Cyber threat data model

High-level model and use cases

Matthew Kellett
DRDC – Ottawa Research Centre

Melanie Bernier
DRDC – Centre for Operational Research and Analysis

Defence Research and Development Canada

Reference Document
DRDC-RDDC-2016-D080
December 2016

IMPORTANT INFORMATIVE STATEMENTS

This document was produced as part of the Cyber Decision Making and Response project under the Cyber Operations S&T portfolio.

Abstract

As part of the Cyber Decision Making and Response project, Defence Research and Development Canada (DRDC) is conducting research on threat characterization and investigating the feasibility of supporting the generation and use of cyber intelligence within an automated computer network defence construct. The purpose of this Reference Document is to take the high-level work done so far on the characterization of threats and to put it in the context of a process for developing cyber intelligence and potentially predicting future attacks based on those threats. Accordingly, a high-level cyber threat data model is proposed that conceptually should allow us to bridge the gap between defensive cyber operations and intelligence processes. Three applications of the model are discussed: a reactive application that allows for the detection and assessment of attacks on our network with the potential for attribution to previously unknown actors; a proactive application that allows for the prediction of where future attacks may be targeted based on our understanding of the intent of known actors; and an observational application that allows for the automation of our computer network defences based on the observed traits of ongoing attacks.

Résumé

Dans le cadre du projet Prise de décision et intervention en cybernétique, Recherche et développement pour la défense Canada (RDDC) effectue des recherches sur la caractérisation de la menace et étudie la possibilité de soutenir la production et l'utilisation du cyber-renseignement dans le concept de défense automatisée des réseaux informatiques. Le présent document de référence vise à utiliser le travail de haut niveau réalisé à ce jour sur la caractérisation de la menace dans le cadre d'un processus de collecte du cyber-renseignement et de prévoir éventuellement les attaques futures en fonction de ces menaces. Par conséquent, nous proposons un modèle de données de haut niveau sur les cybermenaces qui nous permettrait, de manière conceptuelle, de combler l'écart entre les cyberopérations défensives et les processus concernant le renseignement. Trois applications de ce modèle sont examinées : une application réactive qui permet de détecter et d'évaluer les attaques visant notre réseau et, éventuellement, de les attribuer à des acteurs auparant inconnus; une application proactive qui permet de prévoir les cibles de futures attaques en fonction de notre compréhension des intentions des acteurs connus; et une application observationnelle qui permet l'automatisation des défenses de notre réseau informatique en fonction des caractéristiques observées des attaques en cours.

Table of contents

List of figures

List of tables

This page intentionally left blank.

1 Introduction

1.1 Background

A research study on how to automate computer network defence for the Department of National Defence / Canadian Armed Forces (DND/CAF) is currently being conducted by Defence Research and Development Canada (DRDC) as part of the Cyber Decision Making and Response (CDMR) S&T project. One of the main activities for this project is the Automated Computer Network Defence (ARMOUR) technology demonstrator, which focuses on automating the observe-orient-decide-act (OODA) loop. Concurrently, other parts of CDMR are researching and developing proofs of concepts that will feed into network defence capabilities such as ARMOUR and/or improve existing algorithms. In particular, the threat characterization and threat targets components of the Metrics portion of CDMR look at how to support the generation and use of cyber intelligence within an automated computer network defence construct. The purpose of this Reference Document is to take the high-level work done so far on the characterization of threats and put it in the context of a process for developing cyber intelligence on those threats and potentially predicting future threats. The follow-on research carried out based on this high-level model is aimed at informing the development of requirements for capital projects currently under consideration by Director General Cyber (DG Cyber).

In this document, we look at existing work on identifying the elements required to characterize cyber threats that was developed as part of a larger, recently completed literature review on approaches to threat characterization and models. We then propose a high-level cyber threat data model that conceptually should allow us to bridge the gap between defensive cyber operations and intelligence processes. We discuss two applications of the model: a reactive application that allows for the detection and assessment of attacks on our network, potentially with attribution to previously unknown actors; and a proactive application that allows for the prediction of where future attacks may be targeted based on our understanding of the intent of known actors.

1.2 Cyber threat characterization framework

DRDC recently contracted a literature review on approaches to characterizing and modeling threats in the cyber environment [1]. The literature survey found that no single existing approach addressed all the requirements that a cyber threat model would need to support the DND/CAF. Based on existing literature, the report proposes a cyber threat characterization framework [1, Section 5.0], which is summarized below and further outlined in Annex A. The framework, which is partly based on existing work by DRDC [2], lays the initial foundation for threat characterization and forms the basis for the cyber threat data model outlined in this document. See the contract report [1] for more information on the characterization of threats in the cyber environment, in general, and on the cyber threat characterization framework, specifically.

The following top-level cyber threat characterization categories are taken from [1] (more detail in Annex A and [1]):

- **Adversary**
  - Type
  - Motivation (hostile, non-hostile)
  - Commitment
  - Resources
- **Asset**
  - Profile
  - Container
  - Vulnerability
- **Attack**
  - Delivery mechanisms
  - Tools
  - Automation
  - Actions
- **Effect**
  - Cyber effects
  - Effect on military activities

2 Cyber threat data model

The cyber threat characterization framework from [1] consists of only data elements, while the data model that we propose here includes the relationships between those elements. The elements and their relationships reflect the processes necessary to generate cyber intelligence and, taken to their natural extreme, predict future attacks. The proposed cyber threat data model is outlined visually in Figure 1. The threat (red line) versus the effect (blue line) distinguishes between two segments of the model. The effect / blue part of the model at the lower right is the "what" part that represents those aspects of an incident that can be directly observed on a network. These effects can be real in that they have actually happened, or intended in that they were supposed to happen. The threat / red part of the model at the upper left is the "who" part that represents what an incident implies about the person or group that caused it and which may change over time as our understanding of the threat changes.

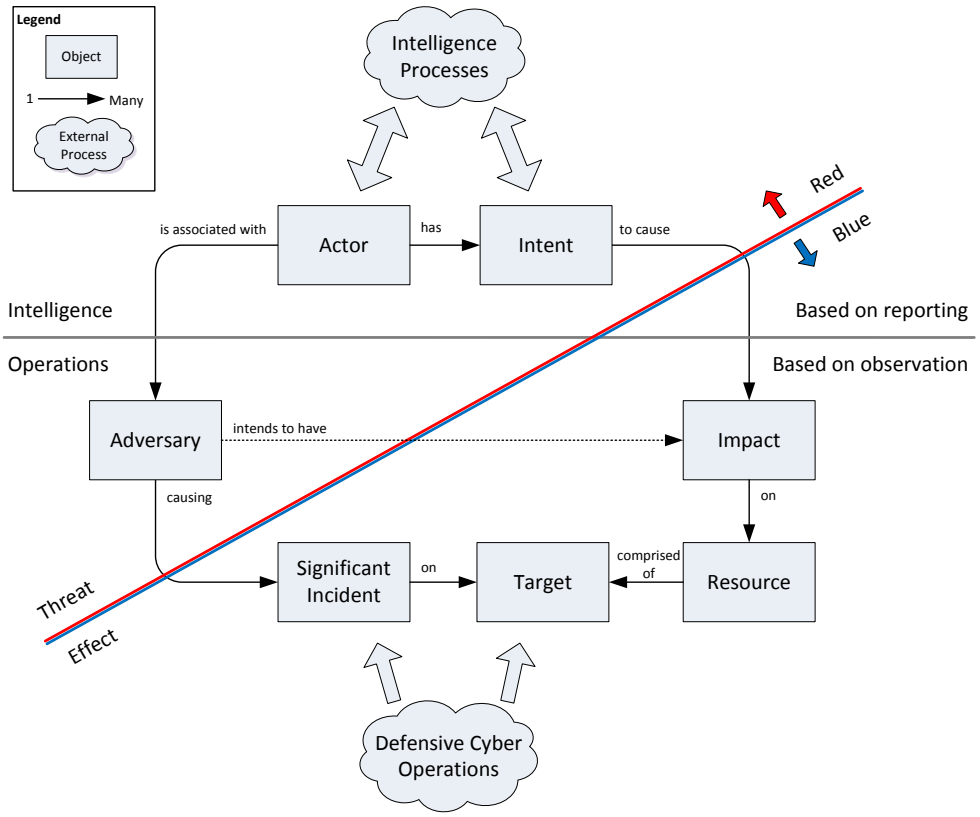

Figure 1: Proposed cyber threat data model.

The full cyber threat data model comprises a series of one-to-many relationships represented by subsets of the model as expressed by the following statements and in Figures 2 and 3:

1. (Threat driven) An ***actor*** is associated with ***adversary(ies)*** causing ***significant incident(s)*** on ***target(s)***.

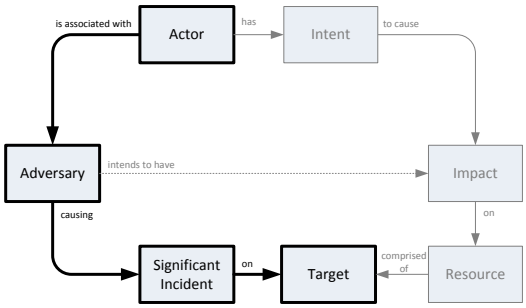

Figure 2: The threat-driven segment (Statement 1) of the cyber threat data model.

2. (Effect driven) An ***actor*** has ***intent(s)*** to cause ***impact(s)*** on ***resource(s)*** comprised of ***target(s)***.

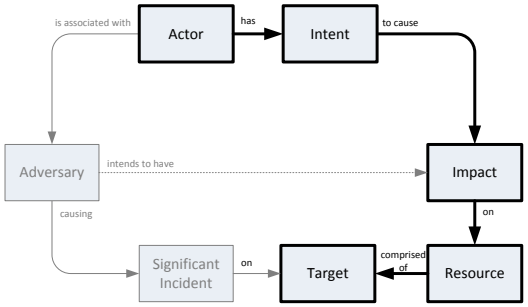

Figure 3: The effect-driven segment (Statement 2) of the cyber threat data model.

At the purely observational/operational subset of the model, these relationships are represented by the following statement and in Figure 4:

3. (Observational) An ***adversary*** causing ***significant incident(s)*** on ***target(s)*** intends to have ***impact(s)*** on ***resource(s)*** comprised of ***target(s)***.

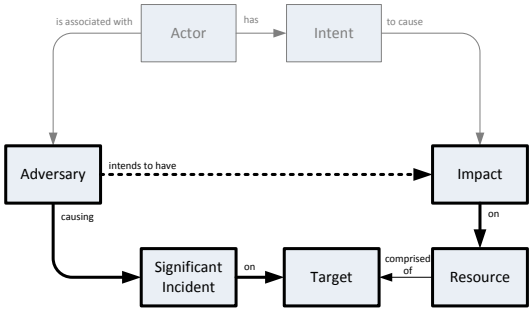

Figure 4: The observational sub-model (Statement 3) of the cyber threat data model.

We distinguish between the full model and the observational sub-model due to the classification of the information likely to be associated with each. If the model is restricted to information that comes only from defensive cyber operations being carried out on the network being monitored—the observational sub-model—its information can be stored at the same classification as the network. When intelligence reporting is added in the full model as part of integrating it with

higher-level intelligence processes, the model's information must then be stored at the maximum classification of the reporting, which in most cases will be higher than that of the network. The practical consequence of this difference is that the observational sub-model is likely to run in near real time on the target network, while the full model is likely to pull a copy of the observational sub-model to a high-side network periodically for human-in-the-loop processing.

Next we examine the three statements and the relationships they describe between the objects in the model. The heart of all three statements is the ***significant incident***-[on]-***target(s)*** relationship. A target can be any discrete part of a computer network infrastructure, including but not limited to an internet protocol (IP) address. A significant incident is an attempt, usually malicious, to compromise a target. Significant incidents can vary from malicious packets and phishing emails sent to the target to malicious websites visited by the target that download malicious content.

In Statement 1, the threat-driven statement, an adversary is some external entity that is observed causing significant incident(s) using tactics, techniques, and procedures (TTPs) that are related. An actor is known to be associated with an adversary(ies) based on intelligence reporting.

Because the ***actor***-[is associated with]-***adversary(ies)*** relationship is not based on observation, there can be a large amount of uncertainty in the relationship. An adversary may meet the reported characteristics of a number of actors, especially when a new adversary appears. The certainty of the relationship must be captured and managed by the model. When new information becomes available, the data model should allow for a new certainty to be assigned to the relationship, while maintaining the history of previously recorded relationships and certainties.

In Statement 2, the effect driven statement, a ***resource***-[is comprised of]-***target(s***). Where a target is a tangible part of the computer network, a resource may be at any layer of the cyber environment (geographical, physical network, logical network, cyber persona, persona). For instance, a phishing email target may be the mail server, but the resource targeted can be the mail server itself, the person to whom it is addressed, or the group to whom that person belongs. A resource can be made up of other resources, but this type of abstraction should be reserved to situations where it is clearly useful, such as linking together resources associated with an otherwise disparate group (all the people on the same course getting phishing emails).

The consequences of a significant incident can lead to an ***impact***-[on]-***resource(s)***. An impact is any observed, suspected, or intended result of a significant incident. For instance, the impact of a phishing email incident could be the exfiltration of data from a resource. Even if the phishing email incident is stopped from reaching its target, we can record that an exfiltration was the intended impact of the significant incident. Intent is the desire of an actor to have certain impacts and is based on intelligence reporting about the known or suspected goals of that actor. Based on general intelligence reporting, the intent(s) of a particular actor can also be inferred from the impacts that actor is having on the network and how related those results are to the actor's known goals. Note that the ***intent***-[to cause]-***impact(s)*** and ***actor***-[has]-***intent(s)*** relationships have the same uncertainty as the ***actor***-[is associated with]-***adversary(ies)*** relationship discussed above and will need to be managed by the model in a similar fashion.

In Statement 3, the observational statement, we short circuit the data model and depend only on observed behaviour. We infer that an ***adversary-***[intends to have]-***impact(s)***. This relationship is

useful for short term responses and may be useful in informing the determination of the ***actor***-[has]-***intent(s)*** relationship in the full model.

Returning to the cyber threat characterization framework from [1], we show in Table 1 how the categories from the framework might be matched up with the objects of the proposed cyber threat data model. These proposed categories can be used to fill in the data required for each element of the model.

Table 1: Potential mapping of the threat characterization framework to the data model.

Data Model Objects	Threat Characterization Categories
Actor	Adversary/Type
	Adversary/Commitment
	Adversary/Resources
Adversary	Adversary/Motivation (non-hostile)
	Attack / Delivery mechanisms
	Attack/Tools
	Attack/Automation
Significant Incident	Attack/Actions
Intent	Adversary/Motivation (hostile)
	Effect / Cyber effects
Impact	Effect / Effect on military activities
Resource	Asset/Container
	Asset/Vulnerability
Target	Asset/Profile

3 Use cases: Bottom-up and top-down

3.1 Bottom up: Generating cyber intelligence

We can use the full cyber threat data model to generate cyber intelligence from the bottom up. At each step, we work our way backwards up the threat-driven and effect-driven statements to determine the actor carrying out the attack and the intent of the attack. We present a simple example here. The statements are adjusted to make them grammatical while still maintaining their original meaning.

Let us consider a phishing email (***significant incident***) sent to a user's mailbox (***target***). When these "facts" are first entered into the model, they are assigned to the generic adversary and generic resource, respectively. Each element of the model uses a generic version as a placeholder while waiting for further information or processing, leading to the following statements:

1. (Threat driven) A ***generic actor*** is associated with a ***generic adversary*** causing a ***phishing email incident*** on ***user A's mailbox***.

2. (Effect driven) A ***generic actor*** has ***generic intent*** to cause ***generic impact*** on a ***generic resource*** comprised of ***user A's mailbox***.

3. (Observational) A ***generic adversary*** causing a ***phishing email incident*** on ***user A's mailbox*** intends to have ***generic impact*** on a ***generic resource*** comprised of ***user A's mailbox***.

We now walk through how the process might proceed using the proposed data model as new information is discovered. We start with the heart of the model, the ***significant incident***-[on]-***target*** relationship. In this simplified example, a spear-phishing email[1] has arrived in an employee's mailbox.

Table 2: Significant incident example record.

Significant incident	
Attack type	Spear-phishing
Timestamp	Yesterday, 5:35 EST
Originator	sender2390842
Origin domain	badmailserver.com
Origin IP	192.168.3.234
Payload	Unknown (In forensics)

[1] A phishing attack is an email campaign where the emails sent are designed to get a large number of recipients to click on a malicious link or open a malicious attachment. A spear-phishing attack is a phishing attack specifically designed to appeal to a single recipient or small group of recipients by using targeted and believable information.

Table 3: Target example record.

Target	
Asset	Mailbox
User	Bloggins, J 234
Email	j.bloggins@ourdomain.ca
Server	exchange02.ourdomain.ca
Server IP	172.21.42.124

The forensics investigation is conducted and determines that the payload of the spear-phishing email is a new virus. The purpose of the virus is to exfiltrate any documents on the target's computer related to specific keywords. An adversary is created to represent incidents that use this version of the virus.

Table 4: Adversary example record.

Adversary	
Attack vector / delivery mechanism	Email
Distinguishing characteristic	virus2353

A resource is created to represent the documents on the email recipient's computer or computers that contain the specific keywords.

Table 5: Resource example record.

Resource	
Location	Target recipient's computers
Resource type	Documents
Distinguishing info	<virus2353 keywords>

An investigation of the target recipient determines that she is a scientist who recently travelled to a conference where she presented her research related directly to the virus keywords. Luckily, she was not fooled by the spear-phishing email and did not infect her computer. An impact is created to represent the loss of information on the research topic.

Table 6: Impact example record.

Impact	
Type	Loss of information
Real or intended?	Intended
Topic	Research on <topic>
Significance	Low

So far, all the information that has been generated in the model is based on observation. We can use this for the observational subset of the model (i.e. Statement 3) and it can be held at the classification of the network. The information generated so far can also be uploaded into a higher-level environment, so that further analysis can be done on the full model.

In the full model, an analyst looks at the adversary and intended impact. She may also look in more detail at the forensics report and investigative report for additional information that may be at a higher level of classification than the observational subset of the model. Based on her analysis, she determines that there are two possible actors that could be behind the adversary. Neither is known to be interested in the research topic, so this incident represents new behaviour. She cannot distinguish between the two, so she adds both into the data model and a new intent to learn information on the research topic.

Table 7: Actor example record—FALLING WHALE.

Actor	
Intrusion set	FALLING WHALE
Confidence level	50%

Table 8: Actor example record—SENTIENT PETUNIA.

Actor	
Intrusion set	SENTIENT PETUNIA
Confidence level	50%

Table 9: Intent example record.

Intent	
Type	Interception/Steal
Topic	<topic>
Importance	Low or unknown

If further incidents are attributed to the same adversary, it may lead to more information on exactly which actor is responsible. In the meantime, the analyst can write an assessment to warn that the research topic is now of interest to a state actor.

3.2 Top down: Predicting future attacks

We may also be able to predict future attacks on our networks by using the data model from the top down approach. When we have reporting on actors and their strategic goals, we should be able to use those to predict adversaries and targets in some cases.

For this section, we use the same example as in Section 3.1 but omit the tables. Working on the higher side, the analyst may have learned from outside reporting that another country is now interested in the research topic.

Let us say that it is the country associated with the threat actor known as FALLING WHALE. The team behind FALLING WHALE is known to specialize in both phishing emails and the exfiltration of data from target networks. The analyst can now create an adversary based on a potential phishing email campaign and an impact based on the exfiltration of data on the research topic.

Through an investigation, it is determined that a small group of researchers works on the research topic. A resource is created to represent the group and their management. The targets associated

with the resource are these employees' mailboxes. Based on the top-down use of the model, these mailboxes can now be watched for signs of phishing-emails and the researchers themselves can be warned to be on the lookout for anything suspicious. When the targeted researcher gets the email from the first example, the organization will already be prepared.

Let us now consider the generic statements about what we know:

1. (Threat driven) A ***known actor FALLING WHALE*** is associated with ***adversaries using phishing emails*** that may intend to cause ***receipt of phishing emails*** in ***employees' mailboxes***.

2. (Effect driven) A ***known actor FALLING WHALE*** has known ***goal to gather information on conference topic*** by causing ***exfiltration of data*** from ***conference attendees*** comprised of ***employees' mailboxes***.

3. (Observational) An ***adversary using phishing emails*** may intend to cause ***receipt of phishing emails*** in ***employees' mailboxes*** and likely intends to ***exfiltrate data*** from ***conference attendees*** comprised of ***employees' mailboxes***.

4 Conclusion

We have taken inspiration from the cyber threat characterization framework [1] and developed a cyber threat data model that can be used to connect defensive cyber operations with intelligence operations. We have shown how data about attacks on targets can be transformed using the data model into cyber intelligence from our networks. We have also shown how reporting on actors and their intent can be used to predict future attacks and close the gap between attacks and attribution.

Before we can integrate the data model into automated computer network defence capabilities, there are a number of challenges that need to be addressed. We need to further develop the details of the data model as well as the procedures required to take advantage of the model's probabilistic and predictive abilities. Finally, we need to develop scenarios or use existing ones to test the completeness of the detailed model and the effectiveness of its related procedures for handling real world situations.

This page intentionally left blank.

References

[1] Magar, A. (2016), *State-of-the-Art in Cyber Threat Models and Methodologies* (DRDC-RDDC-2016-C132), Sphyrna Security, Kanata, Ontario.

[2] Bernier, M. (2013), *Military Activities and Cyber Effects (MACE) Taxonomy* (DRDC CORA TM 2013-226), Defence Research & Development Canada – Centre for Operational Research and Analysis.

This page intentionally left blank.

Annex A Cyber threat characterization framework

The cyber threat characterization framework proposed in [1, Section 5.1]. Cyber threat is broken down into Adversary, Attack, Asset, and Effect.

- **Adversary**
  - Type
    - Script kiddies, newbies, novices
    - Hacktivists, political activists
    - Cyber punks, crashers, thugs
    - Insiders, user malcontents
    - Insider/outside collusion
    - Coders, writers
    - Black hat hackers, professionals, elite
    - Cyber terrorists
    - Nation-states
  - Motivation
    - Hostile
      - Ideological
      - Financial
      - Revenge
      - Ego
      - Psychotic
      - Coercion
    - Non-hostile
      - Mistakes
      - Errors
      - Omissions
  - Commitment
    - Intensity
    - Stealth
    - Time
  - Resources
    - Personnel
    - Knowledge
    - Access
- **Attack**
  - Delivery mechanisms
    - Local access
    - Remote delivery
    - Distributed delivery
    - Social engineering
  - Tools
    - Physical attack
    - Information exchange
    - User command
    - Script or program
    - Autonomous agent
    - Tool
    - Distributed tool
    - Data tap
  - Automation
    - Automatic
    - Semi-automatic
    - Manual
  - Actions
    - Probe
    - Scan
    - Flood
    - Authenticate
    - Bypass
    - Spoof
    - Read
    - Copy
    - Steal
    - Modify
    - Delete
- **Asset**
  - Profile
    - Features
    - Quality
    - Characteristics
    - Value
  - Container
    - Hardware
      - Workstations
      - Servers
      - Network devices
      - Storage devices
      - Mobile devices
    - Software
      - Applications
      - Operating systems
      - Virtual images
    - Objects
      - Paper
    - People
      - Employees
      - Partners
      - Contractors
  - Vulnerability
    - Previously known
      - Design
      - Implementation
      - Configuration
    - Previously unknown
      - Design
      - Implementation
      - Configuration
- **Effect**
  - Cyber effects
    - Interruption
    - Modification
    - Degradation
    - Fabrication
    - Interception
  - Effect on military activities
    - Deny
    - Degrade
    - Disrupt
    - Destroy
    - Digital espionage

This page intentionally left blank.

DOCUMENT CONTROL DATA

(Security markings for the title, abstract and indexing annotation must be entered when the document is Classified or Designated)

1. ORIGINATOR (The name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g., Centre sponsoring a contractor's report, or tasking agency, are entered in Section 8.)

DRDC – Ottawa Research Centre
Defence Research and Development Canada
3701 Carling Avenue
Ottawa, Ontario K1A 0Z4
Canada

2a. SECURITY MARKING (Overall security marking of the document including special supplemental markings if applicable.)

UNCLASSIFIED

2b. CONTROLLED GOODS

(NON-CONTROLLED GOODS)
DMC A
REVIEW: GCEC DECEMBER 2013

3. TITLE (The complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S, C or U) in parentheses after the title.)

Cyber threat data model: High-level model and use cases

4. AUTHORS (last name, followed by initials – ranks, titles, etc., not to be used)

Kellett, M.; Bernier, M.

5. DATE OF PUBLICATION (Month and year of publication of document.)

December 2016

6a. NO. OF PAGES (Total containing information, including Annexes, Appendices, etc.)

22

6b. NO. OF REFS (Total cited in document.)

2

7. DESCRIPTIVE NOTES (The category of the document, e.g., technical report, technical note or memorandum. If appropriate, enter the type of report, e.g., interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.)

Reference Document

8. SPONSORING ACTIVITY (The name of the department project office or laboratory sponsoring the research and development – include address.)

DRDC – Ottawa Research Centre
Defence Research and Development Canada
3701 Carling Avenue
Ottawa, Ontario K1A 0Z4
Canada

9a. PROJECT OR GRANT NO. (If appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant.)

05ac

9b. CONTRACT NO. (If appropriate, the applicable number under which the document was written.)

10a. ORIGINATOR'S DOCUMENT NUMBER (The official document number by which the document is identified by the originating activity. This number must be unique to this document.)

DRDC-RDDC-2016-D080

10b. OTHER DOCUMENT NO(s). (Any other numbers which may be assigned this document either by the originator or by the sponsor.)

11. DOCUMENT AVAILABILITY (Any limitations on further dissemination of the document, other than those imposed by security classification.)

Unlimited

12. DOCUMENT ANNOUNCEMENT (Any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in (11) is possible, a wider announcement audience may be selected.))

Unlimited

13. ABSTRACT (A brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual.)

As part of the Cyber Decision Making and Response project, Defence Research and Development Canada (DRDC) is conducting research on threat characterization and investigating the feasibility of supporting the generation and use of cyber intelligence within an automated computer network defence construct. The purpose of this reference document is to take the high-level work done so far on the characterization of threats and to put it in the context of a process for developing cyber intelligence and potentially predicting future attacks based on those threats. Accordingly, a high-level cyber threat data model is proposed that conceptually should allow us to bridge the gap between defensive cyber operations and intelligence processes. Three applications of the model are discussed: a reactive application that allows for the detection and assessment of attacks on our network with the potential for attribution to previously unknown actors; a proactive application that allows for the prediction of where future attacks may be targeted based on our understanding of the intent of known actors; and an observational application that allows for the automation of our computer network defences based on the observed traits of ongoing attacks.

Dans le cadre du projet Prise de décision et intervention en cybernétique, Recherche et développement pour la défense Canada (RDDC) effectue des recherches sur la caractérisation de la menace et étudie la possibilité de soutenir la production et l'utilisation du cyber-renseignement dans le concept de défense automatisée des réseaux informatiques. Le présent document de référence vise à utiliser le travail de haut niveau réalisé à ce jour sur la caractérisation de la menace dans le cadre d'un processus de collecte du cyber-renseignement et de prévoir éventuellement les attaques futures en fonction de ces menaces. Par conséquent, nous proposons un modèle de données de haut niveau sur les cybermenaces qui nous permettrait, de manière conceptuelle, de combler l'écart entre les cyberopérations défensives et les processus concernant le renseignement. Trois applications de ce modèle sont examinées : une application réactive qui permet de détecter et d'évaluer les attaques visant notre réseau et, éventuellement, de les attribuer à des acteurs auparant inconnus; une application proactive qui permet de prévoir les cibles de futures attaques en fonction de notre compréhension des intentions des acteurs connus; et une application observationnelle qui permet l'automatisation des défenses de notre réseau informatique en fonction des caractéristiques observées des attaques en cours.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (Technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus, e.g., Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

threat; cyber operations; data model; threat characterisation; threat model; intelligence processes; cyber intelligence; defensive cyber operations