



INSTITUTE FOR DEFENSE ANALYSES

An Ontology for the Embedded System TTP Matrix

Rachel Kuzio de Naray, Project Leader
Brian A. Haugh
Steven P. Wartik

With contributions from Noah T. Plymale

March 2022

IDA Document P-32935

Log: H 2022-000004

Approved for public release; distribution is unlimited.
Cleared for public release by the DoD Office of
Prepublication Review 4/20/2022, Case 22-S-1548



The Institute for Defense Analyses is a nonprofit corporation that operates three Federally Funded Research and Development Centers. Its mission is to answer the most challenging U.S. security and science policy questions with objective analysis, leveraging extraordinary scientific, technical, and analytic expertise.

About This Publication

This work was conducted by the IDA Systems and Analyses Center under contract HQ0034-14-D-0001, Project AX-1-3100, Technical Analysis for the Director, Developmental Test, Evaluation and Assessments. The views, opinions, and findings should not be construed as representing the official position of either the Department of Defense or the sponsoring organization.

Approved for public release; distribution is unlimited.

Acknowledgments

The authors would like to thank the IDA committee, Dr. Stephen M. Ouellette (chair), Dr. John S. Hong, Dr. Allyson M. Buytendyk and Dr. Christine R. Bucher, for providing technical review of this effort.

For More Information

John S. Hong, Project Leader
jhong@ida.org, 703-845-2564

Stephen M. Ouellette, Director, SED
souellet@ida.org, 703-845-2443

Copyright Notice

© 2022 Institute for Defense Analyses
730 East Glebe Road, Alexandria, Virginia 22305 • (703) 845-2000.

This material may be reproduced by or for the U.S. Government pursuant to the copyright license under the clause at DFARS 252.227-7013 (a)(16) [Jun 2013].

Rigorous Analysis | Trusted Expertise | Service to the Nation

INSTITUTE FOR DEFENSE ANALYSES

IDA Document P-32935

**An Ontology for the
Embedded System TTP Matrix**

Rachel Kuzio de Naray, Project Leader

Brian Haugh

Steven Wartik

With contributions from Noah T. Plymale

Executive Summary

Embedded systems are crucial to Department of Defense (DoD) operations.¹ They are parts of weapons, transportation, and observation platforms (including satellites and autonomous drones), all critical to DoD operations. They perform key functions in direct control of platforms as well as in their communications. Therefore, embedded systems must be designed and developed to resist and easily recover from cyberattacks, which could threaten their control and operations resulting in loss of information, equipment, and lives. Assessments of the cyber resilience of DoD embedded systems are therefore an essential part of their design and development.

Cyberattacks on Information Technology (IT) systems and internet-facing systems have received more attention than embedded systems cyberattacks, largely due to IT systems' larger attack surfaces through connections to wide area networks (WANs). However, even isolated embedded systems face cyber threats through routes such as software updates, communications systems, and connections to electronic maintenance systems. Hence there is a need for modeling potential cyberattacks on embedded systems.

To address this need, the MITRE Corporation developed the Embedded System Tactics, Techniques and Procedures (TTP) Matrix (ESTM). ESTM is a matrix of categories of tactical objectives (tactics) and cyber techniques for achieving those objectives as applied to embedded systems. It adapts the approach of the MITRE ATT&CK[®] matrices to embedded systems.²

This matrix of tactics and techniques has now been transformed into a formal ontology³ to support sharing information and reasoning about cyberattacks on embedded systems. The ESTM Ontology provides a formal logic-based model of the relationships between ESTM tactics and techniques using the Web Ontology Language (OWL) in a format consistent with related ontologies for cyberattacks on enterprise and mobile systems. It is generated automatically from a spreadsheet-based version of ESTM.

¹ An embedded system is a microprocessor-based computer hardware system with software designed to perform a dedicated function, either as an independent system or a part of a large system. At the core is an integrated circuit designed to carry out computation for real-time operations.

² The MITRE Corporation developed MITRE ATT&CK[®] as a knowledge base of cyber adversary behavior and taxonomy for adversarial actions against enterprise and mobile networks.

³ An ontology is a set of concepts in a subject area or domain that categorizes the concepts, shows their properties, and states relationships between them.

The ESTM Ontology could be used in cyber risk assessments (CRAs), such as Cyber Table Top (CTT) exercises, to model possible cyberattacks on embedded systems to assess their vulnerabilities and devise mitigation strategies. During an exercise, CTT participants could query this ontology to identify attack techniques that might be useful. After an exercise, the sanitized results could be stored in an ontology-compliant knowledge base for future access and analyses.

Contents

1.	Introduction	1-1
A.	Background	1-2
1.	The MITRE ATT&CK Knowledge Base.....	1-2
2.	Ontological Representation of ATT&CK	1-5
3.	PIVOT	1-6
B.	Structure of this Paper	1-8
2.	ESTM Ontology Modeling Approach	2-1
A.	Ontology Structure	2-1
B.	Tactics and Techniques in Relation to Imported Ontologies	2-1
C.	Cyber Embedded System Tactics.....	2-4
D.	Cyber Embedded System Techniques.....	2-5
E.	Relationship between Tactics and Techniques.....	2-6
F.	Annotations	2-6
3.	Generating the ESTM Ontology.....	3-1
4.	Using the ESTM Ontology.....	4-1
5.	Summary.....	5-1
	Appendix A. References	A-1
	Appendix B. Acronyms and Abbreviations	B-1

Figures

Figure 1-1. ATT&CK Matrix for Containers	1-3
Figure 1-2. ATT&CK Matrices	1-4
Figure 1-3. MAMO Dependencies	1-6
Figure 2-1. ESTM Ontology import structure	2-1
Figure 2-2. Objective class hierarchy	2-2
Figure 2-3. Cyber Objective in ESTM Ontology.....	2-3
Figure 2-4. Cyber Attack class hierarchy	2-4
Figure 2-5. Cyber Embedded Technique as Cyber Attack subclass.....	2-4
Figure 4-1. OACRA V1.0 relationship to MAMO	4-1
Figure 4-2. OACRA V1.1 relationship to MAMO and ESTM Ontology	4-1
Figure 4-3. OACRA V1.1 larger import structure.....	4-2

Tables

Table 1-1. Namespaces used in this paper	1-9
Table 2-1. ESTM and ATT&CK Tactics.....	2-5
Table 2-2. ESTM Ontology Technique Counts	2-5
Table 3-1. Input Spreadsheet Content.....	3-1

1. Introduction

In information science, an ontology expresses, for some area of interest, concepts, their properties, and their relationships (Buytendyk 2020). This document describes an ontology that models cyberattacks on embedded systems. An embedded system is a computer system embedded within a larger system, either mechanical or electronic, to perform a dedicated function.⁴ Embedded systems are crucial to Department of Defense (DoD) operations. They are part of weapons, transportation, and observation platforms (including satellites and autonomous drones), all critical to DoD operations. Some examples of embedded systems in DoD platforms include automatic braking systems for ground vehicles, flight control systems of aircraft, guidance and seeking sensor systems for missiles, and self-defense systems for aircraft carriers and combat ships.

Cyberattacks on Information Technology (IT) systems and internet-facing systems have received more attention than cyberattacks on embedded systems. This situation is perhaps because non-embedded systems are more likely to have a direct connection to a wide area network (WAN) and therefore present an easier-to-find, easier-to-access attack surface. By contrast, embedded systems operate (if architected according to best practices) on a local area network (LAN) that connects to other networks and external systems only through a carefully configured and constantly monitored firewall.

However, there have been successful cyberattacks (and hacking demonstrations) on embedded systems. Examples include:

- The DHS hack of a Boeing 757 in 2016.⁵
- Hack of an F-15 fighter jet in August 2019.⁶
- Chinese hack of a Tesla Model X, turning on the brakes remotely, etc.⁷

⁴ Adapted from https://en.wikipedia.org/wiki/Embedded_system

⁵ See <https://www.aviationtoday.com/2017/11/08/boeing-757-testing-shows-airplanes-vulnerable-hacking-dhs-says/>

⁶ See <https://nationalinterest.org/blog/buzz/hacked-how-f-15-fighter-can-be-taken-own-cyber-attack-132042>

⁷ See <https://www.usatoday.com/story/tech/2017/07/28/chinese-group-hacks-tesla-second-year-row/518430001/>

- Top 25 auto cybersecurity hacks.⁸

Attacks on IT systems and web servers have the potential to cause considerable inconvenience, but are seldom life-threatening (although their indirect effects may be, as when a hospital IT system shutdown results in patient death). Embedded systems control active platforms and equipment whose failure can directly cause the loss of life and damage or incapacitation of equipment. Cyberattacks that specifically target embedded systems merit special attention and careful study by the DoD.

In this paper, we describe the development of an ontology for sharing information on and reasoning about cyberattacks on embedded systems and discuss how it can be used to support cyber risk assessments (CRAs) such as Cyber Table Top (CTT) exercises.

A. Background

1. The MITRE ATT&CK Knowledge Base


The MITRE Corporation has long recognized the importance of enumerating, cataloging, and describing tactics and techniques employed to conduct cyberattacks. Starting in 2013, MITRE began developing MITRE ATT&CK[®]. ATT&CK is “a knowledge base of cyber adversary behavior and taxonomy for adversarial actions across their lifecycle.”⁹ ATT&CK is typically presented as a matrix of types of cyber objectives and attacks on a platform. The matrix’s columns are attackers’ tactical objectives,¹⁰ termed *tactics*. Each cell under a column head (a tactic name) represents a means of achieving the objective of that column and is termed a *technique*.

Figure 1-1 shows an example ATT&CK matrix. The matrix in question concerns cyberattacks on containers. It has eight tactics (Initial Access, Execution, etc.). Cells below each tactic heading list techniques adversaries might use to achieve the tactic’s objective. For example, Figure 1-1 shows the “Exploit Public-Facing Application” technique in the “Initial Access” column. This means MITRE has concluded an adversary may try to gain initial access to a container by exploiting public-facing applications.

⁸ See <https://www.forbes.com/sites/stevetengler/2020/06/30/top-25-auto-cybersecurity-hacks-too-many-glass-houses-to-be-throwing-stones/?sh=47cfdad97f65>

⁹ See <https://attack.mitre.org/resources/faq/>, “What is ATT&CK?”

¹⁰ The ATT&CK web pages use “goal” rather than “objective.” For the purposes of this paper these words are synonymous, and we use “objective” because it appears in an ontology we use. See Figure 2-2.

Techniques are not tactic-specific. In Figure 1-1, the “Valid Accounts” technique is used for four tactics: Initial Access, Persistence, Privilege Escalation, and Defense Evasion. Techniques in boxes have sub-techniques. The number of sub-techniques is the subscript after the technique name. On the website, clicking the  image displays the sub-techniques.

Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Discovery	Impact
3 techniques	4 techniques	4 techniques	4 techniques	6 techniques	2 techniques	3 techniques	3 techniques
Exploit Public-Facing Application	Container Administration Command	External Remote Services	Escape to Host	Build Image on Host	Brute Force (3)	Container and Resource Discovery	Endpoint Denial of Service
External Remote Services	Deploy Container	Implant Internal Image	Exploitation for Privilege Escalation	Deploy Container	Unsecured Credentials (2)	Network Service Scanning	Network Denial of Service
Valid Accounts (2)	Scheduled Task/Job (1)	Scheduled Task/Job (1)	Scheduled Task/Job (1)	Impair Defenses (1)		Permission Groups Discovery	Resource Hijacking
	User Execution (1)	Valid Accounts (2)	Valid Accounts (2)	Indicator Removal on Host			
				Masquerading (1)			
				Valid Accounts (2)			

Image Source: <https://attack.mitre.org/matrices/enterprise/containers/>

Figure 1-1. ATT&CK Matrix for Containers

Figure 1-1 is one of many matrices in ATT&CK. ATT&CK originally consisted of two matrices, one focused on Windows-based enterprise systems and the other on mobile devices. ATT&CK now consists of three system categories: enterprise systems, mobile devices, and industrial control systems (ICS). Enterprise systems and mobile devices are deemed sufficiently complex to be organized into hierarchies, as shown in Figure 1-2.¹¹ Each node (e.g., Enterprise, Mobile, ICS) corresponds to a matrix that comprises a set of tactics and corresponding techniques. Each child node, whose name is indented below its parent’s name (e.g., PRE and Cloud are children of Enterprise; Office 365 is a child of Cloud), corresponds to a matrix that comprises a subset of its parent matrix’s tactics, and a subset of the tactics’ corresponding techniques. Figure 1-1 shows that the Containers matrix has eight tactics. The matrix for Enterprise, the parent matrix of Containers, has 14 tactics.¹² MITRE’s analysis has thus concluded that 14 tactics are relevant to an enterprise system—any kind of enterprise system—and that of those 14, only eight are relevant to Containers.¹³

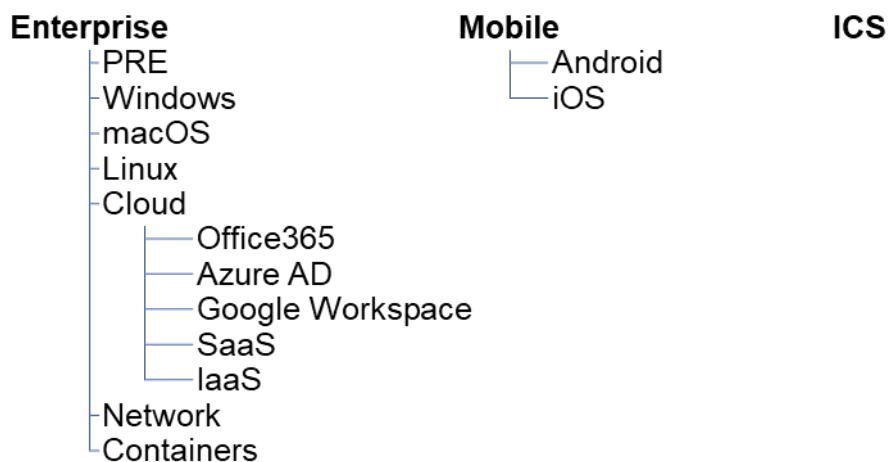


Figure 1-2. ATT&CK Matrices

¹¹ The image is derived from <https://attack.mitre.org/matrices/>

¹² See <https://attack.mitre.org/matrices/enterprise/>

¹³ MITRE does not consider tactic “Resource Development”, which is defined as “techniques that involve adversaries creating, purchasing, or compromising/stealing resources that can be used to support targeting” (see <https://attack.mitre.org/tactics/TA0042/>) relevant to containers. This is justifiable, because targeting a container requires access to the system that hosts the container, and this access would require resource development.

2. Ontological Representation of ATT&CK

The MITRE Corporation's ATT&CK knowledge base is freely available online.¹⁴ Users can search and query this online version to retrieve descriptions of terms. The online term descriptions are unstructured text. Although useful to humans, they are not well suited to automated processing. For example, Section 1.A.1 states that a single technique may be used in multiple tactics. This means that some tactics are associated with identically named techniques. An examination of the ATT&CK knowledge base reveals that sometimes descriptions of techniques with the same name differ from tactic to tactic. Are such techniques truly the same? Automated text processing systems are not at the stage where they can reliably answer this question.

To improve the potential for using ATT&CK concepts in automated reasoning systems, three organizations—CUBRC,¹⁵ the Institute for Defense Analyses (IDA), and the Intelligence and Army Intelligence and Information Warfare Directorate (I2WD)¹⁶—worked together to develop the MITRE ATT&CK Matrix Ontology (MAMO). MAMO expresses the concepts in ATT&CK using the Web Ontology Language (OWL).¹⁷ It organizes concepts hierarchically as classes, relates classes using properties, and establishes paradigms for rigorously expressing knowledge about cyberattacks.

ATT&CK describes categories of cyberattack objectives and the techniques used to achieve them; MAMO supports creating and reasoning about instances of cyberattacks fitting these categories. ATT&CK is concerned with cyberattacks on enterprise systems, mobile devices, and industrial control systems. MAMO, being a representation of ATT&CK, focuses on expressing concepts about those three system categories. Furthermore, to address the issue raised at the beginning of this subsection, OWL has the expressive power to let MAMO's designers explicitly differentiate between a technique that applies to multiple tactics and a set of different techniques that have the same name.

MAMO is a domain-specific ontology. It does not attempt to be general-purpose. It uses the dependency structure illustrated in Figure 1-3. MAMO concentrates on a specific area of importance (cyberattacks). MAMO builds on the Cyber Ontology, a more general domain-specific ontology dealing with concepts in the cyber domain. The Cyber Ontology builds on the Common Core Ontologies (CCO), a set of mid-level ontologies “designed to represent and integrate taxonomies of generic classes and relations across all domains of

¹⁴ See <https://attack.mitre.org/>

¹⁵ See <https://cubrc.org/>

¹⁶ See https://c5isr.ccdc.army.mil/inside_c5isr_center/i2wd/

¹⁷ See <https://www.w3.org/TR/owl2-overview/>

interest.”¹⁸ CCO in turn builds on the Basic Formal Ontology (BFO), an upper-level ontology expressing the most general kinds of classes and relationships.¹⁹ In each case, one ontology builds on another by extending and specializing concepts. CCO is more specific than BFO; the Cyber Ontology is more specific than CCO; and MAMO is more specific than the Cyber Ontology. This common ontology design pattern allows a more specific ontology to take advantage of definitions in the more general one. Thus, CCO follows BFO’s logical model and adds its own axioms. It is more specific than BFO, providing a broad range of commonly used concepts that can be extended into multiple specific domains, facilitating interoperability between them using its shared concepts. The Cyber Ontology follows CCO’s logical model and adds axioms; and MAMO follows the Cyber Ontology’s logical model and adds axioms. All axioms from BFO, CCO, and the Cyber Ontology apply to MAMO, and any concept expressed in these first three ontologies is also expressed in MAMO.

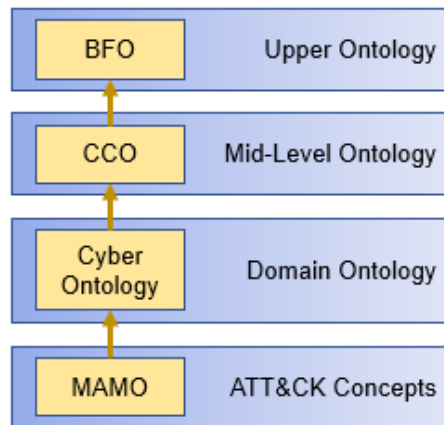


Figure 1-3. MAMO Dependencies

3. PIVOT

As Sections 1.A.1 and 1.A.2 discuss, ATT&CK and MAMO address cyberattacks in enterprise systems, mobile devices, and industrial control systems. They do not address embedded systems, whose secure functions are essential in executing DoD missions. Embedded systems perform key functions in direct control of weapons, observation, transportation, and communication platforms. Given the importance of embedded systems to DoD, analysis of potential embedded system cyberattacks is essential to their effective design, development and deployment.

¹⁸ See <https://github.com/CommonCoreOntology/CommonCoreOntologies>

¹⁹ See <https://basic-formal-ontology.org/>

The MITRE Corporation has done a study of embedded system cyberattacks that it refers to as Platform Independent Vectors of Techniques (PIVOT), as described in Zuniga et al. (2020). (The acronym is motivated by the fact that many attacks on embedded systems rely on first obtaining access to, and then installing malware on, a non-embedded system. This malware then finds a way to install itself on the embedded system. The process of malware moving itself from one system to a different target system is known as pivoting.)

Like each ATT&CK matrix, PIVOT can be represented in a matrix, with columns corresponding to tactics and cells under those columns corresponding to techniques used to achieve the tactical objectives of the column. This matrix, developed and maintained by MITRE, is called the Embedded System Tactics, Techniques and Procedures (TTP) Matrix (ESTM).

The Office of the Director for Developmental Test, Evaluation, and Assessments (D,DTE&A) directed IDA to express ESTM concepts in an ontology. The ontology would express embedded system tactics and techniques, just as MAMO expresses enterprise system, mobile device, and ICS tactics and techniques. It could be used in parallel with MAMO, and along with the Ontology for Attacks in Cyber Risk Assessments (OACRA),²⁰ to extend the knowledge base to a new area, one that is critical to military operations. An attacker who seizes control of an embedded system on a weapons platform can not only disrupt a mission, but might be able to direct lethal force toward friendly troops.²¹

D,DTE&A was motivated to create OACRA and the ESTM ontology because they recognize the value in having ontologies available to support CTT exercises.²² In particular, (1) in preparation for and during an exercise, CTT participants could query an ontology to identify attack techniques and targets that might be useful,²³ and (2) after an

²⁰ IDA developed and delivered OACRA to D,DTE&A in 2020. OACRA builds on MAMO, expressing concepts particularly important in CTTs. In particular, CTTs have been designed to capture data on the attack steps that might be used in a cyberattack. In addition to the cyber tactics and techniques captured in MAMO, OACRA includes an attack name, description, assumption, time, possible outcome, result, and cyber target. Buytendyk et al. (2020) explains OACRA in detail.

²¹ The concepts in ESTM are not entirely new to ATT&CK. Many industrial control systems contain embedded systems. For that reason, both the ESTM and ICS matrices have tactics “Impair Process Control” and “Inhibit Response Function”, neither of which appear in other ATT&CK matrices. However, whereas ESTM focuses on embedded systems, ICS focuses on systems that may contain embedded systems. MITRE’s analyses conclude different techniques are used to achieve an objective. For example, for “Impair Process Control”, ESTM has tactic “Manipulate Instrumentation and/or Controls”, which is not in ICS. Conversely, ICS has “Brute Force I/O”, which is not in ESTM.

²² D,DTE&A uses CTT exercises to identify possible cyberattacks, estimate their risks, and devise mitigations.

²³ MITRE’s web pages for ATT&CK have a search capability. However, searches are limited to free text. An ontology provides a logic-based structure that supports automated reasoning and question-answering (e.g., answers to queries such as “show me all tactics that can be achieved using the Phishing technique.”)

exercise, CTT participants could enter their results and observations into a knowledge base. (They would probably sanitize results to avoid identifying vulnerabilities of specific systems.)²⁴ Such a knowledge base could be useful to system designers seeking to identify potential attacks at the beginning of system design, as well as throughout system development. This, of course, is the objective of the DevSecOps development approach: “shift security left”. Identifying the potential tactics and techniques that may be used in cyberattacks against a system facilitates early identification of suitable mitigations that can protect a system against such attacks.

B. Structure of this Paper

The remainder of this paper explains the ontology derived from PIVOT’s Embedded System TTP Matrix, known as the ESTM Ontology. It explains the ontology’s modeling approach, which is identical to that of OACRA, as described in Buytenchyk et al. (2020). It also describes how the ESTM Ontology was automatically generated from ESTM, and how it can be re-generated in response to ESTM modifications or extensions. It ends with discussion of how the ESTM Ontology can be used, and highlights the value in having ontologies available to support CRAs, and in particular CTTs.

We note the following terminology and notation conventions used in this paper:

- “ESTM” refers to the MITRE-developed matrix of tactics and techniques for cyberattacks on embedded systems.
- “ESTM Ontology” refers to the OWL ontology IDA generated from that matrix.
- Similarly, “MAMO” refers to the ontology developed from the MITRE-developed ATT&CK matrices.
- This paper presents ontology elements from the ESTM Ontology and other ontologies. These elements are shown in the Calibri font. They often have a form similar to:

cco:elucidation

The text before the colon is the element’s namespace,²⁵ qualifying the element’s name (the portion after the colon) to distinguish the entire element from other

²⁴ This paper includes MITRE’s description of ATT&CK as a knowledge base, which is accurate when understood as an informal knowledge base composed of textual knowledge about categories of tactics and techniques. However, ATT&CK is not a formal ontology, which would map directly into a logical formalism. MAMO, in contrast, is a formal ontology that represents the ATT&CK knowledge base. The results of CTT exercises can be captured as instances of MAMO classes and relationships and can therefore constitute a formal knowledge base in the context of MAMO.

²⁵ See <https://www.w3.org/TR/xml-names/> for definition and discussion of namespaces.

elements that might have the same name but a different namespace. Table 1-1 lists the namespaces and their meaning.

- Sometimes the name in an ontology element is enclosed in single quotation marks:

`cco:'definition source'`

This is a convenience notation using a label for the element, which improves readability when the actual name is more suited to automated processing than to human consumption.

Table 1-1. Namespaces used in this paper

Namespace	Description	URI
cco	Common Core Ontologies	http://www.ontologyrepository.com/CommonCoreOntologies/
rdfs	RDF Schema	http://www.w3.org/2000/01/rdf-schema#

(This page is intentionally blank.)

2. ESTM Ontology Modeling Approach

A. Ontology Structure

Figure 1-3 shows the ontology structure for MAMO. The ESTM Ontology has an analogous structure that builds directly on concepts formalized in the Cyber Ontology, which builds on CCO, which builds on BFO. Figure 2-1 illustrates the ESTM Ontology structure. Each arrow in the figure indicates that the ontology at its tail imports the ontology at its head. (CCO actually comprises 12 ontologies and BFO two, so the figure is a simplification.)

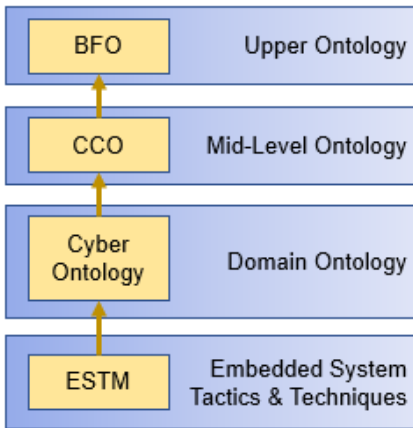


Figure 2-1. ESTM Ontology import structure

B. Tactics and Techniques in Relation to Imported Ontologies

The ESTM Ontology follows the approach used in MAMO to express tactics and techniques relevant to cyberattacks on embedded systems. A tactic is an objective, Objective being a concept expressed in CCO. Figure 2-2 shows the concept hierarchy.²⁶

²⁶ This and other figures are screen captures from the Protégé ontology editor. They use class labels (the `rdfs:label` annotation), not URLs. Labels are more readable than URLs, especially considering BFO 2.0's use of numeric identifiers for class names.

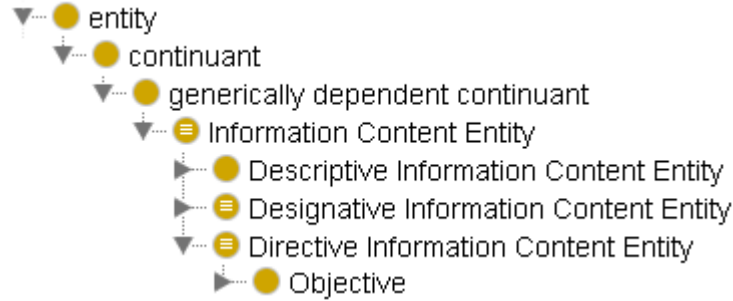


Figure 2-2. Objective class hierarchy

It may be helpful to briefly explain Objective and its ancestors in the hierarchy:²⁷

- An entity is anything that exists, has existed, might exist, or will exist.
- A continuant is an entity that exists through time.
- A generically dependent continuant is a continuant whose existence depends on one or more other entities.
- An Information Content Entity is a generically dependent continuant whose existence depends on some entity that conveys information.
- A Directive Information Content Entity is an Information Content Entity whose propositions prescribe some other entity.
- An Objective is a Directive Information Content Entity that prescribes some end state to be achieved.

The Cyber Ontology, with its focus on the cyber domain, extends CCO’s Objective class with a subclass named ‘Cyber Objective’, thereby grouping tactics related to cyberattacks in a single class hierarchy.²⁸ The Cyber Objective class is the ancestor of all categories of tactics defined in ATT&CK. MAMO defines two subclasses of ‘Cyber Objective’: ‘Cyber Enterprise Tactic’ and ‘Cyber Mobile Tactic’, for tactics relevant to enterprise systems and mobile devices, respectively. The ESTM Ontology incorporates a

²⁷ This short description does not intend to explain all the nuances of these classes and how to use them. See Arp (2015) for a book-length exposition of the top three classes. See the documentation in <https://github.com/CommonCoreOntology/CommonCoreOntologies.git> for information on the other classes.

²⁸ “Cyber Objective” is actually the class’s label, not its name. This paper uses labels, which are more readable. However, referring to them as labels everywhere seems unnecessarily pedantic.

'Cyber Embedded Tactic' class as a corresponding cyber-objective class for embedded systems. See Figure 2-3.




Figure 2-3. Cyber Objective in ESTM Ontology

In MAMO, a technique is a kind of cyberattack. Figure 2-4 shows the hierarchy. Here is a brief explanation. Again, the details are well beyond the scope of this paper.

- An entity is, as defined above, anything that exists, has existed, might exist, or will exist.
- An occurrent is an entity that unfolds itself in space, time, or both.
- A process is an occurrent that happens within some temporal interval and relates to some material entity.
- An Act is a process in which some agent (roughly, an animal or organization) has a role.
- An Intentional Act is an Act in which an agent has a causative role and is prescribed by some Directive Information Content Entity.
- An Act of Artifact Employment is an Intentional Act involving some artifact to achieve the Directive Information Content Entity's prescriptions.
- A Cyber Attack is an Act of Artifact Employment in which an Agent uses a computer to engage in a malicious act.

Therefore, a cyber technique is a process, something that occurs in space and time; has an agent (the attacker) who plays a causative role; and is carried out as prescribed by some 'Directive Information Content Entity', which gives the act intention. Because a MAMO technique involves using some computing artifact, a technique is further categorized as an 'Act of Artifact Employment', and more specifically as a 'Cyber Attack', defined in the Cyber Ontology.




Figure 2-4. Cyber Attack class hierarchy

MAMO classes for enterprise system attack techniques and mobile device attack techniques are subclasses of ‘Cyber Attack’. The ESTM Ontology incorporates a corresponding class for embedded system techniques (Figure 2-5).




Figure 2-5. Cyber Embedded Technique as Cyber Attack subclass

C. Cyber Embedded System Tactics

In the ESTM Ontology, class ‘Cyber Embedded Tactic’ has 13 subclasses, one for each tactic PIVOT identifies. Of these 13, 11 are also tactics associated with enterprise system and mobile device attacks. For example, all three cyber tactic categories have an “initial access” tactic. “Initial access” is, however, simply a phrase that describes an objective. There is no implication that initial access to an embedded system means exactly the same thing as initial access to an enterprise system or mobile device. For this reason, ESTM Ontology tactic classes are separated according to the type of system they target. This is achieved by ending the ESTM tactic class names with the string “EmbeddedTactic”, and ending their labels with the phrase “Embedded Tactic”. These suffixes follow a practice in MAMO to promote readability: someone viewing the ontology easily recognizes the category to which a tactic belongs.

As mentioned above, many (but not all) ESTM and ATT&CK tactics overlap. Table 2-1 shows the tactics for each category (sans category suffix). Notice that the first column has two tactics, “Impair Process Control” and “Inhibit Response Function”, that do not appear in columns 2 and 3. These tactics are unique to embedded systems, which, unlike enterprise systems and mobile devices, perform process control and have (physical) response functions.

Table 2-1. ESTM and ATT&CK Tactics

Embedded System	Enterprise System	Mobile Device
Collection	Collection	Collection
Command and Control	Command and Control	Command and Control
Credential Access	Credential Access	Credential Access
Defense Evasion	Defense Evasion	Defense Evasion
Discovery	Discovery	Discovery
Execution	Execution	Execution
Exfiltration	Exfiltration	Exfiltration
	Impact	Impact
Impair Process Control		
Inhibit Response Function		
Initial Access	Initial Access	Initial Access
Lateral Movement	Lateral Movement	Lateral Movement
		Network Effects
Persistence	Persistence	Persistence
Privilege Escalation	Privilege Escalation	Privilege Escalation
		Remote Service Effects

D. Cyber Embedded System Techniques

The ESTM Ontology has 217 classes that describe techniques for attacking embedded systems. All are direct subclasses of Cyber Embedded Technique. Table 2-2 shows the number of techniques for each tactic.

Table 2-2. ESTM Ontology Technique Counts

Tactic	# Techniques
Collection	21
Command and Control	11
Credential Access	10
Defense Evasion	39
Discovery	13
Execution	18
Exfiltration	11
Impair Process Control	16
Inhibit Response Function	14
Initial Access	15
Lateral Movement	14
Persistence	28
Privilege Escalation	7

As with tactics, ontology class names for techniques have a suffix “EmbeddedTechnique”, which distinguishes them from similarly named techniques for enterprise systems and mobile devices. Labels have the suffix “ Embedded Technique”.

The ESTM Ontology differs from MAMO in that all embedded system cyberattack techniques are represented as direct subclasses of class ‘Cyber Embedded Technique’. In MAMO, a multi-level hierarchy of cyberattack techniques exists. This hierarchy comes from ATT&CK. If a future version of ESTM defines a technique hierarchy, the ESTM Ontology should evolve to reflect that hierarchy.

E. Relationship between Tactics and Techniques

The ESTM and its ontology have been designed so that similar techniques associated with different tactics all have distinct names, which often use part of the corresponding tactic name to distinguish them. Using such distinct technique names, there is a one-to-many relationship between tactics and techniques. The ESTM Ontology formalizes this fact. Each technique class in the ESTM Ontology asserts its relationship to a single tactic. This assertion is done using a subclass existential restriction using the ‘has objective’ property. For example, technique ‘C2 Fallback Channels Embedded Technique’ has the restriction:

● **'has objective' some 'Command and Control Embedded Tactic'**

The related technique ‘Exfiltration Fallback Channels Embedded Technique’ is distinguished by the “Exfiltration” prefix and has the corresponding restriction:

● **'has objective' some 'Exfiltration Embedded Tactic'**

This one-to-many relationship is more restrictive than MAMO. For example, MAMO class ‘Access Token Manipulation Enterprise Technique’ asserts:

● **'has objective' some
('Defense Evasion Enterprise Tactic' or 'Privilege Escalation Enterprise Tactic')**

The implication is that in MAMO the same technique can support achieving two different tactics. As Section 2.D discusses, every technique in the ESTM and its ontology is considered unique, partly in terms of its tactic.

F. Annotations

Classes defined in the ESTM Ontology have a standard set of annotations. These annotations follow the paradigm established by CCO. Every class should have:

- A label (annotation property `rdfs:label`), giving a natural-language name for the class.

- A definition (annotation property `cco:definition`), which is usually a one-sentence description of the class. Definitions derive from the following template:

A *superclass* in which *description*.

Here, *superclass* is the class's parent, and *description* is a phrase that succinctly explains characteristics possessed by all members of the class. For example, the ESTM Ontology's definition of C2 Fallback Channels Embedded Technique is:

A **Cyber Embedded Technique** in which **an adversary uses a fallback or alternate communication channel if the primary channel is compromised or inaccessible in order to maintain reliable command and control and to avoid data transfer thresholds.**

Stating the superclass eliminates any need to mention characteristics common to its superclasses, as well as aiding understanding by a user reading the definition outside an ontology editor such as Protégé.

- An (optional) elucidation (annotation property `cco:elucidation`), a statement of arbitrary length presenting details about the class that cannot easily fit into the (short) definition annotation.
- A definition source (annotation property '`cco:definition source`'), stating the document, website, etc., from which information on the concept derives. In the ESTM Ontology, the definition source is always Zuniga (2021).
- A statement of the ontology (annotation property '`cco:is curated in ontology`') in which the class's definition appears. In the ESTM Ontology, this is always:

<http://www.ontologyrepository.com/CommonCoreOntologies/Domain/ESTM>

which is the namespace of the ESTM Ontology. As Figure 2-1 shows, the ESTM Ontology transitively imports many ontologies. Sometimes a user wants to know the ontology that declares an entity, and this property defines it.

The source documentation (Zuniga 2021) invariably had a one-sentence definition for tactics and techniques, but it did not always provide additional information about them. The ESTM Ontology has no elucidation annotation for classes corresponding to ESTM techniques that lacked such additional information. The ESTM did not include additional information for any of its tactics, so the ESTM Ontology has no elucidations for any tactic classes.

(This page is intentionally blank.)

3. Generating the ESTM Ontology

The ESTM Ontology has been generated using a purpose-built Python program developed by IDA. The program’s input is an Excel spreadsheet containing semi-structured information on tactics and techniques. The output is a file containing an OWL representation of that information.

The spreadsheet is derived from Zuniga (2021). Its contents were manually extracted from that document, and transformed into rows and columns. The document has prose descriptions of each technique. These descriptions were transformed, by hand, into the definition/elucidation format established in CCO.

Table 3-1 shows the spreadsheet’s contents. The program generates, for each row of the spreadsheet, a class expressing the technique with annotations for the technique name, definition, and elucidation, along with canned values for the definition source and the curating ontology (see Section 2.F). It also generates a class each time it encounters a previously unseen tactic. Finally, it asserts a subclass restriction on the technique class, as described in Section 2.E.

Table 3-1. Input Spreadsheet Content

Item	Description
Tactic ID	An alphanumeric identifier that is unique to each tactic name.
Tactic Name	A natural language label that describes the tactic.
Tactic Definition	A succinct, one-sentence definition of the tactic.
Technique ID	An alphanumeric identifier that is unique to each technique name.
Technique Name	A natural language label that describes the technique.
Definition	A succinct, one-sentence definition of the technique.
Elucidation	An arbitrary-length explanation of the technique beyond what the definition states.

The program does not include the tactic ID or the technique ID in the ontology. At this time there is no foreseen use for these values. In MAMO, IDs are used in resolvable URLs (e.g., <https://attack.mitre.org/tactics/TA0009> resolves to the web page that describes “Collection”, the tactic with ID TA0009), but ESTM has no corresponding website.

Zuniga (2021) does not contain descriptions of tactics. The tactic definitions were written for the purposes of the spreadsheet to support ontology generation. Like technique definitions, they are in the CCO format described in Section 2.F.

In the documents IDA originally received from MITRE, each technique had a unique identifier, and sometimes different identifiers had the same technique name. For example, there were two techniques named “Fallback Channels”. One, with identifier EST000172, had tactic Command and Control. The other, with identifier EST000186, had tactic Exfiltration. Their definitions differed, meaning the concept of a Fallback Channels technique is interpreted differently depending on the tactic in question. This posed an ontology design conundrum. Should a single technique have multiple definitions? Should it have a single definition that incorporates multiple meanings? Should there be multiple techniques with the same name (or in the ontology, the same label)? Or should there be multiple techniques with different names? IDA and MITRE jointly decided the ESTM Ontology should have multiple techniques with different names because their details varied depending on the associated tactic. The source documents were rewritten to make each technique name unique. The technique with identifier EST000172 was renamed C2 Fallback Channels, and the technique with identifier EST000186 was renamed Exfiltration Fallback Channels.

4. Using the ESTM Ontology

MAMO is used to model cyberattacks, and the ESTM Ontology has been designed to be used in the same manner. Both ontologies model (cyber) techniques and (cyber) tactics, but for two different categories of systems: (1) embedded systems for ESTM, and (2) enterprise systems and mobile devices for MAMO. The ESTM Ontology is intended to express the same kinds of concepts, and support the same kinds of reasoning, as MAMO.

The initial release of OACRA (V1.0) imports MAMO directly along with the Targets ontology²⁹ (Figure 4-1).




Figure 4-1. OACRA V1.0 relationship to MAMO

Since the CRA community using OACRA needs to address cyberattacks on embedded systems, OACRA V1.1 has been updated to directly import the ESTM Ontology along with MAMO while keeping the ESTM Ontology independent of MAMO (Figure 4-2).




Figure 4-2. OACRA V1.1 relationship to MAMO and ESTM Ontology

²⁹ The Targets ontology was developed for OACRA to identify classes of commonly used targets of cyberattacks in CTTs.

Both MAMO and the ESTM Ontology directly import the Cyber Ontology and indirectly import the CCO and BFO (Figure 4-3).³⁰




Figure 4-3. OACRA V1.1 larger import structure

The cybersecurity community can now use OACRA to model embedded system attacks in the same way it has been using MAMO to model attacks for enterprise and mobile systems.

Achieving this structure has involved moving one class ('oacra:Cyber Objective') and one property ('cco:has objective') from MAMO into the Cyber Ontology. These additions to the Cyber Ontology were planned in coordination with the chair of the Common Core Cyber Ontology Working Group (C3O WG). However, they were not implemented due to resource constraints for managing the C3O. While awaiting renewal of C3O WG activity, we have created a version of the C3O for OACRA that includes this class and property. This has simplified the import structure of the ESTM ontology, so that it is independent of MAMO in accord with its intended interpretation.

³⁰ This continues to simplify the structures of the CCO and BFO, each of which consists of multiple ontologies.

5. Summary

The MITRE Corporation's continuing investment in ATT&CK, its knowledge base of adversarial cyberattack tactics and techniques, is a noteworthy effort in helping the world plan and prepare for the many bad actors and ever-evolving cyber threats they unleash. The ATT&CK website is a valuable resource that should be part of modern software and system development.

MITRE did not intend ATT&CK to describe every possible cyberattack on every possible system. In particular, ATT&CK does not cover embedded systems. Embedded systems are particularly important to DoD: they exist in many, perhaps most, platforms that DoD fields and uses in its missions. Also, MITRE developed its ATT&CK website to support browsing and free text queries: the kinds of interactions a human would make. The website is inarguably useful, but it does not present a formal, rigorous model of cyberattacks, and does not allow the kinds of queries and analysis such a model would support.

MITRE recognizes both gaps. It has developed PIVOT, a set of tactics and techniques that pertain to embedded systems. Although these tactics and techniques are not officially part of ATT&CK, MITRE has also created ESTM, a spreadsheet-based representation of PIVOT in the matrix form used by ATT&CK.

A formal machine-processable model of ATT&CK has been developed in the MAMO. MAMO is an OWL representation of ATT&CK and can be the basis of SPARQL queries, which can be more complex and more specific than free text searches. OWL also offers the opportunity to formally express semantics that are only implied by textual descriptions.

IDA, in cooperation with MITRE, has created an ontological representation of ESTM. The ESTM Ontology is generated automatically from the ESTM spreadsheet; because ESTM is evolving rapidly, automatic generation saves time and increases confidence that the ontology accurately reflects the spreadsheet.

IDA has designed the ESTM Ontology to be structurally analogous to MAMO. Both extend the Cyber Ontology and the Common Core Ontologies. Both fit into OACRA, IDA's ontology to be used during CRAs. By providing the ESTM Ontology, IDA has extended the scope of CRAs to include embedded systems. IDA believes the ESTM Ontology will prove useful during CTT exercises that involve embedded systems, thereby increasing CTT utility throughout DoD.

(This page is intentionally blank.)

Appendix A. References

- Arp, R., Smith, B., and Spear, A., *Building Ontologies with Basic Formal Ontology*. The MIT Press, Cambridge, MA, 2015.
- Buytendyk, A., Haugh, B., and Kuzio de Naray, R., *Ontology for Attacks in Cyber Risk Assessment*. Institute for Defense Analyses Paper P-16400, Alexandria, Virginia, October 2020.
- Kushner, D., “The Real Story of Stuxnet.” *IEEE Spectrum*, 26 February 2013. Available online at <https://spectrum.ieee.org/the-real-story-of-stuxnet>.
- Zuniga, M. et al., *Embedded System TTP Matrix Technique Descriptions*, MP210460, The MITRE Corporation, 2021.
- Zuniga, M., Miller, R., Cannon, D. Robertson, A. and Esbeck, K., *Platform Independent Vectors of Techniques (PIVOT) Concept*, The MITRE Corporation, MITRE Technical Report MTR200611, Annapolis Junction, Maryland, 2020.

(This page is intentionally blank.)

Appendix B. Acronyms and Abbreviations

BFO	Basic Formal Ontology
C3O WG	Common Core Cyber Ontology Working Group
CCO	Common Core Ontologies
CRA	Cyber Risk Assessment
CTT	Cyber Table Top
D,DTE&A	Director, Developmental Test, Evaluation & Assessments
DoD	Department of Defense
ESTM	Embedded System TTP Matrix
I2WD	Intelligence and Information Warfare Directorate
ICS	Industrial Control System
IDA	Institute for Defense Analyses
IT	Information Technology
LAN	Local Area Network
MAMO	MITRE ATT&CK Matrix Ontology
OACRA	Ontology for Attacks in Cyber Risk Assessments
OWL	Web Ontology Language
PIVOT	Platform Independent Vectors Of Techniques
TTP	Tactics, Techniques and Procedures
WAN	Wide Area Network

(This page is intentionally blank.)

REPORT DOCUMENTATION PAGE

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION

1. REPORT DATE 03-2022	2. REPORT TYPE Paper	3. DATES COVERED	
		START DATE	END DATE

4. TITLE AND SUBTITLE An Ontology for the Embedded System TTP Matrix
--

5a. CONTRACT NUMBER HQ0034-19-D-0001	5b. GRANT NUMBER	5c. PROGRAM ELEMENT NUMBER
5d. PROJECT NUMBER AX-1-3100	5e. TASK NUMBER	5f. WORK UNIT NUMBER

6. AUTHOR(S) Wartik, Steven, P.; Haugh, Brian, A.; Plymale, Noah, T.; de Naray, Rachel Kuzio
--

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Institute for Defense Analyses 730 East Glebe Road Alexandria, Virginia 22305	8. PERFORMING ORGANIZATION REPORT NUMBER P-32935 H 2022-000004
--	---

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Ms. Sarah Standard Cybersecurity/Interoperability Technical Director Director, Developmental Test, Evaluation and Assessments (D,DTE&A), OUSD R&E	10. SPONSOR/MONITOR'S ACRONYM(S)	11. SPONSOR/MONITOR'S REPORT NUMBER
--	---	--

12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT <p>The MITRE Corporation's ATT&CK[®] matrices enumerate tactical objectives (tactics) in cyberattacks, and the techniques agents use to achieve those objectives. The ATT&CK matrices cover enterprise systems, mobile systems, and industrial control systems. MITRE recently developed the Embedded System Tactics, Techniques and Procedures Matrix (ESTM) to apply ATT&CK concepts in embedded systems. Embedded systems are crucial to Department of Defense operations. ESTM provides a framework to support analysis of embedded system cybersecurity.</p> <p>To allow using ATT&CK concepts in automated reasoning systems, IDA worked with the Army Information Intelligence and Warfare Directorate and CUBRC to create the MITRE Attack Matrix Ontology (MAMO), an ontological representation of ATT&CK. IDA has recently worked with MITRE to create an ontological representation of ESTM, thereby extending automated reasoning about cybersecurity into the domain of embedded systems. Like MAMO, the ESTM Ontology builds on existing ontologies, inheriting their semantics and design paradigms.</p> <p>The ESTM Ontology can be used in Cyber Table Top exercises to model possible cyberattacks on embedded systems and devise mitigation strategies. During an exercise, participants could query this ontology to identify attack techniques that might be useful. After an exercise, sanitized results could be stored in a knowledge base for future access and analyses.</p>
--

15. SUBJECT TERMS Ontology; Embedded System; OACRA; Cyberattack; Cyber Risk Assessment; MAMO; Cyber Table Top; ATT&CK; Cybersecurity; Tactic; Technique

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified		

19a. NAME OF RESPONSIBLE PERSON John Hong	19b. PHONE NUMBER 703-845-2564
---	--