA), while the right contains the products of Complex Threat Analysis (CTA).

4.1 Attack Surface Analysis (ASA)

During a ReACT assessment, ASA identifies which areas of an organization's core system should be tested for vulnerabilities and which require additional protection. This information is also relevant during

the ATAC process, although in this case the focus remains on the target system information. There are three components necessary to complete ASA during the ATAC process:

- Functional Security Translation
- Development of the Functional Security or Configuration Baseline for the target system
- Development of the Communications Map

These components are completed using the information available in publically available advisories and open source intelligence reports or information gathered following a security breach. Again, while this information can be presented in the ATAC Functional Security Matrix (FSM), that is not a requirement.

4.2 Attack Timeline

The attack timeline is chronological listing of the various events of the attack. It may include information about how an adversary moved within a compromised system or network, the delivery date of the payload, and any changes made to the system or network. Similar to ASA, the timeline is developed from publically available information such as advisories and open source intelligence reports or information gathered following a security breach. Timeline information may be represented on a standard chronology or the ATAC Lifecycle.[2]

As noted in Heuer and Pherson's seminal work, *Structured Analytical Techniques for Intelligence Analysis*, timelines aid in the identification and correlation of events. "These techniques also allow you to relate seemingly disconnected events to the big picture to highlight or to identify significant changes or to assist in the discovery of trends, developing issues, or anomalies."[3] When considering the order of events for a historical incident, details will often be muddy at best. The timeline helps the analyst solidify what is known versus unknown.

4.3 Attack Path Model

In the ATAC process, the attack path model is closely related to the attack timeline. While analysts use the attack timeline in order to determine ***when*** an adversary moved in a system, the attack path model is concerned with the ***how*** of that traversal. The attack path model alters the perspective of security professionals, thereby promoting the development of more effective mitigation strategies. INL researchers have included space within the ATAC FSM for attack path modeling. However, this is just a suggested display for data.

4.4 Predictive Attack Path Analysis (PAPA)

Predictive attack path analysis (PAPA) , which marks the beginning of complex threat analysis in the ATAC process, shares many characteristics with attack path modeling. Both focuses on the methods used by an adversary to transverse a system, however, PAPA goes beyond attack path modeling by using limited data to extrapolate these movements. Additionally, PAPA can be used either following an incident or preemptively by an organization to identify how future attacks are likely to manifest themselves. It is not necessary that an analyst organize this information in the FSM. Analysts should use whatever methods are most appropriate for them.

4.5 Forecasting Attack Technology Report

The identification of emerging attack technology is by far the most complex component of the ATAC process. It requires that the in-house analyst or external SME not only have a complex understanding of

---

[2] The ATAC Lifecycle is discussed in greater detail in Appendix II of this document.

[3] Richards J. Heuer Jr and Randolph H. Pherson, *Structured Analytic Techniques for Intelligence Analysis*, (Washington, D.C.: CQ Press, 2011), 52-54.

an organization's core systems, but also substantial information and understanding of the factors which will influence the future development of that technology (i.e. cost, popularity of product, market competition, personal interest, etc.).

The simplest and most accessible form of forecasting is basic extrapolation.[4] Analysts collect open source information related to the attack technology and incident including prevalence of similar events, the prevalence of attack technologies, style of attacks, information regarding initial breaches, and data on the adversary's supporting infrastructure (command and control channels, servers, purchased and rented services, etc.). This information is then used to estimate future trends in attack technology. It is not necessary that an analyst organize this information in the FSM. Analysts should use whatever methods are most comfortable or convenient for them.

4.6 ReACT Response Plan

The final product of the ATAC process (as well as the ReACT Methodology) is the ReACT Response Plan. (This can be completed after simple or complex threat analysis.) This piece is intended to address specific issues identified throughout the course of ATAC. The ReACT Response Plan assists organizations with the development of policies that will prevent security breaches going forward.

5. Results and Discussion - ATAC Case Study

As part of the ATAC proof of concept project, the INL team conducted an onsite assessment with an industry partner. This case study was intended to prove value added for industry going forward. In particular, the industry partner was highly appreciative of INL's efforts to develop a process which would allow for improved use of existing threat information. This case study also assisted with the development of a more complete ATAC process. Finally, the feedback from the industry partner helped to shape the next steps for the ATAC project.

5.1 Summary of Case Study at Industry Partner Site

During the course of the proof of concept projects, INL researchers conducted an onsite assessment with a large energy utility company. The primary purpose of this visit was to further develop the ATAC process as well as determine the feasibility of moving from ReACT to ATAC as previously discussed. Additionally, the onsite meeting was used to gather business sensitive information the industry partner did not feel comfortable sharing via alternative means.

In many instances, the ATAC process will address information gathered during a ReACT assessment. This was also the case during the industry case study, when key information was gathered during the onsite assessment for later ATAC analysis. Onsite work was conducted over three and a half business days. However, the timeframe for the ATAC process is highly dependent on a number of different factors including an organization's threat assessment capabilities, an organization's existing security posture, the complexity of the core system being analyzed, and an analyst's current knowledge of the core system. Wherever possible this summary includes time estimations for future assessments.

During this onsite, INL researchers conducted several meetings with parties familiar with both the business goals and the functional requirements of the core system selected (the Energy Management

[4] As noted by Robert Clark, authors tend to develop their own vocabulary when discussing trend analysis. Clark distinguishes three predictive techniques, extrapolation, projection, and forecasting. [Robert M. Clark, *Intelligence Analysis: A Target-Centric Approach* (2nd edition), (Washington, D.C.: CQPress, 2007), 195-213.]

System (EMS)).[5] Industry partner participants included representatives from the following groups: IT management, EMS management, network operations, and the security operations center.

INL researchers conducted initial meetings in order to introduce the ReACT and ATAC process to the industry partner. These meetings were also used to determine the time commitments of key individuals during the week. Additionally, INL researchers used this time to direct a group discussion to define critical impact based on the organization's business and mission goals. During the case study, these meetings were relatively brief (1-2 hrs.). However, these meetings become less critical when conducting in-house assessments, depending on an organization internal understanding of the business goals.

Several meetings conducted in the first two days focused on developing a detailed understanding of the functionality and operations of the EMS necessary for the ReACT assessment. In essence, this work focused on gathering the necessary information to develop the functional baseline. These meetings were orchestrated by the industry partner for the benefit of INL researchers (2-3 hrs.). However, when conducting an in-house assessment, these meetings may become unnecessary. In many cases, an analyst may have sufficient knowledge of the system or be able to access internal documentation to develop the functional baseline independently. While these meetings can be limited, it is advisable that an analyst still budget 2-3 hours for the development of the functional baseline. This is because the ReACT assessment is most successful as part of a collaborative effort to identify technology needs and functionality.

During the onsite assessment, one full day was devoted to developing the communications maps and gathering existing security posture information (~8 hrs.). The collection of communication and security posture information provides asset owners with the data necessary to perform ASA. (INL researchers also used time during the meetings to develop a list of critical components, services, and data for use during the ATAC evaluation.) The length of these meetings was primarily dictated by the complexity of the EMS. An in-house analyst more familiar with this system may have been able to limit this component to a half-day (~4 hrs.), but some additional meeting requirements are likely. This is due not only to the complexity of the system, but also the dependent relationships between the EMS and other management systems. It is unlikely that an analyst would independently have access to all the communication and security information necessary to complete this section of the assessment.

Following the completion of the onsite, INL researchers conducted attack surface analysis, root cause failure analysis, and cluster analysis on the EMS (the final components of the ReACT assessment). This part of the process took approximately 4 hours, but it is likely the most difficult to be completed by industry. In particular, attack surface analysis is the most intensive, and is based on experience (for the advanced user) or research (for the less experienced user). Likely time requirements for this process are six to twelve hours, but more in depth analysis may take up to forty hours. ASA, RCFA, and cluster analysis provide asset owners with key system information which feeds a ReACT Response Plan. Additionally, the information gathered during a ReACT assessment is used as part of the ATAC process.

Also after the completion of the onsite, INL researchers completed the Complex Threat Analysis (CTA) for the industry partner.[6] This part of the process took approximately four hours. This is the section of the ATAC processes that is the most likely to benefit from SME involvement, not because of the time differential, but because of the level of difficulty associated with the analysis. An estimation for the minimum time needed would be 40 hours.

[5] In this case the core system was selected prior to onsite meeting. Through teleconferences and discussions, the industry partner-led group selected the Energy Management System (EMS).

[6] INL researchers had originally included the CTA of ATAC as part of the work scope for the onsite meeting. Unfortunately due to scheduling conflicts with the in-house threat analysts, this portion of the case study had to be rescheduled.

5.2 Feedback Summary

Work with industry partners provided substantial feedback:

- Development of detailed and extensive instructions would assist with the training of onsite teams by industry.
- Real-time examples of the ATAC process seem to be the most successful when training new users.

5.3 Lessons Learned

During the course of this proof of concept project, INL researchers identified the following lessons learned:

- Work with industry suggests that ATAC process will be more successful after the completion of a ReACT assessment. While the ATAC process can still be used independently from ReACT, the existing capabilities of most organizations support a relatively fluid transition from ReACT to ATAC.
- The industry partner for the onsite assessment maintains a highly sophisticated security posture, and is exceptionally capable in their information gathering methods. Given this reality, additional case studies would be beneficial to baseline an organization with fewer resources and in order to further develop the ATAC process.
- Complex threat analysis of ATAC remains the most difficult component for industry to implement. While even experienced threat analysts face difficulties when developing predictions or forecasts, additional case studies may yield new methods for this process. Future efforts should focus on the development of indicators which can be adopted by industry in order to improve these efforts.

5.4 Resources Required

When moving from ReACT to ATAC, an industry-based analyst will first complete a ReACT assessment in order to begin ATAC. This in turn will require substantial information regarding the core system.[7] Since the first steps of the ReACT process rely on the gathering of this information, the only preparation necessary is the identification of where this information can be found. Unfortunately, the needs of the analyst will vary from organization to organization. For example, some groups may keep a central database with detailed information regarding the core system. In this case, the analyst will be able to amass substantial data independently. However, if the core system is exceptionally complex (or if such a database does not exist) then an analyst will have to conduct interviews with various teams. There is no easy response to answer the question of who needs to be involved, as this answer will also vary. However, likely participants include members of the core system management, IT management, network operations, and the security operations center.

The ATAC process is best completed by a trained threat analyst. However, if organizations do not have the option of hiring a threat analyst full time, much of ATAC (especially STA) can be completed by any individual with knowledge of the core system. In fact, key portions of the ATAC process (for example the development of a ReACT Response Plan) will be more successful with an in-house analyst rather than an external SME.

[7] Examples of the information necessary include applications and services on the core system, key data stores for work, any connections to the system (neighboring databases, subsystems, etc.), and ingress and egress traffic (inter- and intranet).

CTA currently represents the most challenging and least defined portion of the ATAC process. With substantial analytical experience, threat analysts (SMEs) are the most likely to excel here. However, an SME acting as a consultant will fail to be able to use the ATAC process to its full conception. This is because SMEs will never be able to match the potential knowledge base and comprehension of system interaction of an in-house analyst. Long-term development of industry's analytical capabilities represents the best option for successful implementation of the complete ATAC process.

5.5 Potential Applications

Regardless of the sophistication of the user group, the ATAC process provides the foundations for industry-relevant threat analysis. The inexperienced user should be able to conduct at least STA based on the introduction of new cyber threat information. This in turn will assist with an organization's ability to react and lower their organizational risk. For the more sophisticated groups, the ATAC process can be adopted for use with the trend analysis associated with PAPA and forecasting of attack technology. Successful application of CTA will provide organizations with increased situational awareness of the threat environment.

6. Going Forward

6.1 Next Steps

The primary goal of the proof of concept project was to prove ATAC's feasibility and benefit to industry. The complexities of this project dictate that further development is a requirement prior to wide adoption. In particular, INL and asset owners identified several areas for further development which would improve the value of ATAC:

- **Conduct additional case studies:** As mentioned in the lessons learned section of this document, additional case studies are necessary in order to better define the ATAC process. This remains INL's primary recommendation moving forward. In particular, additional case studies should be carried out among industry partners with differing levels of cyber security capabilities. Examples of recommended case studies would involve small, medium, and large utilities.
- **Development of detailed training materials**: The industry partner expressed an interest in the creation of additional training materials. The ATAC process provides analytical foundations for users with limited experience. However, currently the adoption of the ATAC process is difficult without the assistance of an SME. The best way to train individuals is likely to be through detailed and expansive workbooks. These pieces should not only provide the conceptual overview of the process, but also include several examples of the breakdown of threat information with the ATAC process.
- **Development of an attack path catalog and attack process encyclopedia**: When working with ATAC, SMEs bring experienced-based knowledge. The most effective way to match that knowledge base is through the collection of common or historical attack paths.
- **Creation of an automated process**: The industry partner expressed an interest in the creation of an automated process which would assist with the analysis. The development of an automated process for use with the ATAC process would be more challenging than its ReACT counterpart. Additionally, due to the changing face of threat, automation may have short term benefits. Still, a program which walks users through the steps might decrease the amount of time needed for ATAC.

- **Real-time threat sharing**: Industry partners identified the development of a real-time alert system using the ReACT and ATAC processes as an area of interest. While such an undertaking is currently outside the scope of this project, additional consideration should go into potential information sharing mechanisms.

7. Conclusions

INL has succeeded in developing a process which benefits asset owners through the development of an organized method for the evaluation of threat intelligence. By focusing on the operational impact of threat, organizations are better equipped to develop effective defense strategies, even facing situations with limited or imperfect information.

The ATAC proof of concept project sought to develop a process for assessing emerging cyber threats for use by industry. Central to this goal, was the development of methods which addressed the concerns of industry. In order to ensure the value added of this process, ATAC was reviewed by several industry partners and vetted during an onsite assessment. Following the onsite assessment, ATAC received substantial positive feedback regarding its potential usefulness for industry members. Through collaboration with industry partners, INL researchers identified several areas for further development prior to wide industry adoption. INL has developed suggestions for the next steps for ATAC deployment.

Appendix I – ATAC Process

This section describes the current ATAC process developed as a result of the proof of concept project. As depicted in the following diagram, there are four basic steps or sections for the ATAC process, threat identification, simple threat analysis (STA), complex threat analysis (CTA), and the ReACT Response Plan. These steps are further dissected below.

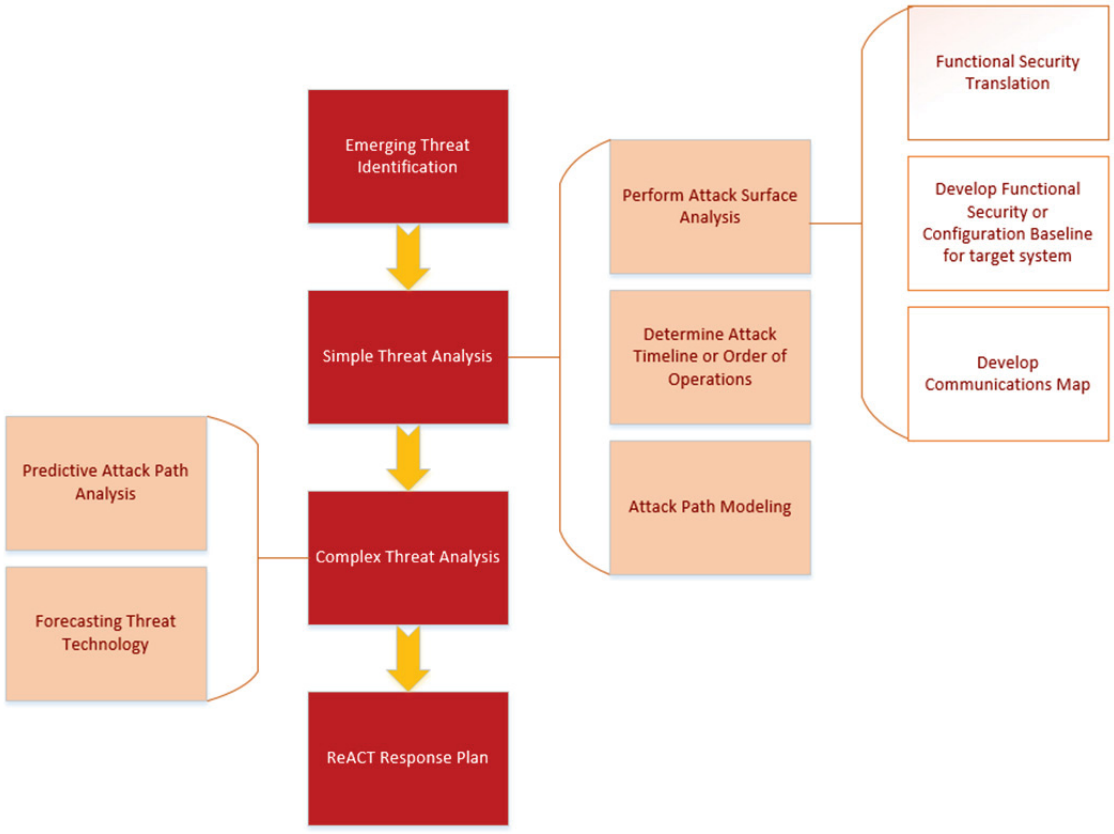

Figure 1. The ATAC Process.

Emerging Threat Identification

The identification of a cyber threat marks the beginning of the ATAC process. Threats can be identified in a number of ways; security teams may locate something on their internal networks, or perhaps more likely, they may receive an alert from outside of the organization. Regardless of the method of identification, security professionals can use the ATAC process to solidify their own security posture and harden their systems.

Step 1. Identify the Type of Threat Intelligence Being Evaluated

The type of threat intelligence being evaluated must be identified in order to understand if the technical threat(s) described affect the organization's cyber security risk. Categorizing the type of technical threat intelligence being analyzed provides contextual understanding of how the technical threat affects attack surface and what can be done to defend against it or detect activity related to its use in the organization's environment.

In most cases, the type of threat intelligence describes attack technology used in one stage of the ATAC Life Cycle or that has a single function. This type of threat intelligence contains technical details regarding what is known about the attack technology in terms of the Attacker's Trifecta components.

The most common sources of single stage technical threat intelligence are:

- Disclosure of a vulnerability, configuration error, etc.;[8]
- Release of exploit code that is used for the initial exploitation of a cyber resource and to provide a point of entry (PoE) on a vulnerable target; and,
- Reports of malware that is used to perform attack operations not related to PoE. In ATAC, the malware framework, its plug-ins, and installation techniques (droppers vs. loaders) are considered individually in order to ensure response plans can adequately address the technical threat.

Sometimes, the threat intelligence describes attack campaigns or is derived from a historical review of an organization's cyber security incidents. Attack campaigns or multi-incident reviews generally focus on how an adversary performed the work, bypassed existing security, and how different attack technology or techniques were used across the life cycle of the attack.

The most common sources of multi-stage technical threat intelligence are:

- Reports of attack campaigns by groups like the Department of Homeland Security's Industrial Control Systems – Cyber Emergency Response Team (ICS-CERT), security companies, or research teams; and,
- Historical reviews of security incidents experienced by the organization performing the ATAC evaluation.

Defenders must perform the following tasks in order to identify the type of threat intelligence and what attack technology:

1. Ascertain the source of the threat intelligence, i.e. vendor security update, disclosure by an individual researcher, etc.
2. Determine whether the attack technology described was or would be used in a single ATAC life cycle stage or in multiple stages.
3. If the attack technology described is single stage, complete the steps 3a. through 3c.
    a. Decide what type of technical security information is being described.
        i. Vulnerability or configuration weakness,
        ii. Exploit code,
        iii. Attack technique used for a single purpose, e.g. pivot technique like Pass-the-Hash; or,
        iv. Payload (RAT, Trojan, rootkit, etc.) or malware related to a payload (C&C channel, dropper, loader, etc.).
    b. Determine what task the attacker would use the attack technology or technique to perform, i.e. (initial exploitation, pivoting, C&C communications, etc.).
    c. Identify which stage of the ATAC Life Cycle it would be used in.
4. If the attack technology described is multi-stage, then complete steps 4a. through 4d.
    a. List all attack technology and techniques described.
    b. Decide what type of technical security information is being described for each item on the list.
        i. Vulnerability or configuration weakness,
        ii. Exploit code,
        iii. Attack technique used for a single purpose; or,

[8] The NVD CVE definition of vulnerability or exposure was taken from http://cve.mitre.org/about/terminology.html.

iv. Payload malware (RAT, Trojan, rootkit, etc.) or malware related to a payload (C&C channel, dropper, loader, etc.).

c. Determine what task the attacker would use the attack technology or technique to perform, i.e. (initial exploitation, pivoting, C&C communications, etc.) for each item on the list.
d. Identify each stage of the ATAC Life Cycle the different pieces of malware would be used in.

Simple Threat Analysis

Following the identification of a threat, the Simple Threat Analysis (STA) is comprised of three steps: Attack Surface Analysis (ASA), the development of an attack timeline (or order of operations), and attack path modeling. These steps are further described below:

Step 1. Performing Attack Surface Analysis

ATAC analysis breaks down how adversaries use different kinds of attack technology to perform a range of tasks throughout the life cycle of the attack. In ATAC, attack technology and the technical threat presented to an organization are expressed in terms of the functional security layers (FSL) and ATAC FSM.[9]

Use of ASA in ATAC helps defenders gather make decisions about the threat introduced by attack technology, including:

- How the technical threat affects attack surface exposure, the probability an attack using the specific attack technology would succeed, and cyber security risk;
- Whether or not a response is required;
- The timeline for a response if one is required; and,
- The technical and operational parameters a response plan must take into account.

As in ReACT, ASA provides the basis from which response plans can be developed. In ATAC, though, ASA is used to identify potential targets within the organization's network and how threat is introduce via the Attacker's Trifecta.[10]

Step 1a. Performing a Functional Security Translation and Developing the Functional Baseline of the Target System

The functional security translation, as shown in Figure 2, lays the groundwork for the remainder of the ATAC process. This is a simple organization step which involves the collection of any information related to the threat. Technical information organized by functional security layer within the Functional Baseline of the Target column of the ATAC FSM. Non-technical information should be included within an analyst's notes. Not only does this step ensure that the information can easily be found and understood, it also highlights known and unknown information for the analyst.

The steps for gathering the technical configuration data of the target system are included below. However, in many cases it is likely the analyst will have to deal with imperfect or incomplete information.

1. Identify which layers of the FSL are explicitly targeted by the exploit code or malware.
2. Create a list of alternative or related software could potentially be affected.

---

[9] Additional information regarding functional security layers can be found in Appendix III.

[10] The Attacker's Trifecta is composed of three parts: a vulnerability (CVE) or weakness (CWE) on the target system, an exploit code specific to the vulnerability, and an attack path.

3. Identify which layers of the FSL could be targets on the potential targets.
4. Record all results in an FSL table.

Included below is a simple example of the process. Analysts took this information directly from a vulnerability report. When completing the functional security translation, it is imperative that analysts take note and evaluate the quality of the information. While in this case the strength of the information was high, this is not always the case. Put simply, not all sources are created equal. Additionally, some information may require additional verification. This information does not need to be formally recorded in the FSM, but additional notes will assist with the remaining analysis.

Functional Security Layer	**Functional Baseline of Target**
UR&R	
Network	**TCP/IP**
Firmware	
Operating System	**Microsoft Windows (all versions)**
Virtualization	
Applications	**Browser components or plug-ins (ActiveX controls)**
	Internet Explorer
	Windows Explorer
	WellinTech Kingview
	.NET Framework
	Visual Studio or IDE (debugging)
Cloud, hosted, or vendor services	
Custom code	
Data & Data Stores	

Figure 2. The Functional Security Translation.

Step 1b. Developing Communications Map of the Target System

Following the development of the functional baseline of the target system, analysts should proceed to the communication map development. The communication map describes the ports, services, and protocols used for network communications on the target system. In order to be successful, the adversary requires not only vulnerability-specific exploit code but also having a means of delivery. Developing a communications map of the cyber resources allows defenders to see what attack paths are open to an adversary.

In this step, the communications paths are added to the FSL chart created for the functional baseline, as shown in Figure 3. Doing so allows the disparate types and sources of information regarding the resource to be gathered together in an easily understandable fashion. The steps for mapping the communications paths of critical cyber resources are described below:

1. Record the actual attack paths identified in the threat intelligence in the attack path map.
2. Identify any potential attack paths not explicitly called out in the threat intelligence.
    a. Research the software targeted by the attack technology for commonly used, alternative communications paths the software could use.
    b. Note these paths as potential attack paths on the FSL.

If the analyst has access to the targeted system, the following steps may provide additional assistance:

1. Perform an inventory of all physical communications components on each critical cyber resource, noting the status of each (enabled, disabled).
2. Gather host-based data regarding the normal ports and services running on each cyber resource.
    a. Reviewing network statistics (netstat) data from operating system;
    b. Recording the ports and services open on the system throughout the week;
    c. Performing a network scan (if possible) to confirm the open ports and services; and,
    d. Comparing all of the host-based ports and services configuration data to identify any anomalies with what was expected and what is actually available on the host.
3. Gather network-based data regarding the normal ports and protocols used by each cyber resource.
    a. Capture ingress and egress traffic with a packet capture tool (WireShark) for 1-2 hours at varied periods of time a minimum of 4 times.
    b. Note ingress and egress network traffic patterns and track expected communications by:
        i. Source IP or MAC address;
        ii. Destination IP or MAC address;
        iii. Protocol used to communicate;
        iv. Port and service used for communications; and,
        v. Any errors resulting from the communications.
4. Record the ports, services, and protocols from the network traffic baseline in the FSL chart.

Functional Security Layer	Functional Baseline of Target	Communications Map		
		Protocol	Services	Ports
UR&R				
Network	TCP/IP			
Firmware				
Operating System	Microsoft Windows (all versions)	TCP, UDP	RPC over HTTP	80
		TCP, UDP	RPC Endpoint Mapper (DCOM, Client/Server, etc.)	135
		TCP, UDP	SMB via NetBIOS API	137
		UDP	SMB via NetBIOS API	138
		TCP	SMB via NetBIOS API	139
		TCP, UDP	RPC over HTTP	443
		TCP, UDP	SMB, named pipes	445
		TCP, UDP	RPC over HTTP	593
Virtualization				
Applications	Browser components or plug-ins (ActiveX controls)	Unknown	Kingview custom protocol (if any)	Unknown
	Internet Explorer	TCP, UDP	HTTP	80
	Windows Explorer	TCP, UDP	HTTPS	443
	WellinTech Kingview	TCP, UDP	HTTP (proxy, load balancer)	8080
	.NET Framework			
	Visual Studio or IDE (debugging)			
Cloud, hosted, or vendor services				
Custom code				
Data & Data Stores				

Figure 3. Communications Map Development on the ATAC FSM.

Step 2. Attack Timeline Development

Following communication map development, analysts should determine, as best they can, the timeline of the attack. This should include a listing of the various events of a previous attack which includes any information about how an adversary moves within a compromised system or network, the delivery date of the payload, and any changes to the system or network. This information can be represented on a standard timeline or through functional use at certain stages of either the ATAC or ReACT Life Cycles.

In the event of a system breach, the system management logs will provide a great amount of information, but analysts should bear in mind that the attack may not progress linearly (i.e. an adversary may entertain breaks, employ different tools, or return to a previous system). Essentially, the attack plan for an adversary is not a static thing, but something that can evolve daily. Additionally, adversary groups may employ elaborate evasion techniques, destroying evidence on previously accessed systems. Because of this, security professional or analysts considering system logs should be on the lookout for anything that appears out of the ordinary, paying close attention to time stamps. Timelines assist the analytical process, by identifying key information about an attack including style (techniques and tradecraft, i.e. evasion methods and attack evolution) and other potential points of entry. For the preemptive analytical effort, this information can be invaluable.

Step 3. Attack Path Modeling

The final step of simple threat analysis involves the creation of an attack path model. This step builds upon the information gathered during the communication map (of the target system) development. The attack path model takes into account the communications map and the attack timeline in order to articulate a more complete view of an adversary's movements within a network. In many cases, analysts will have to deal with imperfect information and will only be able to estimate the attack path for an attack. However, even in these instances attack path modeling is a vital component of the ATAC process which facilitates a reexamination of an attacker's workflow.

Complex Threat Analysis

As mentioned earlier in this document, CTA represents the most challenging aspects of the ATAC process.

Step 1. Predictive Attack Path Analysi*s*

ATAC can be used to improve defenders' consumption of threat intelligence and to help them plan an effective response to a known technical threat. However, other analysis techniques can be used to provide insight regarding technical threat. One of these is PAPA, which is an analytical process for anticipating the most likely attack paths an adversary would leverage to compromise a targeted cyber resource given its exposed attack surface.

PAPA is used when defenders want to improve the defensive and detection capabilities around their most high value cyber resources. The attack surface, which is composed of the functional baseline, communications paths, and existing security measures, of the target systems is considered from the attacker's perspective.

Analysts can use PAPA in order to relate a target's attack surface exposure to the Attacker's Trifecta. This allows groups to anticipate how the adversary will run an attack. It is performed after ReACT assessment has been completed. To perform PAPA for a cyber resource, complete the following steps:

1. List all communications paths available to the target in an ATAC FSM.
2. Record all software and services accessible via the communications path by FSL in the ATAC FSM.
3. List all known vulnerabilities or weaknesses at each FSL.

a. Research the known vulnerabilities to see if exploit code or an attack technique exists for them.
b. If so, make note of the recommended security measures available to counter each exploit or attack technique.

4. Log the existing defensive security measures for each FSL.
    c. Identify any FSL that does not have at least 3 defensive and 2 detection controls.
5. Compare the recommended security measures for known vulnerabilities to existing security measures at each FSL.
    a. Identify any FSL whose existing security measures do include the recommended security measures.
6. Prioritize potential attack paths based on the:
    a. Number of vulnerabilities on each FSL;
    b. Gaps between existing and recommended security measures for each vulnerability; and,
    c. Variation from the recommended baseline (3 defensive, 2 detection measures) for existing security at each FSL.

Step 2. Forecasting Threat Technology

In Forecasting Threat Technology, analysts anticipate what types of attack technology or techniques are most likely to be used against high value cyber resources based on a core system's functional baseline and available communications paths. Known threats or attack techniques are evaluated vis-à-vis the functional baseline and communications of high value targets. Included below are recommended steps for the completion of an attack technology forecast:

1. Record all software by FSL in an ATAC FSM.
2. List all communications paths for the components of the functional baseline.
3. Research known attack technology and techniques available for each component in the functional baseline.
    a. Note which:
        i. Components of the functional baseline are affected;
        ii. Layer of the FSL the vulnerable components are located on; and,
        iii. Attack paths are used to run the attacks.
    b. Note the recommended security measures for each identified attack technology or technique.
4. Log the existing defensive security measures for each FSL.
    a. Identify any FSL that does not have at least 3 defensive and 2 detection controls.
5. Compare the recommended security measures for known vulnerabilities to existing security measures at each FSL.
    a. Identify any FSL whose existing security measures do include the recommended security measures.
6. Prioritize potentially effective attack technology or techniques based on the:
    a. Number of identified examples for each component of the functional security baseline;
    b. Number of components on each FSL with publicly available attack technologies or techniques;
    c. Gaps between existing and recommended security measures for each vulnerability; and,

d. Variance from the recommended baseline (3 defensive, 2 detection measures) for existing security at each FSL.

Developing the ReACT Response Plan

The final step of the ATAC process is the development of a ReACT Response Plan. This step is intended to make the necessary changes to the security plans in order to reduce attack surface, correct any root cause failures, and address any gaps. Included below is a general strategy for the development of a ReACT Response Plan:

1. Review ReACT Defensive FSM for each cyber resource for FSL layers that do not have at least three defensive controls in each of the four SDLC life cycle stages. Add additional controls as necessary.
2. Review ReACT Detection FSM for each cyber resource for FSL layers that do not have at least detection controls in each of the four SDLC life cycle stages. Add additional controls as necessary.

When developing a ReACT Response Plan it is helpful to keep in mind the goals of an adversary. The primary goal of any adversary is the execution of remote arbitrary code and the elevation of privileges from unauthorized or limited access to administrative. In order to be successful, an adversary's attack requires three basic requirements:

- A vulnerability (CVE) or weakness (CWE) on the target system;
- An exploit code specific to the vulnerability; and,
- An attack path.

A defender only needs to limit an adversary's ability to use one of these three components. This technique of limiting the attacker's access to required resources is referred to as attack surface reduction. By reducing the attack surface, the defender accomplishes the following:

1. Attack surface is reduced, minimizing the number of ways an attacker could compromise the target system.
2. The probability of a successful attack occurring on a target system or resource is reduced because the adversary's options for running an attack have been decreased.
3. Enterprise risk is lowered because the probability of a successful breach via cyber means has declined.

Appendix II – Explanation of the ATAC Lifecycle

Kill Chain Operations and the ATAC Lifecycle:

As presented in Lockheed Martin's 2010 research, a kill chain model is a systematic method for disrupting an adversary's efforts.[11] In the course of their work, the Lockheed Martin team developed a new kill chain model for intrusions based on the necessary steps for payload delivery (illustrated in Figure 1).

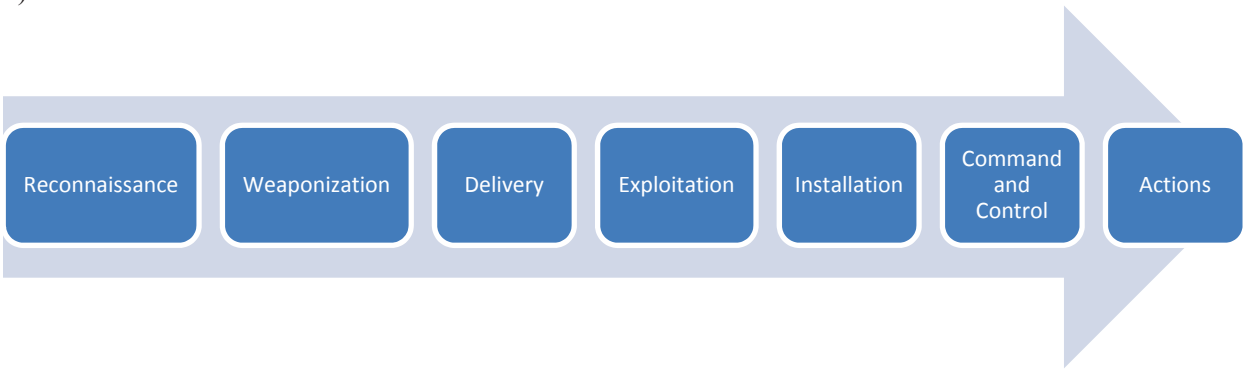

Figure 1: The Lockheed Martin Intrusion Kill Chain.

Lockheed Martin's Intrusion Kill Chain has been widely adopted in the cyber security industry, in part due to the emergence of advanced persistent threats (APT), a challenge which required a reworking of traditional incident response. By using a kill chain model to describe the various phases of an attack, the Lockheed Martin researchers hoped to increase the understanding of elaborate attacks, and thereby increase the success rate of intrusion detection and mitigation. However, Lockheed Martin's Intrusion Kill Chain provides little guidance on security posture development. At the same time it complicates the workflow of an attack and implies a single adversary. In short, it requires that security professionals understand an attack at a level unnecessary for defense, and makes the model inaccessible and difficult to implement within their own organizations.

The ATAC methodology attempts to address these concerns by refocusing cyber security efforts on impact-driven threat analysis within the context of the ATAC Lifecycle (as depicted in Figure 2). This lifecycle forms the foundations for intelligence-driven mitigation, by providing the basis for consistent and repeatable attack analysis. Without this consistency, the sharing of threat information outside an organization, is at best limited and at worst insignificant.

[11] Eric M. Hutchins, et al., "Intelligence-Driven Computer Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains," (paper presented at the 6th International Conference on Information Warfare and Security, 2010). Retrieved from http://www.lockheedmartin.com/content/dam/lockheed/data/corporate/documents/LM-White-Paper-Intel-Driven-Defense.pdf The process is known as a kill chain since disruption anywhere along the process can result in the failure of an attack.

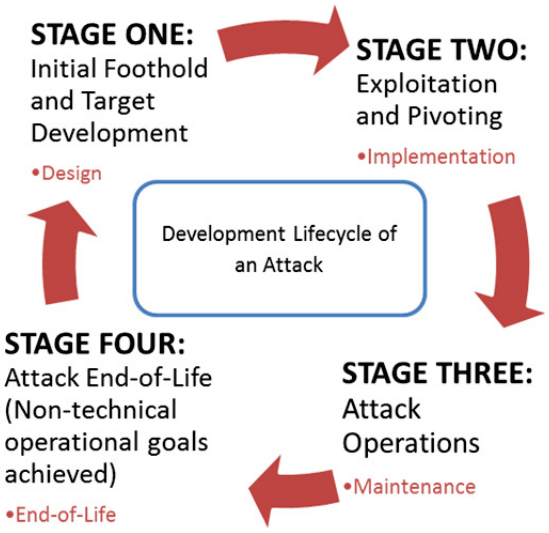

Figure 2: The ATAC Lifecycle. System Development Lifecycle concepts are indicated in red.

The Development Lifecycle of an Attack addresses adversaries from a project management perspective. Attack groups, both simple and complex, rely heavily on business development concepts in order to ensure efficiency of operations. They manage financial and staff resources, distribute work according to skill and experience, and are constantly looking to upgrade their technical systems to improve efficiency. This perspective is useful when considering mitigation strategies. Any disruption of an adversary's work flow can have a significant impact on the progress of an attack. The use of ReACT and ATAC will show security professionals how to disrupt these attacks. (The ReACT Process provides companies and practitioners with a standardized, methodical approach for assessing and mitigating cyber security risks.)[12]

[12] Bri Rolston, "The Response Analysis & Characterization Tool (ReACT)," (Idaho National Laboratory, 2013).

Appendix III – Explanation of the Functional Security Layers

The functional security layers (FSL) are a fundamental component of ReACT. The FSL tool is simply a data schema that allows large amounts of technical data to be gathered methodically and represented in an easily-understood format.

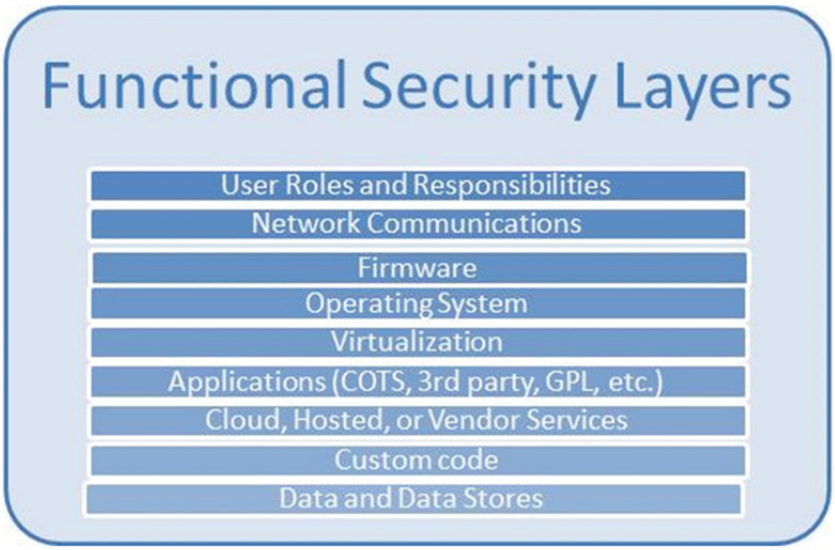

The advantage of evaluating core systems (and their technology) according to FSL is that individual components and attack vectors are less likely to be overlooked. If information regarding core systems' components or existing security posture is not gathered, security issues may not be recognized during the ReACT process. Failure to identify those issues may provide an opening for an attacker, thereby increasing an organization's risk exposure.

Appendix IV – ATAC Operational Security Baseline Template = Functional Security Matrix

Functional Security Layer	Functional Baseline of Target	Attack Path Model		
		Protocol	Services	Ports
UR&R				
Network				
Firmware				
Operating System				
Virtualization				
Applications				
Cloud, hosted, or vendor services				
Custom code				
Data & Data Stores				

Bibliography

Clark, Robert M. *Intelligence Analysis: A Target-Centric Approach* (2nd edition). Washington, D.C.: CQPress, 2007.

Heuer, Richards J. and Randolph H. Pherson. *Structured Analytic Techniques for Intelligence Analysis*. Washington, D.C.: CQ Press, 2011.

Hutchins, Erin M., et al. "Intelligence-Driven Computer Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains." Paper presented at tge 6th International Conference on Information Warfare and Security, 2010. http://www.lockheedmartin.com/content/dam/lockheed/data/corporate/documents/LM-White-Paper-Intel-Driven-Defense.pdf

Rolston, Bri (2013). "The Response Analysis & Characterization Tool (ReACT)." Idaho National Laboratory, 2013.

ATAC and ReACT Data Definitions

Advanced Persistent Threat (APT):

A group with the intent, means, and capability to sustain a long-term attack against an adversary.

Attack Technology, Analysis and Characterization (ATAC) tool:

A process by which individuals or organization can methodically characterize threats and develop technical response plans to manage cyber risk presented by threats.

Applications:

One of the Functional Security Layers (FSL), applications refers to any commercial off-the-shelf (COTS) applications that are default applications which come with a system or those loaded onto a system by users. Unlike custom or proprietary software, the source code of the application is NOT owned or managed by the core system owner. Examples of applications include remote server management, internet browser, database, media player, or web server software.

ATAC Lifecycle:

A lifecycle model that characterizes the stages of an attack into four stages: Target Development, Exploitation and Pivoting, Attack Operations, and Attack End-of-Life. In order to execute a successful attack, an adversary must perform all these stages, although the steps necessary to meet goals at each stage may vary from attack to attack. The ATAC Life Cycle mirrors the Software/System Development Life Cycle, but it represents how an attacker or attack team must manage the work flow and planning associated with an attack.

Attack End-of-Life (EoL):

The final stage of the ATAC Lifecycle during which the operational goals of an adversary are met. In the Attack EoL, technical work is peripheral to achieving strategic goals like data exfiltration, etc.

Attack Operations:

The third stage of the ATAC Lifecycle which occurs after an adversary has gained entry to a network. In this stage, the adversary establishes a foothold on the network that allows him to manage compromised systems remotely, establish command and control communications (C&C), or perform evasive maneuvers to escape detection. He or she will also begin identifying target systems, information, and credentials necessary to meet their operational goals.

attack methodology:

The combination of attack technology and techniques employed by an attacker in order to gain unsanctioned access to or to compromise a system. Attack methodology tends to be unique to an individual or team because the selection of specific tools, techniques, and methodologies combined to perform an attack will vary based on the adversary's or team's skills, preferences, and work flow management.

Attack Path:

The combination of network protocols, host-based ports and services, and physical communications components used by an adversary to deliver an exploit for initial access, upload payloads to compromised systems, or to manage C&C networks.

Attack Path Modeling:

The third piece of Simple Threat Analysis, attack path modeling is the process by which analysts identify the attack paths of high value targets and associate them with likely attack techniques.

Attack Surface:

The attack surface of a system is the combination of exposed attack paths which can be used to compromise a system.

Attack Surface Analysis (ASA):

An analytical process which involves evaluating the exposure of an application, system, or network to attack vis-à-vis a comparison of its base technology build, available attack paths, existing security posture, and known attack methodology. Simply put, attack surface analysis is determining the gap between what is well protected and what is not. Versions of attack surface analysis are conducted during both the ATAC and ReACT processes.

Attack Technique:

The type of attack used to perform key tasks like Elevation of Privilege (EoP)during the attack but not specific to a specific attack technology. Example attack techniques include SQL injection, heap spraying, reverse proxy communications, C&C beaconing, etc. and do not refer to specific pieces of malware or exploit code used in attacks.

Attack Technology:

The actual technology used to perform an attack or attack operations. Attack technology refers to the specific toolkits, payloads, and exploits code employed by an attacker to gain unsanctioned access to or to compromise a system.

Attack Timeline:

A chronological listing of the various events of a previous attack which includes information about how an adversary moves within a compromised system or network, the delivery date of the payload, and any changes to the system or network. May be represented on a standard timeline or through functional use at certain stages of ATAC Lifecycle.

Attack Timeline (or Order of Operations) Development:

The first step of Complex Threat Analysis concerned with determining the order of adversary actions and movements within a system or network. When chronicling a past attack, this step involves the creation of an attack timeline; however, when evaluating potential attacks, this step focuses on the attack workflow with an order of operations.

Cluster Analysis:

An optional step, performed in conjunction with Root Cause Failure Analysis, cluster analysis is a breakdown of the root cause failures in order to determine if an organization has a systemic failure or cluster of failures in one aspect of DIME or the Organizational Hierarchy. For example, consistent RCFA indicators in the design stage of an application roll-out may indicate a systemic failure in the way business requirements are used to select or design technology. Cluster analysis is also used to determine if the people, process, or technology sub-components of the FSL contribute to repetitive failures.

Communications Map:

A listing of all the pathways (ports, protocols, physical components) used by a cyber component to communicate within a system or network.

Communications Map Development:

The second step of Simple Threat Analysis (STA) Communications Map Development focuses on the path used by an adversary to deliver exploit code or malware to a vulnerable system. Communications Map Development can also be used to determine the communications paths leveraged by an attacker to perform attack operations, which will typically include both ingress and egress traffic from a compromised cyber resource. This delivery information is organized in the ATAC Functional Security Matrix (FSM) by Functional Security Layer (FSL). For example, a web browser based attack employing malicious Java applets may use ports 80, 8080, and 443, information which would be recorded by the ATAC tool user in the Applications section of the Communications Map Information column of the ATAC FSM.

Complex Threat Analysis (CTA):

Following Simple Threat Analysis (STA), CTA encompasses Predictive Attack Path Analysis (PAPA) and Forecasting Attack Technology. This portion of the ATAC process is intended to provide organizations with additional information for use during preemptive security efforts.

Core System:

Also known as critical work components, a core system is comprised of any software, communications technologies, data, or services necessary for an organization to achieve their mission goals. A core system may be used to perform multiple tasks associated with the mission goals but has only a single technical purpose. An example of a core system might be the Front End Processor (FEP) used to manage multiple, remote endpoints in the process environment.

Core System Identification:

A component of the Impact Driven Risk Analysis (IDRA) methodology, Core System Identification is the process by which groups identify systems necessary for critical workflow. Core systems must be identified by both the function and technical make-up of a process or ICS function. For example, the core systems used to manage the transmission of an oil pipeline might include the remote flow sensors, the RTUs or PLCs managing those sensors, the central ICS management system that aggregates data and process control from multiple PLCs or RTUs, and the historian database used to push the flow data to a billing server.

Critical Impact Identification:

Also known as impact identification. A component of the Impact Driven Risk Analysis (IDRA) methodology, Critical Impact Identification is the process by which groups characterize and rank the effects associated with a loss of function. In terms of cyber security and IDRA, critical impact allows organizations to prioritize their risk and to allocate resources necessary for mitigation efforts.

Critical Workflow:

Any work function which is necessary for an organization to achieve their business goals.

Critical Workflow Identification:

Also known as development of the functional overview. A component of the Impact Driven Risk Analysis (IDRA) methodology, Critical Workflow Identification is a step conducted by organizations prior to beginning a ReACT assessment, the goal of which is to identify core systems and critical impacts associated with the workflow. Put another way, a critical workflow is comprised of the processes necessary to ensure an organization's mission goals. In the context of the ReACT assessment, the Critical Workflow Identification limits the focus of

the assessment to an organization's most important business processes, the critical impacts associated with those processes, and the core systems that could be used by an adversary to realize critical impact.

Custom or Proprietary Software:

One of the Functional Security Layers (FSL), custom or proprietary software (aka "custom code") refers to any software for which the system owner is responsible for maintaining the source code, application, or scripts throughout the Software Development Life Cycle. Examples include remote management scripts used to manage ICS server, ladder logic used by ICS that must be maintained for the process to run, and applications or software written by an integrator or 3rd party.

Cyber Resources:

The various components which make up a core system. Cyber resources may include any combination of user activities, applications, data, services, or systems.

Cyber Risk:

An individual's or group's intentional or accidental exposure to danger, loss, or harm through data, services, computers, networks, or technology.

Data and Data Storage:

One of the Functional Security Layers (FSL), data and data storage refers to the information and information holding components used by a core system. Data storage may include any permanent or temporary storage mechanism used by the read, write, or execute functions for storing, transmitting, or manipulating data. An example includes the temporary cache of sensor data in the ICS application when the sensor data is synchronized between ICS devices.

Emerging Threat Identification:

The preliminary step of the ATAC process, which is characterized by the identification of a new threat (attack technology, adversarial group, etc.).

Exploitation:

The fourth step of the Lockheed Martin Intrusion Kill Chain which is characterized by the execution of malicious code on the target system in order to gain an initial Point of Entry on a target network or to pivot from one foothold to another in an already compromised network.

Exploitation and Pivoting:

The second stage of the ATAC Life Cycle which corresponds to the Implementation phase of the System/Software Development Life Cycle.. This stage encompasses the Lockheed Martin Intrusion Kill Chain steps of Delivery and Exploitation. In this stage, an adversary gains an initial Point of Entry (PoE) to one or multiple systems or pivots to another system or systems using the foothold gained during the initial PoE. It is during this stage that additional target systems may be identified for attack and to further the adversary’s operational goals.

Forecasting Threat Technology:

Part of Complex Threat Analysis (CTA), Forecasting Threat Technology is the process by which organizations identify either the most likely delivery mechanisms for future attacks against specific targets or, alternatively, identify the most probable evolutions of known attack technologies and techniques necessary for bypass and evasion.

Firmware:

One of the Functional Security Layers (FSLs), firmware (or embedded device software) is a combination of special purpose hardware, persistent memory, program code and the data stored on it. Examples of firmware include: BIOS, chipset, video card, programmable logic controller (PLC), or smart meters.

Functional Baseline:

An output of the ReACT assessment, the Functional Baseline is a picture of a core system's technical architecture. Functional baseline information is expressed as a list of the core system's software inventory broken-down by Functional Security Layer (FSL). In order to perform RCFA or Cluster Analysis, the Functional Baseline must be also understood in terms of its use case functionality, i.e. who uses it and why, and how the architecture of the core system meets those business needs.

Functional Overview:

See Critical Workflow Identification.

Functional Security Layers (FSL):

A means of organizing functional baseline information (during the ReACT process) and attack technology attack workflow information (during the ATAC process). The Functional Security Layers provide cyber security professionals with a means of organizing and synthesizing data, and ensure that relevant information is not overlooked during the ReACT and ATAC processes.

Functional Security Layers (FSL) subcomponents:

The subcomponents of the FSL are people, process and technology and are incorporated into the FSL rows in the Functional Security Matrix (FSM). The subcomponents are used primarily for performing RCFA and Cluster analysis during a ReACT assessment.

Functional Security Matrix:

A template used during both ATAC and ReACT assessments to organize information from the FSL, the FSL subcomponents, and the ATAC or ReACT Life Cycles. This information is then used to analyze information regarding either an attacker's use of threat technology throughout the attack work flow or the defender's management of core systems throughout the SDLC.

Functional Security Translation:

The first step of Simple Threat Analysis (of the ATAC Process). A process for organizing known attack technology information within the Functional Security Layers (FSL). The information organized in this step should be highly detailed and address how the attack technology was used by the adversary across the various functions of the ATAC Life Cycle.

Heat map:

An optional product of the ReACT process, a heat map is a visual representation of the cyber risk an organization faces. For the ReACT process, components are graphed along two axes, Impact (the vertical axis) and Probability (typically the horizontal axis).

Hosted and Cloud Services:

One of the Functional Security Layers (FSL), hosted and cloud services refers to any combination of hosted, managed, 3rd party, or cloud services used by an organization to perform a core function in the critical work flow. Hosted or cloud services includes any of the

*as-a-Service offerings (software, platform, or infrastructure) provided by a 3rd party provider over the internet or wide area networks.

Impact Driven Risk Analysis Methodology (IDRAM):

A risk identification and analysis methodology which characterizes risk in terms of what's important to a business and what core systems could be used to impact an organization's ability to perform.

In-house Analyst:

For the purposes of this document, an in-house analyst is an analyst, cyber security professional, etc. who is a full time employee of a company.

Intrusion Kill Chain:

A model developed by Lockheed Martin to describe the actions conducted by an adversary from the conception of an attack through its completion. The model is described as a kill chain to emphasize the interdependency of the steps; disruption anywhere along the Intrusion Kill Chain will result in a failed attack.

Incident Response:

Also known as incident management, it is the process of responding to or managing a functional or security-impacting incident that caused (or may cause) an interruption or a reduction in the quality of an IT or ITC service. ATAC, ReACT, and IDRAM all utilize the ITIL concepts of incident and problem management.

Network Communications:

One of the Functional Security Layers (FSL), network communications must be understood by security professionals as they can provide adversaries with a remotely-accessible attack path to a targeted system. Examples of network communications include: Distributed Network Protocol 3 (DNP3) communications between a control center and an RTU; CompuTrace beaconing from the BIOS of a CompuTrace-protected system over any DSL, Ethernet, wireless, or satellite communications channel; cellular communications from a remote substation to a control center; and Secure Shell (SSH) connection over TCP/UDP port 22 used to manage a server remotely.

Operating System (OS):

One of the Functional Security Layers (FSL), in its simplest form an OS is the software collection that manages resources and provides services for computer programs. An OS can be physical, virtual, embedded or mobile in nature. Examples include: the platform operating system that the Energy Management System (EMS) runs on; an embedded Linux kernel in an RTU, FEP, or PLC; or an Android OS on the mobile phone hosting the web-enabled Human Machine Interface (HMI) application used to manage a SCADA network.

Order of Operations:

Also known as the Order of Attack Operations, this is the general order of operations for a potential attack. This concept is used to understand the workflow of an attack, which are expressed as the Target Development, Exploitation and Pivoting, Attack Operations, and Attack End of Life (EoL) in ATAC. The Order of Operations chronology differs from the attack timeline, which is a chronological order of operations for a past attack.

Payload:

The malicious software, aka malware, which is loaded on compromised systems by the adversary to perform attack operations. Payload may also be known as exploit kits, Trojan horses, or Remote Access Trojans (RATs) and is used to manage attack operations work such as C&C communications, data exfiltration, etc.

Pivoting:

The process by which an adversary identifies new potential targets and moves from one compromised system to another by employing additional attack technology. Pivoting occurs after the adversary has gained an initial foothold on the network through a Point of Entry (PoE) attack and is transitioning to another network segment or target type. When pivoting through a network, attackers generally use a second type of Elevation of Privilege attack technique to move around, not the initial PoE attack used on the PoE systems.

Predictive Attack Path Analysis (PAPA):

Based on a security posture or attack technology, predictive attack path analysis is the process of identifying how future attacks are likely to manifest themselves. Put another way, Predictive Attack Path Analysis is the process of identifying an adversary's likely attack path by observing the connections between core and other systems.

Preemptive Mitigation:

Corrections made to a security posture following Simple Threat Analysis (ATAC Process), intended to provide immediate protections from an emerging threat.

Response Analysis and Characterization (ReACT) tool:

An information schema and analysis methodology which provides organizations with an organized and comprehensive approach for assessing and improving their current security posture.

ReACT assessment:

An output of a Response Analysis and Characterization (ReACT) review, the ReACT assessment determines an organizations existing security posture to be used during attack surface analysis, root cause failure analysis and cluster analysis. ReACT assessments provide the basis for building a work plan to address the root cause failures that allow security weakness or vulnerabilities to exist.

Root Cause Failure Analysis (RCFA):

Following the completion of a ReACT assessment, Root Cause Failure Analysis is the process of identifying the reason behind weaknesses in security posture. During the ReACT process, these failures are categorized by Functional Security Layer (FSL), the organizational hierarchy (people, process or technology) and the System/Software Development Life Cycle (SDLC).

Security Posture:

The tools, policies, and protections currently deployed by a group, individual, or organization in order to address threats and maintain security. For example, an organization may frequently review access control lists to ensure that access is limited to only necessary individuals. In terms of the ReACT assessment, the security posture is used to perform attack surface analysis.

Simple Threat Analysis (STA):

Following the identification of a threat, the Simple Threat Analysis (STA) is comprised of three steps: Attack Surface Analysis (ASA), the development of an attack timeline (or order of operations), and attack path modeling.

STRIDE Threat Model:

A system developed by Microsoft for classifying computer security threats. It is comprised of six categories including: spoofing, tampering, repudiation, information disclosure, denial of service, and elevation of privilege.

Subject Matter Expert (SME):

A SME is an individual with substantial expert knowledge in a particular field or topic. For the purposes of this document, the term SME distinguishes a contractor or external cyber security professional or threat analyst from an in-house analyst.

System (or Software) Development Life Cycle (SDLC):

A series of steps for the development of systems or software, including: design, implementation, maintenance and end-of-life. In the ATAC process, the ATAC Life Cycle is based on the SDLC.

Target Development:

The first stage of the ATAC Life Cycle, which includes the process of gaining an initial foothold and corresponds to the Design phase of the System Development Lifecycle. This stage incorporates the Lockheed Martin Intrusion Kill Chain steps: Reconnaissance and Weaponization. During this stage an adversary selects and researches a target in order to identify the most likely means of access.

Threat Technology:

Threat intelligence or technical threat intelligence is the initial input of the ATAC process. Examples of threat intelligence include, but are not limited to, alerts, RSS feed, and informal sharing.

User Roles & Responsibilities:

One of the Functional Security Layers (FSL), user roles and responsibilities (UR&R) defines the relationship between users, computer systems, network, and data. Examples include requiring VPN tokens and credentials to access a company's network or a WPA/PSK key for a specific wireless network.

Virtualization:

One of the Functional Security Layers (FSL), virtualization is any software loaded on a system that virtualizes a substantial feature of the system. While users can virtualize computer hardware platforms, operating systems, storage devices, or other computer network resources, in the case of the FSL virtualization generally refers to operating systems. Examples of virtualization include the platform and centralized management features for VMWare, Microsoft Hypervisor, or Xen.