A PRACTICAL GUIDE TO ADVERSARY ENGAGEMENT

MITRE | Engage™

This page intentionally left blank

Table of Contents

List of Figures

List of Tables

This page intentionally left blank

1 Introduction

1.1 Cyber Denial, Deception, and Adversary Engagement

Cyber defense has traditionally focused on the use of defense-in-depth technologies to deny the adversary access to an organization's networks or critical cyber assets. In this paradigm, any time the adversary can access a new system or exfiltrate a piece of data from the network, they win. However, when a defender introduces deceptive artifacts and systems, it immediately increases ambiguity for the adversary. Is the system the adversary just accessed legitimate? Is the piece of data the adversary just stole real? Questions such as these begin to drive up the cost to operate, while driving down the value of the adversary's cyber operations.

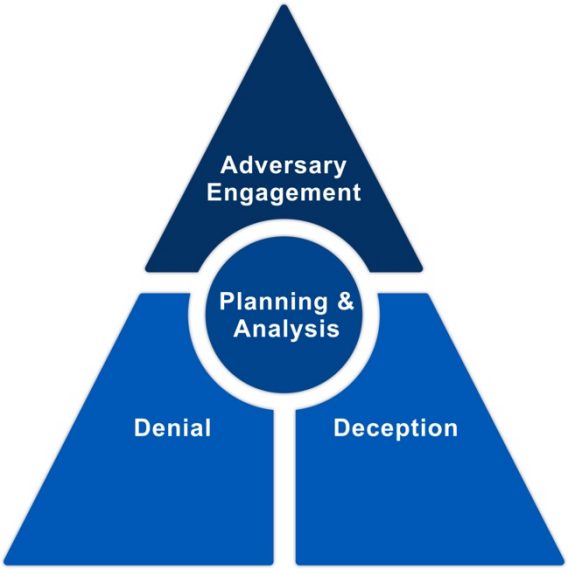

Figure 1: Adversary engagement is the use of cyber denial and deception, in the context of strategic planning and analysis to drive up the cost, while driving down the value of the adversary's operations.

Cyber Denial is the ability to prevent or otherwise impair the adversary's ability to conduct their operations. This disruption may limit their movements, collection efforts, or otherwise diminish the effectiveness of their capabilities. In *Cyber Deception* we intentionally reveal deceptive facts and fictions to mislead the adversary. In addition, we conceal critical facts and fictions to prevent the adversary from forming correct estimations or taking appropriate actions. When cyber denial and deception are used together, and within the context of strategic planning and analysis, they provide the pillars of *Adversary Engagement*.

The main goals of adversary engagement can be any combination of the following: to expose adversaries on the network, to elicit intelligence to learn more about the adversary and their Tactics, Techniques, and Procedures (TTPs), or to affect the adversary by impacting their ability to operate. Adversary engagement operations provide opportunities for the defender to demonstrate tools, test hypotheses, and improve their threat models, all with the added benefit of negatively impacting the adversary.

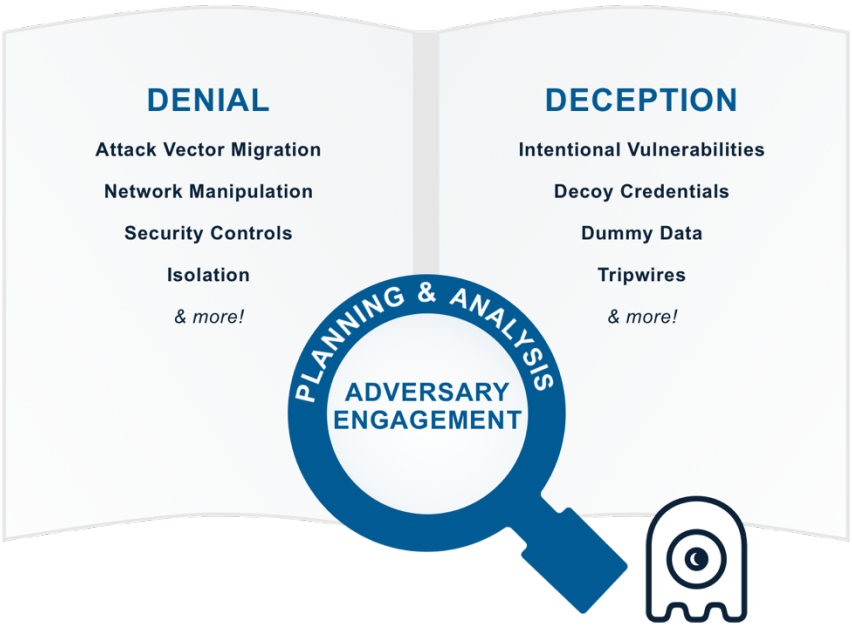

Figure 2: Operators viewing denial and deception activities through the lens of planning and analysis can identify opportunities for adversary engagement.

Denial, deception, and adversary engagement technologies are far from novel. Honeypots, or dummy systems set up for the sole purpose of studying adversary interactions, have been in use for decades. However, they often suffer from the shortcomings of being unrealistic, easily signatured, and general-purpose. In contrast, the engagement environments used in adversary engagement operations are carefully tailored, highly instrumented environments designed on an engagement-by-engagement basis. Often, these environments are seamlessly integrated into a production network. However, these environments need not be complex or even highly technical. Suppose an organization is worried about intellectual property theft. This organization could set up a series of shared directories across the corporate network that are accessible and contain apparently sensitive, but fake, corporate data. If these directories are accessed, the defender will receive a high-fidelity alert. Anyone in the organization can access the information; however, if no one in the organization has a business need to access this sensitive data, any user activity associated with the data warrants a review. These simple decoy directories are an example of utilizing adversary engagement as part of a larger intellectual property protection strategy.

Adversary engagement is an iterative, goal driven process, not merely the deployment of a technology stack. It is not enough to deploy a decoy and declare success. Rather, organizations must think critically about what their defensive goals are and how denial, deception, and adversary engagement can be used to drive progress towards these goals. Unlike other defensive technologies, such as antivirus (AV), adversary engagement technologies cannot be considered “fire and forget” solutions. In the previous decoy directory example, while a relatively trivial operation, the defender still needed to think carefully about their goals when designing the engagement.

Let’s continue to explore this example to understand how goals shaped and drove the operational design. The company was worried about a particular threat. In this case, the threat was a malicious insider. They

identified critical cyber assets of value to their adversary. Specifically, they were concerned about the valuable intellectual property associated with the latest R&D project. Using these facts, the organization was able to create decoy directories that would be appealing to the adversary, by filling the decoys with fake results from the latest R&D tests. The organization also identified key locations likely to be noticed by anyone poking around the network. In this case, they chose the CEO's private SharePoint. With few resources, this exemplar company was able to deploy deceptive artifacts that will likely result in high fidelity alerts to insider threat behavior. As a reminder, this should not be the totality of this company's insider threat protection program. Rather, this adversary engagement operation should augment the larger strategic goals of the program.

While it is easy to imagine the value of successfully deceiving the adversary over a long period of time, given the current state of adversary engagement operational tooling, especially automated and at scale, the adversary may eventually realize the deception. However, this may still provide defensive benefits. For example, the adversary may now doubt the validity of any compromised system on your network. Additionally, the defenders learn what steps the adversary takes to verify a system and can use that information both to identify malicious activity and to better tailor environments in future operations.

1.2 Getting Started with Adversary Engagement

Now that we have explored the concepts of cyber denial, deception, and adversary engagement, we will describe how to get started integrating adversary engagement processes and technologies into your defensive strategy. We will start by exploring the Engage Matrix, the foundation to your adversary engagement strategy. The Engage Matrix is a shared reference that bridges the gap between defenders and decision makers when discussing and planning denial, deception, and adversary engagement activities. At its core, the Matrix empowers the defender to identify their adversary engagement goals, and then use those goals to shape their operational activities. Building on this foundation, we will start to create an operational template that you can adapt to fit your own organization.

It is important to recognize that adversary engagement operations are scalable. This guide will walk through a template for more complicated team operations. However, the basic idea of an adversary engagement operation is simple, as illustrated in the decoy directory example. Additionally, while deception for detection use cases are often associated with low complexity deceptions, they are not the only simple adversary engagement opportunities. All types of operations can also be used effectively, even by teams with limited experience or resources. For example, take a computer, connect it to a network, run the adversary's malware, and study the result. You could start an elicitation operation with a laptop and a non-enterprise Wi-Fi connection. As your adversary engagement goals and capabilities mature, operations can grow quickly in complexity. However, it is important to remember that, as you learn, so will your adversaries. That single laptop probably won't get you far, but it will give you a chance to get your feet wet and begin to experiment with how adversary engagement strategies could inform your organization's defensive strategy.

2 MITRE Engage™ Matrix

MITRE Engage was created to help the private sector, government, and vendor communities plan and execute adversary engagement activities by providing a shared reference to bridge the gaps across these communities. A core offering of MITRE Engage is the Engage Matrix, which was built based on MITRE expertise in adversary engagement and our knowledge of adversary behavior observed in the real world. The Engage Matrix is intended to enable the discussion and planning of adversary engagement, deception, and denial activities to drive strategic cyber outcomes.

Prepare	Expose		Affect			Elicit		Understand
Plan	**Collect**	**Detect**	**Prevent**	**Direct**	**Disrupt**	**Reassure**	**Motivate**	**Analyze**
Cyber Threat Intelligence	API Monitoring	Introduced Vulnerabilities	Baseline	Attack Vector Migration	Isolation	Application Diversity	Application Diversity	After-Action Review
Engagement Environment	Network Monitoring	Lures	Hardware Manipulation	Email Manipulation	Lures	Artifact Diversity	Artifact Diversity	Cyber Threat Intelligence
Gating Criteria	Software Manipulation	Malware Detonation	Isolation	Introduced Vulnerabilities	Network Manipulation	Burn-In	Information Manipulation	Threat Model
Operational Objective	System Activity Monitoring	Network Analysis	Network Manipulation	Lures	Software Manipulation	Email Manipulation	Introduced Vulnerabilities	
Persona Creation			Security Controls	Malware Detonation		Information Manipulation	Malware Detonation	
Storyboarding				Network Manipulation		Network Diversity	Network Diversity	
Threat Model				Peripheral Management		Peripheral Management	Personas	
				Security Controls		Pocket Litter		
				Software Manipulation				

Figure 3: The MITRE Engage Matrix bookends engagement activities with strategic planning and analysis, to ensure that every action is goal driven.

2.1 The Structure of Engage

The Engage Matrix is broken into several components. These components are intended to emphasize that adversary engagement is a process and not something solved merely by deploying a technology stack. It encourages users to consider their strategic and operational goals upfront, and then to align their engagement actions to drive progress towards those goals. We will explore each of these components in the following sections.

2.1.1 Strategic and Engagement Actions

The Engage Matrix is divided into two categories of actions. *Strategic* actions bookend the matrix and ensure that defenders appropriately drive operations with strategic planning and analysis. Additionally, the strategic actions ensure that the defender has gathered appropriate stakeholders (management, legal, etc.) and defined acceptable risk. Throughout the operation, the defender will ensure that all engagements operate within these defined guardrails. *Engagement* actions are the traditional cyber denial and deception activities that are used to drive progress towards the defender's objectives

Prepare	Expose		Affect			Elicit		Understand
Plan	Collect	Detect	Prevent	Direct	Disrupt	Reassure	Motivate	Analyze
Cyber Threat Intelligence	API Monitoring	Introduced Vulnerabilities	Baseline	Attack Vector Migration	Isolation	Application Diversity	Application Diversity	After-Action Review
Engagement Environment	Network Monitoring	Lures	Hardware Manipulation	Email Manipulation	Lures	Artifact Diversity	Artifact Diversity	Cyber Threat Intelligence
Gating Criteria	Software Manipulation	Malware Detonation	Isolation	Introduced Vulnerabilities	Network Manipulation	Burn-In	Information Manipulation	Threat Model
Operational Objective	System Activity Monitoring	Network Analysis	Network Manipulation	Lures	Software Manipulation	Email Manipulation	Introduced Vulnerabilities	
Persona Creation			Security Controls	Malware Detonation		Information Manipulation	Malware Detonation	
Storyboarding				Network Manipulation		Network Diversity	Network Diversity	
Threat Model				Peripheral Management		Peripheral Management	Personas	
				Security Controls		Pocket Litter		
				Software Manipulation				

Figure 4: Strategic Actions are taken to support your operational strategy

Prepare	Expose		Affect			Elicit		Understand
Plan	Collect	Detect	Prevent	Direct	Disrupt	Reassure	Motivate	Analyze
Cyber Threat Intelligence	API Monitoring	Introduced Vulnerabilities	Baseline	Attack Vector Migration	Isolation	Application Diversity	Application Diversity	After-Action Review
Engagement Environment	Network Monitoring	Lures	Hardware Manipulation	Email Manipulation	Lures	Artifact Diversity	Artifact Diversity	Cyber Threat Intelligence
Gating Criteria	Software Manipulation	Malware Detonation	Isolation	Introduced Vulnerabilities	Network Manipulation	Burn-In	Information Manipulation	Threat Model
Operational Objective	System Activity Monitoring	Network Analysis	Network Manipulation	Lures	Software Manipulation	Email Manipulation	Introduced Vulnerabilities	
Persona Creation			Security Controls	Malware Detonation		Information Manipulation	Malware Detonation	
Storyboarding				Network Manipulation		Network Diversity	Network Diversity	
Threat Model				Peripheral Management		Peripheral Management	Personas	
				Security Controls		Pocket Litter		
				Software Manipulation				

Figure 5: Engagement Actions are the actions taken against the adversary that are more often associated with cyber deception.

By dividing the Matrix into Strategic and Engagement categories, we highlight the importance of approaching adversary engagement as a thoughtful and intentional process, guided by the defender's goals.

2.1.2 Engage Goals

The Engage actions are further subdivided into Goals, Approaches, and Activities. Across the top of the Matrix are the Engage *Goals*. Goals are the high-level outcomes you would like your operation to accomplish. The Engage Goals can be adapted to fit a spectrum of use cases and objectives based on needs, resources, and authorities. For example, the deception opportunities available to a corporate defender will differ greatly than those available to a U.S. Cyber Command operator. However, in either case, there are opportunities within each goal to operate safely and legally within the bounds of the user's authorities.

Figure 6: The Engage Goals help defenders drive progress towards strategic outcomes

The Prepare and Understand goals focus on the inputs and outputs of an operation. While the Matrix is linear, it is essential to understand that within and across operations, this process is cyclical. For each operation, you will be continuously iterating and improving as the operation progresses. This iterative cycle will demand that you constantly think about how to Prepare for the next set of operational actions and how to Understand operational outcomes in the context of your larger cyber strategy. This process ensures that you are constantly aligning and realigning your actions to drive progress towards your Engagement Goals. The Engagement Goals are Expose, Affect, and Elicit. These Goals focus on actions taken against your adversary. Let's explore each of these Goals in more detail.

We can *Expose* adversaries on the network by using deceptive activities to provide high fidelity alerts when adversaries are active in the engagement environment. There is often overlap between practices that are considered good cyber hygiene and techniques used to expose adversaries on the network. For example, if the defender collects and analyzes system logs and identifies malicious behavior, that is not an example of using the Engage activity System Monitoring to Expose the adversary. That is just good cyber security practice! However, if a defender places decoy credentials on the network, and monitors the system logs for the use of those credentials elsewhere in the network, that is an example of using Lures and System Monitoring to Expose the adversary using adversary engagement activities. It should be mentioned that, at the time of writing this handbook, most of the commercial vendor offerings focus on deception to Expose adversaries on a system.

We can *Affect* adversaries by having a negative impact on their operations. Affect activities are ultimately about changing the cost-value proposition in cyber operations for the adversary. The defender may want to increase the adversary's cost to operate or drive down the value they derive from their operations. For example, the defender can negatively impact the adversary's on-network operations to drive up the resource cost of doing operations by slowing down or selectively resetting connections to impact

exfiltration. This type of activity increases the adversary's time on target and wastes their resources. To drive down the value of stolen data, a defender could provide an adversary deliberately conflicting information. Providing such information requires the adversary to either choose to believe one piece of data over another, disregard both, collect more data, or continue with uncertainty. All these options increase operational costs and decrease the value of collected data. It is important to note that we limit all Affect activities to the defender's network. We are NOT talking about hack back or any activities in the adversary's space. This distinction is important to ensure that our defense activities are legal! However, this does not mean the impact of our activities are confined to just the time the adversary is operating within our network. For example, if we provide the adversary with false or misleading data, we may impact their future operations by influencing capability development, targeting, and/or infrastructure requirements. This type of information manipulation of our own data is legal and can have longer term impact on adversary behavior.

We can *Elicit* information about the adversary to learn about their tactics, techniques, and procedures (TTPs). By creating an engagement environment that is uniquely tailored to engage with specific adversaries, the defenders can encourage the adversary to reveal additional or more advanced capabilities. To do this, the defender may need to use a combination of documents, browser artifacts, etc. to reassure an adversary and reduce suspicion, while adding enticing data and exploitable vulnerabilities to motivate the adversary to operate. Observing an adversary as they operate can provide actionable cyber

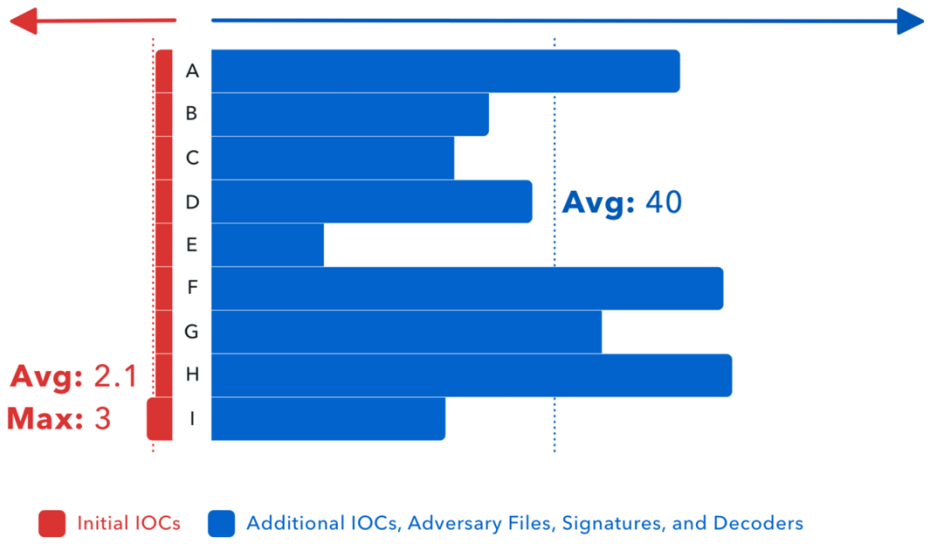

Figure 7: Intelligence gathered before and after nine elicitation operations

threat intelligence (CTI) data to inform the defender's other defenses. Many years ago, when MITRE began our adversary engagement program, we used malware self-infection operations to form the basis of most of our CTI knowledge. As seen in Figure 7, in the set of operations that MITRE ran, we gained on average 40 additional IOCs, adversary files, signatures, or new decoders per operation. While some of these operations requires a higher level of technical maturity (such as writing custom decoders), you can easily get valuable IOCs and adversary files just by running and observing an operation.

At this point, it is important to highlight the difference between these Engagement Goals, your operational objective, and your organization's larger adversary engagement goals. For a given operation, you will

define one or more operational objectives. These objectives are the specific, measurable actions that will enable you to accomplish your organization's larger adversary engagement goals. For example, your goal may be to fill in a known gap in your knowledge of a specific APT. Over the course of one or more adversary engagement operations, you can make progress towards this goal by setting operational objectives such as "Identify at least X many new indicators" or "engage the adversary to obtain a second stage malware file." Goals set direction; objectives take steps in that direction. It is also important to remember that, while we identify three high level Engagement Goals, where adversary engagement fits into your organization's cyber strategy may mean you have organizational goals that fall on a spectrum or build on one another. For example, an operation might allow the adversary to engage freely for a time to elicit new TTPs, before identifying an opportunity to influence or control an adversary's actions to meet a larger strategic goal for your organization. When using the Engage Matrix to plan an operation remember that these categories are intended to provide a language through which your goals can be clearly articulated and not to limit or box in your use of adversary engagement operations.

2.1.3 Engage Approaches

The next row contains the Engage *Approaches*. Approaches let you make progress towards your selected goal.

Figure 8: Engage Approaches are the ways in which you drive progress towards a selected goal

Strategic Approaches help you to focus on the steps you must complete before, during, and after an operation to ensure that your activities are aligned with your overall strategy. Strategic Approaches help ensure that your operations of today inform your operations of tomorrow.

Engagement Approaches help you to identify what actions you would like to take against your adversary and help you to drive progress towards that impact. As seen in Figure 8, Engage outlines 9 approaches to drive progress towards the various goals.

Table 1: Summary of the Engage Approaches

Name	Description
Plan	Identify and align an operation with a desired end-state
Collect	Gather adversary tools, observe tactics, and collect other raw intelligence about the adversary's activity
Detect	Establish or maintain awareness of adversary activity
Prevent	Stop all or part of the adversary's ability to conduct their operation as intended
Direct	Encourage or discourage the adversary from conducting their operation as intended

Disrupt	Impair the adversary's ability to conduct their operation as intended
Reassure	Add authenticity to deceptive components to convince an adversary that an environment is real
Motivate	Encourage the adversary to conduct part or all of their mission
Analyze	Retrospective review of information gained from an operation

2.1.4 Engage Activities

The remainder of the Matrix, as seen in Figure 3 is composed of the Engage *Activities*. Activities are driven by real adversary behavior and are the concrete techniques you use in your approach. You will notice that the same activities often appear under one or more Approach or Goal. Adversary engagement activities can be adapted to fit a spectrum of use cases and objectives based on implementation. For example, a Lure may be used to attract the adversary towards a tripwire (Detection) in one operation, and then used to distract the adversary from a more vulnerable target (Direction) in a second operation.

2.2 ATT&CK® Mappings

When an adversary engages in a specific behavior, they are vulnerable to expose an unintended weakness. In Engage, we look at each ATT&CK® technique to examine the weaknesses revealed and identify an engagement activity or activities to exploit this weakness. By mapping the engagement activities to ATT&CK we can better plan which activities will enable us to reach our strategic objectives.

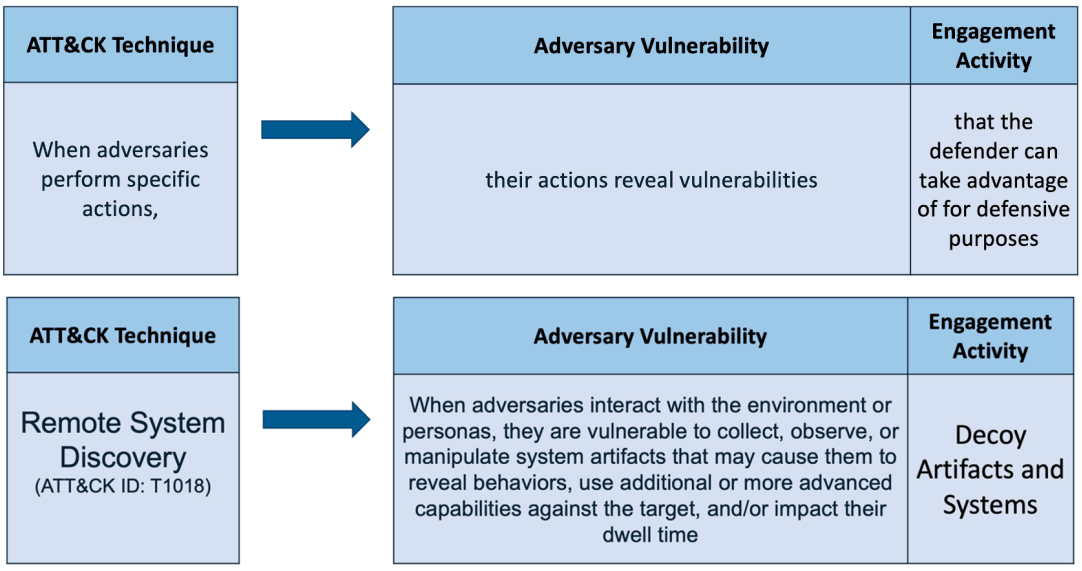

Figure 9: By mapping to ATT&CK, Engage ensures that each engagement activity is driven by real adversary behavior.

2.3 Operationalizing the Engage Matrix

Adversary engagement operations should be an iterative process; continuously refining and changing operational activities based new opportunities. The Engage Matrix is intended to drive discussion and planning activities, not to be a one size fits all outline of adversary engagement operational activities.

Figure 10: Use the MITRE Engage Cycle to integrating Engage into your cyber strategy

Figure 10 outlines the MITRE Engage Cycle and illustrates how the Engage Matrix can be operationalized. This cycle has no defined beginning or end, but for the sake of walking through the model, we will start with *collecting raw data from sensors.* This collection is tool agnostic—it simply refers to your collection methods regardless of how that collection happens. These collection tools can range from low-cost solutions such as collection Windows System Monitor (Sysmon)/Auditd/etc., to vendor Endpoint Detection and Response (EDR) solutions. The next step in the cycle is *analyzing raw data in the context of existing CTI data.* Here you can use tools such as MITRE ATT&CK to contextualize this new data. By analyzing adversary actions and comparing this data to past behavior, the defender can identify patterns, such as common actor TTPs, that offer indications about the adversary's current, and possibly future activities. Armed with this knowledge, the defender can use the Engage Matrix to *identify opportunities for engagement* to meet strategic defensive goals. It should be noted that, as you begin your deception program, you may rely heavily on intelligence feeds, open-source reporting, and/or information sharing groups to learn about the behavior of your target adversary, or the intended target of your deception. As you build your deception capabilities, you can complement this CTI data with insights gained by watching adversaries in your own environment. After you have found opportunities, it is time to *implement your engagements*. At this stage the deceptive assets are deployed and the engagement begins.

As previously discussed, adversary behavior should drive the engagement. Each time the adversary interacts with the engagement environment, the expert defender refines operational activities to manage the operation most effectively. Did the adversary ignore or overlook a deployed lure? The operator may move or change the lure to better encourage engagement. Did the adversary display a new behavior? There may be new opportunities for deception given this new activity. Did the adversary drop a new file? The operator may pause the operation to analyze the file and ensure that future engagement activities are aware of the capabilities of this new piece of malware. As an example, if the adversary is operating in a nonproduction environment for an elicitation operation, the defender may need to make adjustments if a piece of deployed malware has capabilities to capture audio or video. In addition to making refinements during an operation, the expert operator will also make refinement across operations. At the conclusion of each operation, lessons learned, distilled intelligence and other operational outcomes should drive future operations.

This cycle can continue indefinitely; the defender collects and analyzes adversary behavior from an operation, which supplies new engagement opportunities, which yield further data to collect and analyze, and the cycle continues. This active defense allows the defender to be proactive rather than reactive to adversary actions. Whether or not your organization is already collecting and utilizing Cyber Threat Intelligence (CTI), cyber denial, deception, and adversary engagement, and/or other active defense principles the Engage Matrix can be used to begin or mature your strategy.

2.4 Integrating the Engage Matrix into Your Cyber Strategy

While we have explored how the Engage Matrix can be operationalized, we have yet to investigate how the Matrix and by extension, adversary engagement, can be applied within the context of your organization's larger cyber strategy. In this section we briefly zoom out to the 10,000-foot view, to understand the role the Engage Matrix can play in your organization.

Driven by the defender's goals, we designed Engage to complement a traditional cyber defense strategy. The Engagement Goals, Expose, Affect, and Elicit, are not inherently deceptive. And therefore, it is easy to imagine how these goals may already be guiding an organization's security practices. Are you worried about insider threats? Look at the Expose activities to start adding deceptive artifacts around your critical assets. Do you have legacy systems that are no longer able to be updated? Look at the Affect activities to understand how you place decoys to direct adversaries away from these vulnerable systems. Do you feel exhausted by the endless game of CVE whack-a-mole? Look at the Elicit activities to begin to generate your own CTI feed to drive your defense by the adversary's TTPs and not the CVE of the day. Whatever your defensive goals, the Engage Matrix can help you find complementary engagement activities to ensure that if a defense-in-depth approach fails, you remain in control.

While we often consider adversary engagement as a distinct security practice, the most effective and mature implementations are seamlessly integrated into the very culture of an organization. Just as we train our workforce in good cyber hygiene habits, we must train the security community to consider deception as a best practice. In the United States, we have pathologized deception; we consider deception as inherently

negative, sneaky, and dishonest. However, Engage enables defenders to normalize denial and deception activities as routine, essential, and intelligent security practices.

Now that we have explored how adversary engagement can be part of a larger strategic plan, we will zoom back in to understand how to run an engagement operation. The remainder of this handbook will explore the process of adversary engagement and discuss how you can organize a team to begin using adversary engagement as part of your organization's cyber strategy.

3 Elements of an Adversary Engagement Operation

In Section 5, we break down an operation into a 10-Step Process. Leading up to that, it is useful to understand the various elements that go into every operation. Successful adversary engagement operations can be broken down into four components: narrative, environment, monitoring, and analysis. Each component serves specific roles and contributes specific objectives to the stated goal of an operation.

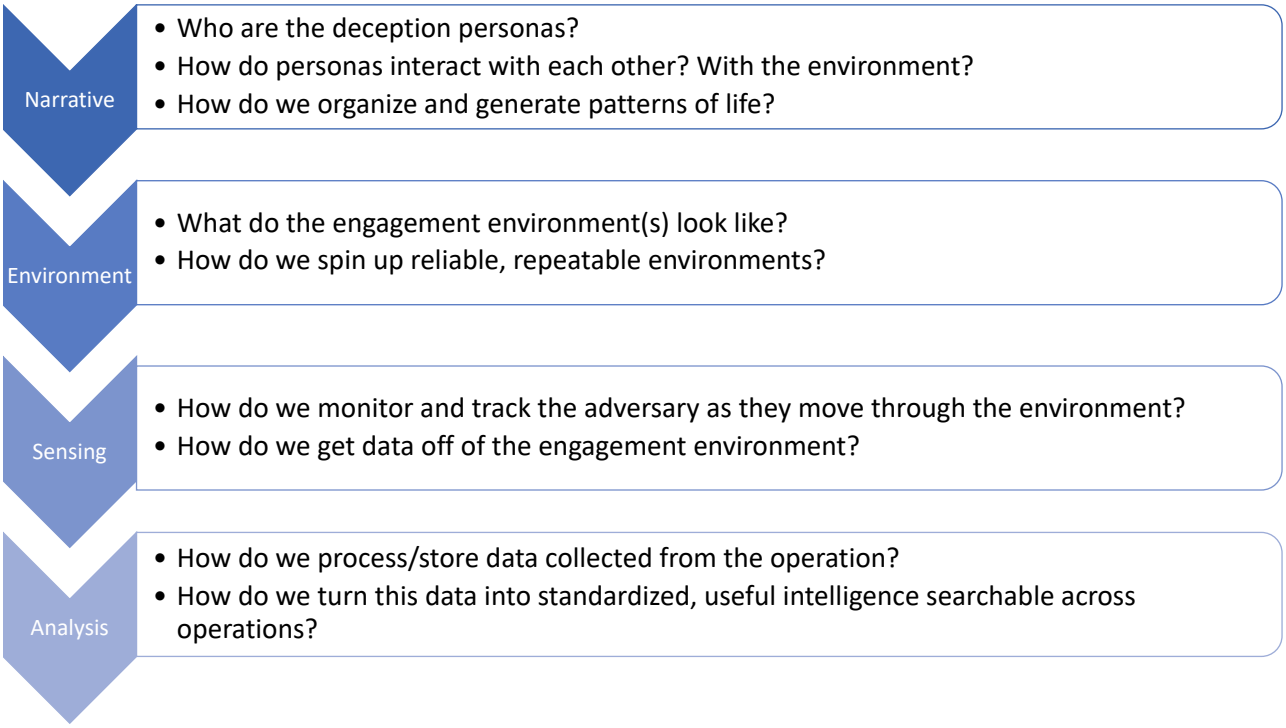

Figure 11: Questions to consider when thinking about adversary engagement operations

3.1 Element 0: Operational Objective

While not strictly an element of an adversary engagement operation, an operational objective is the foundation upon which the other four elements must be built. Due to the creative nature of adversary engagement, it is easy for operations to grow unwieldy or overly complex when the scope of the operation

is not clearly defined though goal setting. The remaining operational elements should always be informed by and aligned with the operational objective.

3.2 Element 1: Narrative

Narrative is the operational element that refers to the overall deception story the operation is intending to portray to the adversary. When creating the narrative, the defender must invariably account for questions such as:

1. What type of network and infrastructure does this victim occupy?
2. What type of documents and communications do any involved personas generate and possess?
3. When are the personas active? How active are they? And what does typical activity look like?

The answers to such questions begin to construct the engagement narrative. There are several ways the engagement narrative is expressed in the engagement environment. One Engage activity that supports the engagement narrative is Pocket Litter. Pocket Litter is data placed on a system to convince the adversary that the system and users are real.

3.3 Element 2: Environment

Environment is the operational element that refers to the setup where the operation will take place. This element includes the hardware/virtual machines, deployed software, networking devices, etc. that make up the environment. This environment may be completely isolated, or it may be integrating within the production environment. When constructing the environment, it is also important to consider how the target adversary will move through the environment. Whether it is a single laptop or a network of computers, it is important to provide pathways through the environment to encourage the adversary to take actions in line with operational goals. For example, if the operational goal is to waste the adversary's resources by maximizing the time spent targeting the fake environment rather than a real environment, then there should be a sufficiently complex network with plenty of appealing pocket litter and lures to encourage the adversary to spend the maximum amount of time in the environment.

It should be noted that there is, of course, a fine line between creating an environment that is so glaringly vulnerable and high value that an adversary immediately recognizes it as a deception and creating an environment so locked-down that the adversary is discouraged from attempting to find and exploit a vulnerability. This distinction is a difficult line to walk and will vary greatly from adversary to adversary. What might scare one adversary away may be the ideal environment for a second adversary. This is yet another example of why setting operational goals and using CTI data to understand your target adversaries is essential. We will discuss this further in Section 5, when we explore how these elements are combined in the process of deception.

3.3.1 Isolated Environments

For some operations the defender may create a synthetic engagement environment separate from the production environment. Such an environment is ideal when a defender is attempting to engage the adversary for information elicitation, as it eliminates the risk to production assets. Isolated environments may be attributed to the defender's organization, or fully unattributed. In the case that the isolated environment is not attributed to the defender's organization, it may provide the defender with opportunities to compare and contrast how the adversary acts in peer environments. Some operations for strategic affects may also benefit from such an isolated environment as it is possible to drive up the resource cost and drive down the value of obtained data. In either case, the defender should use an isolated environment when they are looking to encourage an adversary's activities without risk to real assets.

When creating an isolated engagement environment, the defender must consider what the adversary expects to find. Understanding these expectations is essential to create an environment that will minimize the risk that the adversary detects the deception. By playing into the adversary's cognitive biases, the defender can lower the risk that the deception will be discovered and help reassure the adversary that they are in a legitimate environment. For example, if the adversary expects to find a researcher's computer, including relevant backdated documents may assuage any doubts the adversary may have that this is a real victim. Understanding how much the adversary knows about the environment, and what they expect to find, can also help the defender prioritize their resources. For example, many ransomware samples will not execute if the network does not contain a domain controller. Therefore, setting up a domain controller is an essential task if the goal is to detonate ransomware. As another example, an external threat may not know what the production systems should look like. However, and insider threat will have exquisite knowledge and about what the production network and systems contain. Therefore, much higher levels of realism are required in the latter operation than the former.

3.3.2 Integrated Environments

An alternative approach to cyber deception is to instrument a production environment with deceptive technologies. This technology may include deploying standalone deceptive systems or littering real assets with deceptive artifacts. Deceptive artifacts may include fake credentials, registry keys, introduced vulnerabilities, and/or decoy data. Each of these artifacts can act as a tripwire to produce a high-quality alert as the adversary moves throughout the network. Therefore, integrated environments are often motivated by operational objectives that seek to expose adversaries operating on the network. Additional opportunities exist to use these deployed deceptive assets to impact the adversary's operations by intruding a resource burden, disrupting their ability to operate, or in some way directing the adversary towards some assets or away from others.

When deception is deployed on a production network, there is typically a low false positive rate. Decoys are often data or systems that, while interesting, no one has a business reason to touch. Therefore, even if a non-malicious employee triggers the tripwire, a conversation is warranted to address the overly curious behavior. It is also important to note that this low false positive rate can mean that the time between deployment and alert may be long. It is important to remember that it is the quality of the alert, not the volume of alerts that should be considered when assessing the return on investment of the operation.

When considering the deployment of deceptive artifacts in integrated environments, it is also important to consider if there is already an adversary who is currently in, or was previously in, the environment. As discussed before, meeting the adversary's expectations in terms of which assets are in an environment is crucial to maintaining the deception. If an adversary is already familiar with an environment, they may more easily identify deceptive assets. As a result, deployed tripwires may be rendered ineffective. However, even if the adversary is aware of the deception, they may still suffer a resource cost due to the increased ambiguity of the network.

3.3.3 Relationship Between Pocket Litter and Environment

When constructing the environment, it is essential that the Pocket Litter and environment are both telling a consistent narrative. Inconsistencies can cause the adversary to become suspicious and threaten operational success. When done well, a consistent engagement narrative can be used to establish anchoring bias in an adversary. Anchoring bias is a cognitive bias established when an individual relies heavily on initial information offered, the anchor, to make subsequent decisions. This bias causes all future information that aligns with the anchor to be accepted more readily, while information that does not align with the anchor will be more easily ignored. In addition to establishing anchoring bias, confirmation bias is essential to increasing the odds of operational success. Confirmation bias is the tendency to favor information that supports one's preexisting beliefs or values. Creating a consistent believable environment that supports the creation of this bias may lead the adversary to ignore facts or elements of the environment that are inconsistent or unexpected, as they are contrary to the established confirmation bias.

When the adversary first lands in the environment it is beneficial to leverage anchoring bias and confirmation bias to increase the adversary's tolerance of future weaknesses or discrepancies in the environment. While it is outside the scope of this document to dive more deeply into cognitive biases and how to establish and cultivate them, one way to begin to establish these biases is to ensure that the engagement environment closely resembles a legitimate high value victim to the target adversary. For example, an adversary that targets government entities may be more suspicions if their malware is detonated on a mostly empty computer with no government ties than if they land on a computer filled with documents, communications, and applications that looks to be part of a government network, even if they did not intentionally target this specific victim.

3.3.4 Pocket Litter versus Lures

It is often easy to confuse Pocket Litter and Lures. While both exist in the environment, and often take similar forms, the distinction lies in your intentions. Pocket Litter artifacts are placed on a system to reinforce the legitimacy of the system and/or user. In contrast, Lures are systems and artifacts intended to serve as decoys, breadcrumbs, or bait to elicit a specific response from the adversary. It is important to note that this distinction is related to the engagement designer's intention, and not the adversary's interpretation of the artifacts. While using Cyber Threat Intelligence (CTI) to understand the TA is essential to predict adversary behaviors, it is never possible to know with absolute certainty how an adversary will perceive a given artifact. Did the adversary target a decoy system because of your trail of breadcrumbs or did something else pique their interest? Furthermore, what was intended as Pocket Litter might serve as a Lure and alter the adversary's expected behavior. Whenever we discuss the distinction between Lures and Pocket Litter, we are referring only to the intended purpose of a given artifact.

3.4 Element 3: Monitoring

Operational monitoring is essential to ultimate operational success. Monitoring refers to the series of sensors deployed throughout the environment to record logs, packet capture (PCAP), or other types of collection in the environment. It is essential that the engagement operators have clear visibility in the environment. The number one priority must be to maintain operational security. This safety is maintained by defining a series of Gating Criteria. These criteria represent the conditions that, when met, mean the unnegotiable pause or end of an operation. To capture these criteria, stakeholders create the Rules of Engagement (RoE). These rules represent the actions that are and are not permitted to happen in the environment. Monitoring is essential to ensure that operators are always engaging within the bounds of the RoE. An example rule might be that adversaries are not allowed to use the engagement environment as a platform to target other entities on the Internet. If this is one of the RoE, then it is essential that outside targeting can be quickly identified and stopped. The operators monitoring a given operation must have a clear understanding of the RoE, know how to identify when a rule is violated, and know how to respond when a given rule is violated.

3.5 Element 4: Analysis

Proper analysis is essential to ensure that operational goals are achieved, particularly when these goals relate to intelligence elicitation. Analysis utilizes existing Cyber Threat Intelligence (CTI) and data from the previously mentioned monitoring sensors to categorize, understand, and learn from adversary behavior in the environment. Analytics can be as simple as pushing all data into an ElasticSearch database and using Kibana dashboard to view data. Over time analytics can grow, mature, and scale. The ultimate goal of automated analytics should be to encapsulate the knowledge of an expert cyber threat intelligence analyst, lower the barrier to entry for new analyst, and scale operations to reduce the time required for a given analyst to look at data and make decisions.

4 Adversary Engagement Operational Security (OPSEC)

When planning how adversary engagement may fit into your organization's cyber strategy, it is important to consider OPSEC concerns. First and foremost, you must think through who should be aware of your adversary engagement program. At MITRE, we spent many years keeping a close hold on our use of adversary engagement as part of our InfoSec strategy. In addition to never speaking about our use of adversary engagement publicly, we defined strict guidelines regarding who was involved with operations and how operational data could be stored, shared, and analyzed. It was only after many years of successful operations, that we decided to be more open about our use of adversary engagement. This decision was motivated by a desire to use our experience to help mature and grow the community.

There is a continuum of options between the two extremes of never speaking about adversary engagement and openly presenting at conferences and releasing tools. For example, some organizations may find middle ground by sharing operations with a closed group of trusted partners. No matter where your organization falls on this spectrum, there are pros and cons. By keeping a closed program, you may have

more opportunities to engage unsuspecting adversaries. When adversaries enter an environment without expecting deceptive assets, their natural biases may increase their tolerance for imperfections in the environment and improve the overall believability of the ruse. With an externally known program, you may have the opportunity to deter adversaries through increased ambiguity; if they know deception is in play, they must question everything on your network. As an additional consideration, participating in adversary engagement communities of interest or sharing groups, may expose you to new technologies and research, as well as enable you to compare operational data to complement and enhance your own findings. Whatever you chose, it is important to gather stakeholders and consider your options carefully at the onset of developing your program.

There are additional OPSEC concerns surrounding operational details. It is important to define guidelines around topics such as operational data storage, sharing, and analysis. Additional considerations are required when working through the operational details for insider threat detection programs. While specific OPSEC requirements for operational data are outside of the scope of this document, we hope this brief discussion has highlighted the importance of defining OPSEC guidelines early in the development of an adversary engagement program.

5 The Process of Adversary Engagement

We can combine these four operational elements to outline the process of an adversary engagement operation.

Organizations often fail to adequately plan how and where denial, deception, and adversary engagement will be utilized on their networks. As previously stated, adversary engagement is an iterative, goal driven process, not merely the deployment of a technology stack. Organizations must think critically about what their defensive goals are and how denial, deception, and adversary engagement can be used to drive progress towards these goals. The Engage 10-Step Process helps organizations consider engagement activities as part of just such a process.

The 10-Step Process is particularly important to organizations with limited resources or less mature security programs. Looking at available vendor technologies, it can be tempting to believe that the only solution lies in sophisticated and expensive appliances. However, when you step back and clearly define goals, it may be possible to scope down engagements to tightly align with those goals. Suddenly, even small organizations can get started integrated adversary engagement into their defensive strategies with limited resources!

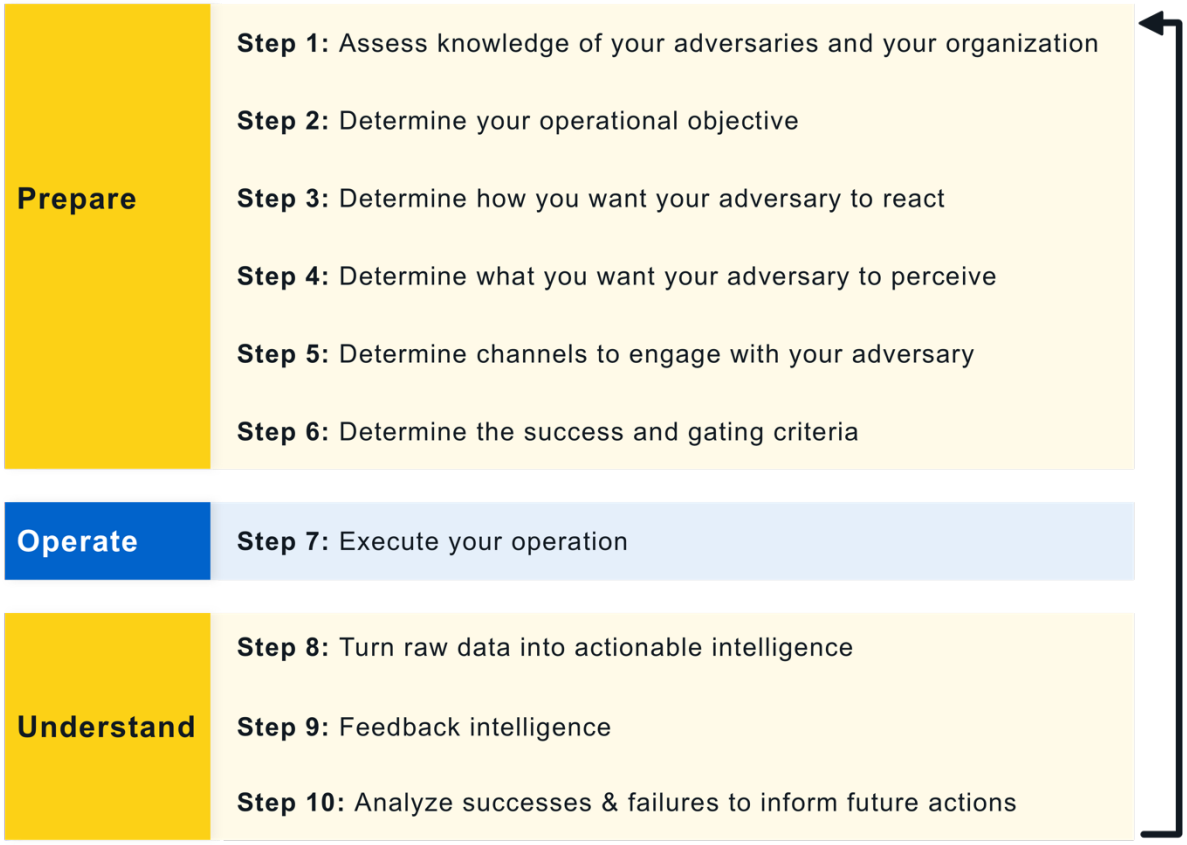

Figure 12: The MITRE Engage 10-Step process of adversary engagement

As seen in Figure 12, the 10 Steps are broken into three categories: Prepare, Operate, and Understand. These three categories mirror the three components of the Engage Matrix. The strategic approaches and activities under the Engage Prepare goal correspond to steps 1-6. The engagement approaches and activities under the Engage Expose, Effect, and Elicit goals correspond to steps 7. And finally, the strategic approaches and activities under the Engage Understand goal correspond to steps 8-10. It is not an accident that the engagement activities fall under the smallest category in the 10-Step Process. While these activities often get the spotlight, strategic actions taken to prepare for and understand the results of your engagement operations are the most important elements of any operation.

This process is not unique to cyber deception. In *The Art and Science of Military Deception*, Barton Whaley wrote a chapter entitled, "The Process of Deception." There, Whaley laid out a ten-step process for creating military deceptions. We refined and repurposed Whaley's work for the cyber domain. We will now explore each step in this process in more detail.

5.1 Step 1: Assess knowledge of your adversaries and your organization

As Sun Tzu said,

> *If you know the enemy and know yourself, you need not fear the result of a hundred battles. If you know yourself but not the enemy, for every victory gained you will also suffer a defeat. If you know neither the enemy nor yourself, you will succumb in every battle.*

This statement rings true in the world of cyber operations. Creating a threat model to understand the risks, vulnerabilities, and strengths of an organization is foundational to planning an effective adversary engagement operation. As part of this model, defenders must identify the organization's critical cyber assets. Likewise, defenders should use Cyber Threat Intelligence to understand the threat landscape.

5.2 Step 2: Determine your operational objective

Armed with an understanding of their own strengths, weaknesses, and threats, the defender should identify their operational objectives. These objectives should reflect the identified priorities. For mature organizations, where adversary engagement is a core pillar of a larger organizational strategy, these objectives should also reflect this larger strategy. These objectives should be specific, measurable actions that will enable the defender to drive progress to larger, more strategic goals. As part of identifying the operational objective, the defender should also identify the Target Adversary, the priority adversary for a given operation. Target Adversaries may be selected for a variety of reasons. A target adversary may represent a threat that historically targets your organization or organizations like yours. The target adversary may represent a gap in your threat intelligence. Or the target adversary may be known to historically utilize TTPs that represent a gap in your current defenses. Regardless of reason, the target adversary allows the defender to focus and prioritize operational activities.

5.3 Step 3: Determine how you want your adversary to react

Now, the defender must decide how they wants the adversary to *react* during the engagement to make progress towards operational objectives. It is important to keep in mind the subtle distinction between what the adversary *think*s and how they *react*. If the defender considers only what they want the adversary to think, it is easy to misjudge how the adversary will react. This mismatch can lead to the adversary reacting unexpectedly (from the defender's vantage point).

5.4 Step 4: Determine what you want your adversary to perceive

Now that the defender has identified how the adversary should react, it is time to consider what the adversary should perceive in the environment to support the operational objective. At this point, the defender must plan for the facts and fictions that will be revealed to and concealed from the adversary. In the book *Cyber Denial, Deception and Counter Deception: A Framework for Supporting Active Cyber Defense*, Kristin E. Heckman, Frank J. Stech, et al., describe this as M-Type and A-Type Deceptions. They define M-Type, or misleading, as deception involving "achieving a reduction of ambiguity, as perceived by the intended target, by building up attractiveness of a wrong alternative." M-Type deception is achieved by revealing facts and fictions and requires careful time and attention to build a series of false signals. In contrast, A-Type, or ambiguity, is deception involving "achieving an increase of ambiguity by surrounding a target with irrelevant information, or to cause confusion based on a lack of certainty." A-Type deception is achieved by concealing facts and fictions.

5.5 Step 5: Determine presentation channels to engage with the adversary

Having identified how the adversary should react and what they should perceive to support the operational objective, the defender must explore the means available to present this effect to the adversary. The engagement environment and engagement narrative will collectively act as the *channels* through which the adversary will engage with these deceptions.

5.6 Step 6: Determine the success and gating criteria

When planning an adversary engagement operation, it is essential that the defender is aware of what both success and failure look like. With every operation there is an element of risk. By setting clear boundaries regarding acceptable versus unacceptable risk, the defender can create clear gating criteria, or points that will trigger the unnegotiable end or pause of an operation. By setting clear Rules of Engagement (RoE) before executing an operation, the defender can avoid any confusion, accidents, or other preventable risks during the operation. Additionally, the defender should identify a clear response protocol in the event of triggering a gating criterion. Every member of the operations team should be aware of these gating criteria and know the appropriate response when triggered. Finally, the successful completion of the operational objective should be clearly understood. If success is poorly defined, it can be easy to waste resources with operations that go on for too long or lose sight of the initial objectives.

5.7 Step 7: Execute the operation

At this point, the operation moves from planning to execution. The defender implements the planned engagement activities and begins to actively engage with the adversary.

Note, the steps that follow are consider Analytical steps. These steps should not be considered only after a gating criterion is reached. Rather, the defender should constantly analyze the engagement operation to iteratively refine implementation details. Throughout the course of the operation, the defender should constantly cycle through planning, execution, and analysis to constantly align progress towards operational objectives.

5.8 Step 8: Turn raw data into actionable intelligence

As the operation proceeds, raw outputs from the engagement should be distilled into actionable intelligence. This process ensures that the results from operational activities are useful to the defender. One key method of distilling intelligence is the use of data analytics. Data analytics allows the defender to map the raw data collected during an operation to the adversary behavior that generated it. To remain useful at scale, automated analytics, such as behavioral analytics, are essential to produce meaningful intelligence. Intelligence produced during this step can be shared inside and outside of the organization, as appropriate, and used to inform threat models and refine operational activities.

5.9 Step 9: Feedback intelligence

The actionable intelligence gained from the operation must be fed back into existing models to inform future decisions. Whenever these models are updated, it is essential to revisit any operational decisions that were made using the old intelligence.

5.10 Step 10: Analyze successes and failures to inform future operations

Whenever a gating criterion is reached, it is essential to analyze the operational successes and failures. This retrospection is an opportunity for the team to review the events of the operation to ensure progress towards operational objectives. It can include a review of the entire operational process from planning, implementation, engagement activity, and impact. In additional to the operation itself, it is an important time to assess the communication and teamwork of the operations team and all contributing stakeholders. While such a review should always occur at the end of an operation, periodic reviews during long-running operations, and especially whenever a gating criterion is reached, are vital to ensure alignment and progress towards the operational objectives.

6 Operational Template

Now that we have explored the high-level process of adversary engagement, we can create an operational template designed to help teams get started running adversary engagement operations. This template is meant to be a guide. When applying this template to your organization, you will need to adapt and refine the structure to fit your organizational needs and resources. It is important to note that, while this template deals is organized around a team, it is very possible to run adversary operations with only a single individual. In that case, consider these various roles and responsibilities as guides to help you identify the various tasks you will need to complete.

6.1 Team Training and Organization

Training and organizing your operations team is an often underestimated challenge when running successful adversary engagement operations. Without a cadre of competent operators, most operations are doomed to fail. Once this cadre is established, the team can function in a steady state, even as individuals may come and go from the team. The challenge comes from building the initial cadre when there is limited expertise or organizational guidelines to lean on to train the team. This challenge is something MITRE struggles with. The following is our current best practices for team training and organization.

6.1.1 Mission Essential Task List (METL) and Mission Essential Tasks (METs)

The METL is a structure by which you identify training requirements and qualifications, establish a team purpose and drive progress to accomplish your goals. The METL concept used by the MITRE Engage team has been adapted from the process used by much larger teams across the U.S. Government.

The METL is made up of a series of Mission Essential Tasks (METs). These tasks are the essential activities that must be completed during the planning, execution, or analysis phases of an engagement operation. Not all adversary engagement activities are considered METs, as some require significantly more specialized skills than others, such as malware analysis and decoder authoring. Therefore, you will never have 100% knowledge of all skills across all the team. However, METs ensure that, where possible,

team members are trained to be proficient in all core competencies. While not everyone is expected to be an expert at every MET, each team member is expected to be able to accomplish each MET. In Table 2, we share the Engage METL Summary used by the Engage team when running self-infection operations for elicitation goals. See the Engage METL Template for a more detailed explanation of this METL template. Adjust this template based on your organizational needs and goals.

Table 2: Template Mission Essential Task List for self-infection elicitation operations

MET	Description
Establish Gating Criteria	Define non-negotiable stops or pauses for an operation. These criteria often form the basis of the Rules of Engagement (RoE) document for the operation
Create Engagement Narrative	Design the persona(s) that will be utilized to interact with the adversary and engagement environment. Design the storyboard of events that will drive the interactions during the operation.
Establish Monitoring System	Identify collection points that will ensure operational safety as outlined in the RoE document
Build Out Victim Windows System	Provision Windows systems consistent with the established engagement narrative
Build Out Victim Linux System	Provision Linux systems consistent with the established engagement narrative
Deploy Monitoring System to Engagement Environment	Buildout the collection system to ensure operational safety
Deploy Persona(s) and Deceptive Assets to Engagement Environment	Buildout the persona(s) and deceptive assets (Pocket Litter and Lures) consistent with the established engagement narrative
Monitor Operational Activity	Observe operational activity within the safety constraints as defined in the RoE
Forensically Investigate Victim Post Operation	Use disk forensics and other IR techniques to gather data from the environment post operation
Analyze Data from Live Operation and Forensic Investigation	Analyze data to distill actionable intelligence
Conduct Open-Source Intelligence (OSINT) Searches Pre/Post Operation	Use OSINT sources to learn about your target adversary's Indicators of Compromise (IOC) and TTPs and to hunt on new IOCs and TTPs you discover during your operation

6.2 Roles and Responsibilities

This section contains a description of the various roles and responsibilities necessary to make a successful adversary engagement operation. Multiple roles may be assigned to one person, or multiple people can

work on a single role. The purpose of defining these roles is to clarify responsibilities. Any changes to these roles during an operation should be communicated to all those involved in the operation.

It is important to note that roles are not meant to be equal in effort over the course of an operation. Some roles are only relevant during a single phase of the operation, whereas others are relevant throughout. Some roles are only responsible for a couple of tasks, whereas others are responsible for many. With this in mind, be sure to staff roles appropriately with individuals who have the necessary bandwidth, seniority, and skillset to maximize likelihood of success of an operation.

6.2.1 Team Lead

The Team Lead is responsible for making decisions that need management approval, require coordination with other entities in the organization, or necessitate visibility across multiple engagement operations. Examples of duties that fall under this role includes signing off on the RoE, approving purchases, and deciding how to handle collateral damage. The Team Lead often holds this role across multiple operations. Their focus is on coordinating adversary engagement operations with other stakeholders in the organization. These stakeholders may include, legal, corporate InfoSec, etc. The Team Lead is the only role that needs to have visibility across multiple operations, to coordinate and deconflict as appropriate.

Responsibilities:

- Approve and expedite procurement of hardware
- Approve and expedite procurement of software
- Author the Rules of Engagement
- Communicate with management about operational risks and the organization's risk tolerance
- Coordinate and deconflict across operations as appropriate

6.2.2 Operational Coordinator

The Operational Coordinator is designated at the start of an operation and is responsible for coordination and ensuring mission success for a single operation.

Responsibilities:

- Facilitate internal communication
- Clearly communicate the Rules of Engagement to all members of the team
- Coordinate and review internal and external reporting
- Capture lessons learned
- Identify areas for improvement
- Ensure that each team member has the tools and training to successfully complete their role
    - Coordinate between the Blue Team and Threat Analyst to ensure monitoring coverage will allow visibility required to uphold the Rules of Engagement

- Coordinate with the Team Lead to deconflict with other simultaneous adversary engagement operations as appropriate
- Ultimate responsibility for operational safety and to initiate appropriate response when gating criteria reached

6.2.3 Threat Analyst

The Threat Analyst is responsible for gathering threat intelligence about the adversary, constructing a persona(s) for the operation, and for attribution analysis of the data obtained from the adversary engagement operation. This may include both during the operation, in a near real time capacity, and postmortem.

Responsibilities:

- Help customize collection system
- Help customize the adversary engagement environment
- Work with the Blue Team to obtain and seed files of interest
- Create Persona Profile of targeted user
- Assist the Blue Team with analyzing C2 communications
- Create new attribution mechanisms based on what is discovered in operation
- Maintain situational awareness of open-source reporting on adversaries of interest
- Author internal and external reporting

6.2.4 System Administrator

The System Administrator is responsible for the initial testing, building, and configuring the computer systems that will be used in the adversary engagement operation. These systems include resources used in the collection system, analysis system, or engagement environment. This will also include troubleshooting any problems and conducting regular maintenance as necessary.

Responsibilities:

- Identify and coordinate procurement of necessary hardware (may require coordination with Team Lead)
- Identify and coordinate procurement of necessary software (may require coordination with Team Lead)
- Build and configure collection system
- Build and configure adversary engagement environment
- Troubleshoot hardware and software problems

- Responsible for ensuring immediate shutdown of the operation when necessary

6.2.5 Operational User(s)

The Operational User(s) are responsible for following the Persona Profile as constructed by the Threat Analyst. This will entail regular use of hosts on the engagement environment to make it look like a live system.

Responsibilities:

- If automated activity tooling is not available, perform activity according to Persona Profile. For long-term or large-scale operations, the Operational User(s) may be a team responsible scheduling who is responsible for activity when.
- If automated activity tooling is available, orchestrate activity across machines according to their respective Persona Profiles

6.2.6 Red Teamer(s)

The Red Team is responsible for testing the collection system, analysis system, and engagement environment before it goes live. They serve to both validate the environment as sufficiently believable and give the team an opportunity to dry-run the operation before interacting with a real adversary.

Responsibilities:

- Work with the operational coordinator to understand the RoE in the context of the red team
- Simulate the adversary on the environment
- Author internal reporting on the actions they take, what they find on the network, and what they find to be suspicious on the network
- Work with the blue team to correlate the actions they took with the actions the Blue Team identified in the analysis system
- Work with the blue team to confirm that the collection system provides adequate visibility to ensure operations remain within the RoE
- Suggest improvements that can be made to the engagement environment to maximize believability and engagement in the context of the operation's goal

6.2.7 Blue Teamer(s)

The Blue Team is responsible for deep technical analysis of the data obtained from the adversary engagement operation. This may include both during the operation, in a near real time capacity, and postmortem.

Responsibilities:

- Help customize collection system
- Help customize adversary engagement environment
- Work with Threat Analyst to obtain and seed files of interest
- Analyze command-and-control (C2) communications
- Extract malware from collections and C2 communications
- Extract malware from system
- Analyze new malware
- Detect collateral damage in Near Real Time (NRT)
- Create new computer network defense (CND) mechanisms, such as signatures, Yara rules, malware signatures, and mitigations, based on what is discovered in operation
- Assist in creating attribution mechanisms based on what is discovered in operation
- Author internal and external reporting

6.2.8 Reverse Engineer

The malware reverse engineer is responsible for working with the Threat Analyst to collect and learn as much as possible about the adversary's TTPs from any malware introduced to the environment either by the operations team or by the adversary and working in coordination with the Blue Team to identify gaps in visibility. This analysis may necessitate the development of protocol decoders and other tools needed to ensure the necessary level of understanding of adversary activity that supports the operational goal. Note, this role can easily be combined with the Threat Analyst if the individual has the appropriate reverse engineering skills.

Responsibilities:

- Reverse engineer any malware that will be used for self-infection
- Report analysis of this malware to Blue Team, System Administrator, Operational User and Threat Analyst to ensure adequate monitoring and an appropriate engagement narrative
- Reverse engineer malware and other files dropped into the environment by the adversary during the operation
- Develop decoders and other tools to allow appropriate visibility into adversary activities
- Assist Threat Analyst in authoring internal and external reporting as it pertains to malware

6.3 Lifecycle Template

In section 5, we explored the process of adversary engagement. In this section, we will build on that process to create template that describes the necessary steps for a successful deception operation. This

template is not intended to be a one-size-fits-all solution. Think of this section as a guide to help you start integrating adversary engagement operations in your organization. This template will show you one Concept-of-Operations (CONOPS) that covers the planning that should be done before beginning an operation, preparation steps that need to be taken before the environment can be constructed, deployment of the environment to the field, daily operations and maintenance, daily monitoring and analysis, and how to gracefully end the operation while capturing lessons learned.

It is important to note that the steps described below are not meant to be uniform in scope and effort. Some may be as simple as filling out a form, others may demand hours of attention and effort. It is also important to note that some steps depend on previous steps being completed, but that is not necessarily the case.

While this section identifies roles and responsibilities associated with each task, it should not be viewed as a complete list, as necessary tasks may vary between operations, and optimal ways of facilitating and coordinating this effort will vary by organization. (Refer to Section 6.2 for further details on the scope of roles and responsibilities.)

For each of the various roles, a single individual should be identified as the responsible party. Even when a responsibility should be worked in coordination with multiple people, the ultimate responsibility for successful completion needs to be assigned to one person. This assignment ensures that that person knows that they are accountable for the specific job to be done and done well. Once again, we will reiterate that this should be considered as only a template for how your team can operationalize the Engage 10 Step Process for adversary engagement. The exact implementation may vary significantly to fit the needs and resources of your organization.

6.3.1 Planning

This phase comprises steps that should be accomplished before starting an operation.

A successful adversary engagement operation relies on a well thought out plan of action. The priorities set during this phase will inform the remainder of operational activities.

6.3.1.1 Build Organizational Threat Model

Role: Team Lead with external stakeholders

Responsibility: Conduct assessments to understand organizational strengths, weaknesses, and critical

Related Step: Assess knowledge of your adversaries and your organization

It is a critical first step to understand the organization's strengths and weaknesses. As part of this threat modeling, it is essential to understand your critical cyber assets. MITRE's Crown Jewels Analysis (CJA)[1] is a process for identifying those critical cyber assets. It is often outside the purview of the adversary engagement team to build the organizational threat model. In this case, the Team Lead should work with the responsible parties to ensure that the operations team has access to the results of this analysis.

[1] https://www.mitre.org/publications/systems-engineering-guide/enterprise-engineering/systems-engineering-for-mission-assurance/crown-jewels-analysis

6.3.1.2 Analyze CTI Data to Understand Threats

Role: Threat Analyst with the Team Lead and external stakeholders

Responsibility: Conduct assessments to understand threat landscape

Related Step: Assess knowledge of your adversaries and your organization

Armed with an understanding of the organizational threat model, the operations team now needs to understand the threat landscape. Using existing CTI data and other intelligence sources, the Threat Analyst can build a model to better understand the threats the organization is facing.

6.3.1.3 Set Operational Objective(s)

Role: Team Lead with the Threat Analyst and external stakeholders

Responsibility: Identify the detection, intelligence elicitation, or adversary management goal(s).

Related Step: Determine your operational objective

Based on the organizational threat model, the threat landscape assessment, and the organization's larger defensive strategy, the operations team must identify operational objectives. It is important to identify the objective(s) up front as this will guide the remainder of the operational setup and activity. Being concrete and specific will lend to a shared understanding across the team. As discussed in Section 3.1, identifying a target adversary is essential to operational success and should inform operational objectives.

It may be the case that outside stakeholders set these objectives. If this decision is made by the operations team, choose this objective in conversation with any partnering organizations. Scope out what observables would indicate success and failure. If working in partnership with general counsel or with law enforcement, understand what kind of information and information handling procedures are necessary for successful legal process (e.g., indictments) as an outcome of the operation.

6.3.1.4 Research the Target Adversary

Role: Threat Analyst in coordination with the Reverse Engineer

Responsibility: Operations are more likely to be successful if they are informed by intelligence. Research the adversary's latest operations, IOCs, and associated TTPs.

Related Step: Determine your operational objective(s)

Understanding the adversary even before this step can help inform goal setting, but once a goal is selected, it is necessary to understand it in the context of an adversary's latest operations and TTPs. This analysis can be accomplished through intel sharing networks, reading open-source reporting, or communicating internally with the organizations Security Operations Center and Incident Response teams.

For example, consider setting up automated analysis pipelines if this group's C2 comms or implants tend to be signaturable. Any amount of work up front like this that can be done before the start of the operation will help the operation run more smoothly. However, setting up such a pipeline risks wasted investment if the target adversary's TTPs change frequently. An informed adversary profile will aid in best determining where to focus upfront analysis.

6.3.1.5 Assign Roles

Role: Operational Coordinator

Responsibility: Assign team roles. The output from this step should be a duty roster reflecting who will be serving in what role for this operation.

Related Step: Logistics setup, not a direct step correlation

As mentioned in Section 6.2, which outlines roles and responsibilities in more detail, roles can be filled by more than one person, and more than one role can be filled by the same person. The most important consideration when assigning roles and responsibilities is whether each individual will have the bandwidth and skills necessary to accomplish all of their responsibilities.

6.3.1.6 Design Engagement Narrative

Role: Threat Analyst, in coordination with the Reverse Engineer, System Administrator and Operational User

Responsibility: Design the engagement narrative.

Related Steps: Determine how you want your adversary to react. Determine what you want your adversary to perceive.

As discussed in Section 3.2, the engagement narrative is the story behind the engagement environment. In other words, this is the story that should elicit a specific set of actions or inactions by the adversary. Storyboarding is the process of creating the deception narrative through a sequence of events, interactions, the persona's pattern of life, etc.

It should include basic details such as what the environment's fictitious purpose is and more specific details such as the infection vector for the environment (e.g., credentials leaked online, phishing email pushed into the environment, etc.). However, aside from essential components, this narrative need not be overly complicated. It will serve as a guiding principle when designing the rest of the operation and may be used to inform certain decision points but convincing the adversary of this story should not be treated as the goal of the operation in and of itself. The narrative is secondary, in support of the goal.

6.3.1.7 Organize Supporting Information

Role: Operational User

Responsibility: Organize information by whether it is fact or fiction and whether the operation must conceal or reveal it.

Related Steps: Determine what you want your adversary to perceive. Determine how you want your adversary to react

Supporting information is what the operation will either conceal from or reveal to the adversary. It comprises all of the information, true or false, that will be used to support the engagement narrative.

Some examples include concealing the fact that the environment is closely monitored, revealing that the environment's IP address falls in the true address space of the organization, revealing fictitious personas of invented end users, and concealing PII of these fictitious end hosts (e.g., social security numbers) in order to keep the barrier to identifying this environment as a hoax sufficiently high. Figure 13 illustrates some examples of this type of information.

	Deception (Reveal)	**Denial (Conceal)**
Facts	**Nonessential information** Things that are validation for the adversary Ex: An email addr of spearfished user	**Essential information** Things you must protect from disclosing to the adversary Ex: Stealthy monitoring of APT
Fictions	**Essential disinformation** What do you want the adversary to believe that is not true Ex: Fake secret product X	**Non disclosable fictions** Things an adversary can potentially see but not use Ex: Encrypted garbage files with no keys

Figure 13: Organization template for the revealed and concealed facts and fictions

6.3.1.8 Design Engagement Environment Network

Role: System Administrator in coordination with the Reverse Engineer

Responsibility: Design deception network diagram

Related Step: Determine channels to engage with your adversary

This network diagram should represent every machine that will be in the engagement environment, how they are networked, and (after the next few steps are completed) where collections occur, where collected data is processed and stored, and where data is fed to a system for generating near real time (NRT) alerts.

This diagram, by completion of the planning phase, should clearly demarcate which of these machines are physical and which are virtual. It should be supplemented with an inventory of what resources are necessary to create this environment.

Do not forget to consider backups when designing the network.

6.3.1.9 Design Personas

Role: Threat Analyst

Responsibility: Design personas for victim and any other relevant end hosts

Related Steps: Determine what you want your adversary to perceive. Determine how you want your adversary to react

Personas serve to support the engagement narrative. In order to make the adversary believe that they are operating on a real network, it is essential to create the illusion that real people are using the network. For any hosts that may fall under scrutiny, assign a work function and some basic personality characteristics to the fictitious person that would be using it. These will be leveraged later to inform what sorts of Pocket Litter appear on the machine and what actions are taken during the op to make the user(s) look active.

6.3.1.10 Determine Gating Criteria

Role: Operation Coordinator in collaboration with Blue Team, Threat Analyst, and external stakeholders

Responsibility: Determine the critical conditions that should guide the operation

Related Step: Determine the gating criteria

The Gating Criteria are the set of events that mean the unnegotiable stop or pause in operational activities. Sometimes, these events include the successful completion of the agreed upon strategic goals. Other times, these events may signify the operation has reached a hard stop. This is often because future operational safety cannot be guaranteed. Or events have occurred that outweigh the agreed upon acceptable risk. Finally, it may just be that if the adversary operates any longer, they may learn something the defender doesn't want them to know or the defender is no longer getting value from the engagement. Multiple parties from the technical operations, threat intel, legal, and management perspectives should be included when Defining Gating Criteria.

6.3.1.11 Write the Rules of Engagement

Role: Team Lead

Responsibility: Write the Rules of Engagement for the operation.

Related Step: Determine the gating criteria

Rules of engagement must be determined for the operation based on the defined gating criteria. Specifically, there must be bright red lines around what is considered timely for detection of collateral damage and what to do when different types of collateral damage are detected. The RoE ought to also include any end conditions (e.g., intelligence goal achieved, adversary has ceased interaction with the environment, adversary is performing validity checks that may tip them off to the inauthenticity of the environment).

Another consideration to address in the RoE is the OPSEC requirements for the operation. How will the operations team send potentially malicious links and IP addresses within the team or more widely across the organization for information sharing? Where is it appropriate to look at potentially malicious files or investigate potentially malicious domains or IPs? How will the team appropriately mark and store this data during and after analysis? If applicable, how will the team feed gathered IOCs into existing threat intelligence feeds to prevent accidents outside the engagement environment. Clear rules regarding this data handling and sharing are imperative to reduce the risk of spills or accidently compromises outside of

the engagement environment. Coordinating the RoE with existing InfoSec teams is also essential to ensure appropriate communication and understanding across the organization.

The RoE is meant to capture foreseeable bright red lines as well as guidelines for how to act under unforeseen circumstances. It is not meant to be exhaustive, and the clarity and legibility of the document should be prioritized.

6.3.1.12 Choose Points of Collection

Role: Operational Coordinator in coordination with the Reverse Engineer, Threat Analyst, Blue Team, System Administrator, and Red Team

Responsibility: Select points of collection for environment monitoring systems (select these based off what the intel output will look like/tools that support the operational goal).

Related Step: Determine the gating criteria

This collection may include both host-based and network-based sensors, traditional security operations tools, and custom capability for finer-grained visibility. A common instinct is to collect everything possible. However, collection, indexing, and search systems are finite resources that typically cannot handle the amount of data and noise generated by a "collect all the things" methodology.

Note, the selection of collection points is nontrivial. Therefore, significant resources and attention should be placed on ensuring that the selected collection points enable adequate monitoring of the engagement environment. What warrants sufficient collection will be highly dependent on operational goals, the RoE, any specific malware or APTs, etc. Therefore, this step should be implemented in coordination with many of the various roles on the operation team. Additionally, as described in Section 6.2.6, the red team can be used to test the environment to ensure sufficient instrumentation.

6.3.1.13 Choose Medium for Near Real Time Notifications

Role: Blue Team in coordination with the System Administrator

Responsibility: Select how to receive notifications about activity in the environment.

Related Step: Determine the gating criteria

Examples include text message, email, use of a messaging service such as Slack or Discord, etc. What is most important is that those responsible for monitoring alerts are comfortable with the medium, any rate limiting will not result in a lack of visibility (relevant if using a messaging service such as Slack or Discord), and that the service meets OPSEC requirements (e.g., ability to turn off content prefetching and/or escape potentially malicious links, end-to-end encryption, etc.). For those organizations with a dedicated Security Operation Center (SOC) and existing Security Information and Event Management (SIEM) system, sending alerts to existing capabilities may be a viable option and enable the organization's existing threat intelligence feeds to be informed of new IOCs to aid in detection elsewhere in the network.

Even if the blue team chooses not to share existing SIEM infrastructure, it is always a good idea to inform existing CTI feeds with IOCs and other alerts from the operation. This information sharing could aid in the

prevention of accidents in the event that someone on the operations team accidently contaminates a clean corporate network with potentially malicious data from an engagement environment.

6.3.1.14 Get Necessary Approvals

Role: Team Lead

Responsibility: Seek necessary approvals within the organization

Related Step: Determine the gating criteria

Make sure to clearly communicate the actions the operations team intend to take as well as any risk associated with the operation to stakeholders in management and legal. Following these conversations, it may be necessary to revise or amend the RoE.

6.3.1.15 Select Collection and Analysis System(s)

Role: Blue Team in coordination with the System Administrator

Responsibility: Select which tools to use at the designated collection points

Related Step: Determine the gating criteria

The Collection System is the set of systems used to gather artifacts and other data from an operation to monitor the engagement. The Analysis System is the set of systems used to store and understand these artifacts and data. Building and maintaining these systems is essential to ensure operations operate within the Rules of Engagement.

The Blue Team should communicate their operational requirements and expectations of what kind of data they will receive. The System Administrator will later be responsible for implementing or installing the discussed solution at each collection point, so they should be aware of any system requirements, constraints, or compatibility issues that may arise.

In the event that there are competing needs in the operation (e.g., familiarity with a certain toolchain vs budget to purchase software), the Operational Coordinator will make the final determination on what is essential for operational success.

6.3.1.16 Gather Deceptive Artifacts

Role: Threat Analyst in coordination with the Reverse Engineer

Responsibility: Gather Pocket Litter, Lures, and other deceptive artifacts in support of the engagement narrative and any personas.

Related Steps: Determine what you want your adversary to perceive. Determine how you want your adversary to react.

The engagement narrative should drive the creation of the deceptive artifacts, particularly Lures. Lures are intended to elicit a specific reaction from the adversary and therefore often require careful planning. Pocket

Litter, on the other hand, exists only to reinforce the engagement narrative. Therefore, Pocket Litter often allows for more flexibility and lower fidelity than Lures.

Gathering artifacts, particularly Pocket Litter, is a step that can bloat to fill as much time as is available. Therefore, it is important to rely on the target adversary profile developed in Section 6.3.1.4. Important high-level considerations include the amount of scrutiny that a given file or process is expected to fall under, whether the adversary's operators usually exfiltrate data indiscriminately in bulk or specific documents relevant to their objectives, etc. This understanding of the target adversary will inform tradeoffs in quality vs quantity, and how much time is necessary to spend on litter in general.

This step lends itself to automation, especially for gathering/generating Pocket Litter. For human language content, such as documents or emails, consider generating content using open-source language models such as GPT-2, or by scraping it from the web. For software, make sure to gather necessary installers and licenses that match the software/host versions that support the narrative. Automation may be more challenging for Lures, as these assets likely need to be specific content relating to a specific topic.

The Threat Analyst will be responsible for selecting deceptive artifacts that will be enticing to the adversary. The Blue Team should be familiar with what assets will be on the machine in order to better identify rogue files and processes on the machine and adversary exfiltration.

This step of gathering Pocket Litter also offers opportunities to establish plausible deniability in the event that the environment is compromised by a security researcher. For example, writing a detailed explanation of the operation, and leaving it in the environment in an encrypted zip file with a sufficiently complex password could be used to prove to a researcher that what appeared to be a security vulnerability in the organization's network was an intentional engagement environment. By providing the password to any security researcher who exfiltrates the encrypted zip, the intentional deception can be revealed. Similar opportunities including creating uncrackable user passwords with phrases such as "honeytoken" or "deception operation" which can be used to reveal post-operation that the environment was part of an intentional deception operation.

6.3.1.17 Acquire Engagement Environment Infrastructure

Role: System Administrators

Responsibility: Acquire any additional machines or infrastructure (e.g., VPS) necessary to support the engagement

Related Step: Determine the channels to engage with your adversary

Consider the engagement narrative and the target both the volume of data the operation anticipates handling as well as the rate at which it will need to process data in order to hit the NRT response targets. Employ necessary OPSEC considerations when purchasing infrastructure.

6.3.2 Deploying

This phase comprises steps to set up a deception environment such that it's ready to go live, and then kicking off the operation.

6.3.2.1 Provision the Network

Role: System Administrator

Responsibility: Provision all hosts on the network and ensure they are connected appropriately.

Related Steps: Determine what you want your adversary to perceive. Determine how you want your adversary to react. Determine channels to engage with your adversary

During this step, the deception network should transform from a diagram to a collection of real (or virtual) machines. Any deviations from the original diagram should be avoided when possible and communicated to the Operational Coordinator and Team Lead when unavoidable.

6.3.2.2 Deploy Deceptive Artifacts

Role: Threat Analyst in coordination with the Reverse Engineer, System Administrator and Operational User

Responsibility: Deploy Pocket Litter, Lures, and other deceptive artifacts onto the deception network.

Related Steps: Determine what you want your adversary to perceive. Determine how you want your adversary to react. Determine channels to engage with your adversary.

This step includes both setting content and software on machines in locations that will be observable to the adversary in support of the engagement narrative, as well as with timestamps that support the engagement narrative.

The System Administrator is responsible for the action of placing or installing artifacts while the Threat Analyst is responsible for advising where and when artifacts should be placed to entice the adversary.

6.3.2.3 Configure Collection System(s)

Role: System Administrator in coordination with the Blue Team

Responsibility: Install, configure, and confirm the functionality of the collection systems(s)

Related Step: Determine the gating criteria

Part of this set up should include verifying that the collections happen as expected through simple test cases. At this point, the Blue Team should consult in order to ensure that the collections meet their expectations and needs.

Part of this set up may also include a fixed duration of operational activity in order to establish a behavior baseline within analytic capabilities that make use of such bespoke network monitoring.

6.3.2.4 Snapshot Baseline

Role: System Administrator

Responsibility: Snapshot and store each machine's baseline state as a point of comparison for analysis and in case there is the need to redeploy.

Related Step: Determine the gating criteria

This step should happen after Pocket Litter has been deployed on each of the machines and collection systems have been configured on end hosts. Backups of each machine's baseline should be stored off of the deception network and should be available on short notice should the need to redeploy the environment arise.

6.3.2.5 Confirm Understanding of the Rules of Engagement

Role: Operational Coordinator in coordination with the Team Lead

Responsibility: Make sure every member of the team is familiar with the final revision of the RoE.

Related Step: Determine the gating criteria

Assuming all team members understand the final version of the RoE without confirmation can cause breaches of the RoE during the course of an operation. This should be avoided. It is the Operational Coordinator's responsibility to confirm that every member of the team is on the same page with allowable actions, necessary actions, and forbidden actions during the operation.

At this point, it is also important to ensure that the team has a well-defined and well understood escalation procedure in the event that a team member feels that the operation has, or is about to, violate a RoE. In most cases, this escalation procedure should outline how the Blue Team or Threat Analyst should engage the Operations Coordinator.

6.3.2.6 Test the Environment

Role: Red Team in coordination with the Blue Team

Responsibility: Perform a red team assessment of the environment.

Related Step: Logistics setup, not a direct step correlation

Budget and time allowing, a Red Team exercise is ideal for checking the believability of the environment as well as the configuration of collections systems and NRT notification systems.

Ideally, the Red Team and Blue Team operate in total isolation of each other until the exercise concludes. Optionally, the Red team may not be informed that the environment is a deception environment, and the Blue team may not be informed that the Red Team is a simulated adversary.

Following the exercise, both teams should sync to map actions taken by Red with actions observed by Blue. If there are blind spots in collections and notifications, or if Red can compromise the environment without Blue's knowledge, each aspect of the operation should be reevaluated before the operation proceeds.

Time permitting, Red's feedback should be used to tweak the environment such that it is more believable and collection system configurations such that visibility is improved. All changes must be properly documented.

6.3.2.7 Reset to Baseline and Deploy

Role: System Administrator

Responsibility: Reset the environment to the latest baseline and start the operation.

Related Step: Execute your operation

Before starting the operation, ensure that the network is in the desired state. Time allowing, this is also a good time to perform one final check that no personally identifiable information exists anywhere on the network.

The start of the operation will look different depending on the chosen infection vector. The final step of deploying the operation is to open the environment to adversary activity, through running malware, connecting it to the internet, putting it in the path of ongoing adversary activity, etc.

6.3.3 Operating

This phase comprises steps that must be taken during an operation in order to maintain it and keep it believable.

6.3.3.1 Generate User Activity

Role: Operational User(s)

Responsibility: To keep the environment looking like it's being actively used, user behavior should be performed or generated on end hosts.

Related Step: Execute your operation

This may take the form of a rotating shift acting as a persona while sitting at the physical host that is part of the operation, or it may take the form of orchestrating automated behaviors across hosts if such automation is available.

6.3.3.2 Monitor and Respond to Near Real Time Alerts

Role: Blue Team

Responsibility: Track activity on the network. Detect and respond to adversary activity.

Related Step: Execute your operation

As is mentioned in more detail in Section 3, the Blue Team is responsible for maintaining situational awareness of all NRT alerts, and especially when there is novel or unexplained activity. The bulk of this will likely be false positives, but understanding what false positives usually look like will allow further tuning of the collection system and enable the Blue Team to better distinguish when alerts look abnormal and demand more attention.

The Blue Team is also responsible for taking any courses of action outlined in the RoE in response to unexplained activity within the allowable timeframe.

6.3.4 Analyzing

This phase comprises steps that must be taken during an operation to maintain situational awareness of adversary behavior in the environment.

6.3.4.1 Maintaining Operational Safety

Role: Blue Team

Responsibility: Analyze data in real time to ensure RoE are upheld

Related Steps: Turn raw data into actionable intelligence

The Blue Team is responsible to ensure that the RoE not violated at any point during the operation. If at any point the blue team feels the operation is at risk of violating a given rule, they are responsible for suspending the operation and following the appropriate escalation procedure before resuming the operation. This responsibility is critical in ensuring operational safety.

6.3.4.2 Analysis to Achieve Desired Operational Outcomes

Role: Threat Analyst in collaboration with the Operational User(s)

Responsibility: Analyze data in real time and postmortem to accomplish operational goals (for example, capturing new TTPs, IOCs, etc.)

Related Steps: Turn raw data into actionable intelligence, Feedback intelligence

It is the Threat Analyst's responsibility to ensure that the operation is continuing to move towards fulfilling operational goals. Depending on adversary activity, this responsibility may require collaboration with the operational user(s) to further engage with the adversary throughout the operation. Additionally, it is the threat analyst's job to ensure that adequate data is collected to accomplish said goals.

6.3.4.3 Reverse Engineer Malware

Role: Reverse Engineer

Responsibility: Identify malware, TTPs, and novel IOCs.

Related Steps: Turn raw data into actionable intelligence, Feedback intelligence

The initial stage malware may need to be analyzed to determine how to successfully execute it in the deception environment. Some malware may have configuration or software dependencies, such as Microsoft Word or .NET runtime. Sophisticated adversaries, particularly targeted threats, integrate environmental checks into their malware.

If the adversary drops later-stage malware in the engagement environment, reverse engineering it may be necessary to maintain situational awareness of what the adversary is doing or might do in the environment.

The amount of time and resources dedicated to reverse engineering should scale according to the complexity of the malware and the organization's risk tolerance. This effort can be continued after the operation ends, provided a sample of the malware is preserved.

6.3.4.4 Decode Command and Control Communications

Role: Reverse Engineer

Responsibility: Decode malware comms in order to be able to interpret them in situ.

Related Steps: Turn raw data into actionable intelligence, Feedback intelligence

Writing a decoder for a given malware sample or family will greatly increase insight into what actions the adversary is taking during an operation and may reveal new IOCs or TTPs.

If the adversary is using strong, properly implemented encryption, this may not be possible without public (e.g. sslbump) or custom (e.g. API calls) monitoring or hooking capabilities.

6.3.4.5 Give Context to Alerts

Role: Threat Analyst in coordination with the Blue Team

Responsibility: Support NRT response by giving context to alerts.

Related Steps: Turn raw data into actionable intelligence, Feedback intelligence

The Threat Analyst should maintain situational awareness of activity on the network in a deeper, more holistic sense than is allowed by the NRT notification system. They should be examining logs, network traffic capture, etc., either raw or in whatever analytic framework is available. They should be prepared to answer questions from the Blue Team as alerts are triggered.

6.3.5 Ending and Improving

This phase marks the end of an operation and suggests points to consider for reflection and improvement of future operations.

6.3.5.1 Pull the Plug

Role: Operational Coordinator in coordination with the Team Lead, Blue Team, and Threat Analyst

Responsibility: End the operation upon reaching a success state, end state, or unanticipated state per the Gating Criteria

Related Steps: Turn raw data into actionable intelligence, Feedback intelligence, Analyze successes and failures to inform future actions

While the Blue Team may end the operation in the event that defined success or end states are met, per the RoE, should the unexpected happen in the course of the operation (as is often the case), it is the Operational Coordinator's responsibility to decide whether continuing or ending the operation is in the proper spirit of the RoE, per legal and management's risk tolerance.

Ending the operation may be automated through programming machines to shut down or disconnect from the network if certain conditions are detected. It can also be as manual as pulling the network cable out of a victim machine if adversarial collateral damage is detected. In any case, all paths to the operation's network should be cut if it is time to end it.

Other relevant considerations at this step are the impression that the adversary should be left with. If the operation is meant to be covert, and the adversary absolutely cannot discover the network to be a ruse, then it may be better to end the operation gracefully (e.g., simulate the behavior of a hardware failure, server crash, or incident response) if the adversary seems to be growing suspicious of the environment. If it doesn't matter whether the adversary's suspicions are raised, it may make sense to keep the operation going until they lose interest or attempt to perform actions outside of the RoE.

These sorts of considerations should be kept in mind while writing the RoE during the planning phase, but since all possible outcomes of an operation cannot possibly be thought through ahead of time, they are worth reiterating here.

6.3.5.2 Review Objective in the Context of Operational Data

Role: Threat Analyst in coordination with the Blue Team

Responsibility: Assess whether the objective of the operation was achieved based on observables coming out of the operation.

Related Steps: Turn raw data into actionable intelligence, Feedback intelligence, Analyze successes and failures to inform future actions

Based on the team's visibility into the environment, it should produce a "yes"/ "no" answer to the question of whether the operation's objective was achieved.

However, it is possible that the answer will look something more like "it's complicated" or "we're not sure." Should this be the case, it is especially important over the course of the next few steps to develop an understanding why this is the case, such that future goals can be chosen, and operations can be designed that will result in more clear outcomes.

6.3.5.3 Conduct After-Action Review

Role: Operational Coordinator in coordination with the Team Leader, Threat Analyst and Blue Team

Responsibility: Assess what lead up to the end of the operation, regardless of whether the objective was achieved.

Related Step: Analyze successes and failures to inform future actions

If the operations team was successful in achieving the operation's goal, take time to review the implications of the new intelligence, or the indicators of cost-imposition. It is very difficult to measure the effectiveness of an operation. What is often more feasible is to show measures of work. Some exemplar success metrics to show these measures of work can include total number of C2 packets sent to the engagement environment, length of time of meaningful engagement, megabytes exfiltrated, novel tools collected, domains/IP addresses/signatures collected, new TTPs identified. Figure out how this value can be communicated back to the rest of the organization, or other partnering organizations.

If the adversary engaged with the environment, but eventually lost interest, review adversary behavior leading up to them abandoning the environment. Were they looking for certain kinds of information or hosts? This can contribute to design considerations of future operations.

If the adversary identified the environment as being a deception environment, review their behavior leading up to, and following, this determination. The more that is understood about how the adversary checks for the legitimacy of the environment, the more authentic looking the environment can be made in future operations. The more the operations team understand about how the adversary acts when they discover themselves subject to a deception, the more intelligence that will have to be fed into counter-deception TTPs.

6.3.5.4 Identify Bottlenecks & Areas for Improvement

Role: Operational Coordinator

Responsibility: Given the understanding of whether or not the operation's goal was achieved and what led to that outcome, identify ways in which the operation could have been better.

Related Step: Analyze successes and failures to inform future actions

This assessment will be unique to the context of each team and operation. There are areas that may be worth considering including the division of responsibilities and team communication, environment features that need to be more believable, arduous or repetitive tasks that may benefit from scripting or automation, administrative paperwork that can be turned into templates for future operations, etc. Note that this conversation should intersect both technical and human domains.

It is very likely that more areas for improvement will be identified than can realistically be tackled. It is recommended that improvements be triaged based on their potential benefit in the context of available effort to implement them.

6.3.5.5 Capture and Communicate Lessons Learned

Role: Operational Coordinator in coordination with the Team Lead

Responsibility: Document outcomes of the operation, including process improvements, intelligence outcomes, and value-added back to the organization's defensive posture. Communicate relevant details up the organization's management chain.

Related Step: Analyze successes and failures to inform future actions

This can take many forms, but the common goal is to make sure that this information exists is captured on paper in addition to in the minds of the team. Lessons learned should outlive any individual's tenure on the team. This may include sharing IOCs with intel sharing partners or briefing the CISO on targeted threats that the organization is currently facing.

7 Future Considerations

The current adversary engagement concept covers operations from simple, standalone computers on a consumer grade Internet connection, to larger, bespoke networks mimicking more involved home networks, small office networks, or enterprise level networks. Future work may consider automating the polymorphic design and provisioning of large networks, automating persona generation, automating natural language content generation, automating behaviors performed by operational users during deception operations, and more.

Outside of the direct scope of running adversary engagement operations, future work may also investigate the weaknesses inherent in one's deception environments in order to identify areas where the environment may be hardened, such as the hiding of virtualization or eliminating its use. Developing tradecraft for identifying deception environments is also useful in the context of countering deception during offensive cyber operations.

8 Acknowledgements

We would like to recognize and thank the following individuals who took the time to peer review this document, or a previous iteration of this document, providing content, feedback, and many edits:

Bill Hill, Director of MITRE Information Security and CISO

Dan Perret, MITRE Principle Cyber Security Engineer

Andrew Sovern, Adversary Engagement Engineer

Dr. Frank Stech, Principal Behavioral Scientist

Adam Pennington, MITRE ATT&CK Lead

Additionally, we would like to thank the members of the MITRE Engage Team

Maretta Morovitz

Gabby Raymond

Stan Barr

Dylan Hoffmann

Bryan Quinn

Brittany Tracy

Leslie Anderson

Olivia Hayward