A Target-Centric Ontology for Intrusion Detection

Jeffrey Undercoffer, John Pinkston, Anupam Joshi and Timothy Finin
University of Maryland, Baltimore County
Department of Computer Science and Electrical Engineering
1000 Hilltop Circle, Baltimore, MD 21250
{undercoffer, pinkston, joshi, finin}@umbc.edu

Abstract

We have produced an ontology specifying a model of computer attack. Our ontology is based upon an analysis of over 4,000 classes of computer intrusions and their corresponding attack strategies and is categorized according to: system component targeted, means of attack, consequence of attack and location of attacker. We argue that any taxonomic characteristics used to define a computer attack be limited in scope to those features that are observable and measurable at the target of the attack. We present our model as a target-centric ontology that is to be refined and expanded over time. We state the benefits of forgoing dependence upon taxonomies, in favor of ontologies, for the classification of computer attacks and intrusions. We have specified our ontology using the DARPA Agent Markup Language and have prototyped it using DAMLJessKB. We present our model as a target-centric ontology and illustrate the benefits of utilizing an ontology lieu of a taxonomy, by presenting a use case scenario of a distributed intrusion detection system.

1 Introduction

Based upon empirical evidence we have produced a model of computer attacks categorized by: the system component targeted, the means and consequence of attack, and the location of the attacker. Our model is represented as a *target-centric* ontology, where the structural properties of the classification scheme is in terms of features that are observable and measurable by the target of the attack or some software system acting on the target's behalf. In turn, this ontology is used to facilitate the reasoning process of detecting and mitigating computer intrusions.

Traditionally, the characterization and classification of computer attacks and other intrusive behaviors have been limited to taxonomies. Taxonomies, however, lack the necessary and essential constructs needed to enable an intrusion detection system (IDS) to reason over an instance that is representative of the domain of a computer attack. Alternatively, ontologies provide powerful constructs that include machine interpretable definitions of the concepts within a domain and the relations between them. Ontologies, therefore, provide software systems with the ability to share a common understanding of the information at issue, in turn empowering the software system with a greater ability to reason over and analyze this information.

As detailed by Allen, et al. [1], and McHugh [24], the taxonomic characterization of intrusive behavior has typically been from the attacker's point of view, each suggesting that alternative taxonomies need to be developed. Allen et al., state that intrusion detection is an immature discipline and has yet to establish a commonly accepted framework. McHugh suggests classifying attacks according to protocol layer or, as an alternative, whether or not a completed protocol handshake is required. Likewise, Guha [9] suggests an analysis of each layer of the TCP/IP protocol stack to serve as the foundation for an attack taxonomy.

As an alternative to a taxonomy, we propose a data model implemented with an ontology representation language such as the Resource Description Framework Schema (RDFS) [28] or DARPA Agent Markup Language [12]. We illustrate the benefits of using ontologies by presenting an implementation of our ontology being utilized by a distributed intrusion detection system. Accordingly, we have specified our target-centric ontology in DAML and have implemented it using DAML-JessKB [19], an extension to the Java Expert System Shell [6].

Because IDS's are either adjacent to or co-located with the target of an attack it is imperative that any classification scheme used to represent an attack be *target-centric*, where each taxonomic character is comprised of properties and features that are observable by the target of the attack. Consequently, our ontology only defines properties and attributes that are observable and measurable by the target of an attack. As a basis for establishing our *a posteriori* target-centric attack ontology, we evaluated and analyzed over 4,000 computer vulnerabilities and the corresponding attack strategies employed to exploit them.

The remainder of this paper is organized as follows: Section 2 presents related work in the form of alternative attack taxonomies as well as presenting related work in the area of ontologies for intrusion detection. Section 3 presents the characteristics of a sufficient taxonomy. Section 4 details the motivation for abandoning taxonomies in favor of ontologies. Our target-centric attack taxonomy is presented in Section 5. Sec-

tion 6 details our implementation and Section 7 provides an example scenario illustrating the utility of the ontology within a distributed intrusion detection system. We conclude with Section 8.

2 Related Work

As previously stated, most of the existing research in the area of the classification of computer attacks is limited to taxonomies. Because a taxonomy is contained within an ontology we address the research in the area of defining intrusion taxonomies before we address ontologies. Accordingly, this section is subdivided, with Subsection 2.1 presenting related work in the area of taxonomies for intrusion detection and Subsection 2.2 presenting related work in the area of ontologies for intrusion detection.

2.1 Related Work: Taxonomies

There are numerous attack taxonomies proposed for use in intrusion detection research.

In [21] Landwehr et al., present a taxonomy categorized according to genesis (how), time of introduction (when) and location (where). They include sub-categories of: validation errors, boundary condition errors and serialization errors, which we incorporate into our ontology as the means of an attack.

During the 1998 and 1999 DARPA Off Line Intrusion Detection System Evaluations [11] [23] [17] Weber provided a taxonomy defining the categories of consequence, to include *Denial of Service, Remote to Local* and *User to Root*, which we incorporate into our work.

Lindqvist and Jonsson [22] state that they *"focus on the external observations of attacks and breaches which the system owner can make"*. Our effort is consistent with their focus.

2.2 Related Work: Ontologies

There is little, if any, published research formally defining ontologies for use in Intrusion Detection.

Raskin et al. [27], introduce and advocate the use of ontologies for information security. In arguing the case for using ontologies, they state that an ontology organizes and systematizes all of the phenomena (intrusive behavior) at any level of detail, consequently reducing a large diversity of items to a smaller list of properties.

In commenting on the IETF's IDMEF, Kemmerer and Vigna [16] state *"it is a but a first step, however additional effort is needed to provide a common ontology that lets IDS sensors agree on what they observe"*.

3 Characteristics of a Sufficient Taxonomy

At this point, a clear understanding of the definition, purpose and objective of a taxonomy is in order. Accordingly, a *taxonomy* is a *classification* system where the classification scheme conforms to a systematic arrangement into groups or categories according to established criteria [32]. Glass and Vessey [8] contend that taxonomies provide a set of unifying constructs so that the area of interest can be systemically described and aspects of relevance may be interpreted. The overarching goal of any taxonomy, therefore, is to supply some predictive value during the analysis of an unknown specimen, while the classifications within the taxonomy offer an explanatory value.

According to Simpson [29] classifications may be created either *a priori* or *a posteriori*. An *a priori* classification is created non-empirically whereas an *a posteriori* classification is created by empirical evidence derived from some data set. Simpson defines a taxonomic character as a feature, attribute or characteristic that is divisible into at least two contrasting states and used for constructing classifications. He further states that taxonomic characters should be observable from the object in question.

Amoroso [2], Lindqvist, et al. [22] and Krusl [20] each have identified what they believe to be the requisite properties of a sufficient and acceptable taxonomy for computer security. Collectively, they have identified the following properties as essential to a taxonomy: Mutually Exclusive, Exhaustive, Unambiguous, Repeatable, Accepted, Useful, Comprehensible, Conforming, Objective, Deterministic and Specific. Accordingly, as an ontology subsumes a taxonomy these characteristics form the underpinnings of our work.

4 From Taxonomies to Ontologies: *The case for ontologies in Intrusion Detection*

Ning et al. [25], propose a hierarchical model for attack specification and event abstraction using three concepts essential to their approach: *System View, Misuse Signature* and *View Definition*. Their model is based upon a thorough examination of attack characteristics and attributes and is encoded within the logic of their proposed system. Consequently, this model is not readily interchangeable and reusable by other systems.

The Intrusion Detection Working Group of Internet Engineering Task Force (IETF) has proposed the Intrusion Detection Message Exchange Requirements [34] which, in addition to defining the requirements for the Intrusion Detection Message Exchange Format, also specifies the architecture of an IDS. The Intrusion Detection Message Exchange Format Data Model (IDMEF) and accompanying Extensible Markup Language Document Type Definition [3] is a profound effort to establish an industry wide data model which defines computer intrusions. IDMEF, however, has its shortcomings. Specifically, it uses XML which is limited to a syntactic representation of the data model which does not convey the semantics, relationships, attributes and characteristics of the objects which it represents.. This limitation requires that each individual IDS interpret and implement the data model programmaticaly.

According to Davis et al. [4], knowledge representation is a surrogate or substitute for an object under study. In turn, the surrogate enables an entity, such as a software system, to reason about the object. Knowledge representation is also a set of *ontological* commitments specifying the terms that describe

the essence of the object. In other words, *meta-data* or data about data describing their relationships.

Frame Based Systems are an important thread in knowledge representation. According to Koller et al. [18], Frame Based Systems provide an excellent representation for the organizational structure of complex domains. Frame Based Languages, which support Frame Based Systems, include RDF, and are used to represent ontologies. According to Welty et al. [33], an ontology, at its deepest level, subsumes a taxonomy. Similarly, Noy and McGuinness [26] state the process of developing an ontology includes arranging classes in a taxonomic hierarchy.

In applying ontologies to the problem of intrusion detection, the power and utility of the ontology is not realized by the simple representation of the attributes of the attack. Instead, **the power and utility of the ontology is realized by the fact that we can express the relationships between collected data and use those relationships to deduce that the particular data represents an attack of a particular type**. Moreover, specifying an ontological representation decouples the data model defining an intrusion from the logic of the intrusion detection system. The decoupling of the data model from the IDS logic enables non-homogeneous IDS's to share data without a prior agreement as to the semantics of the data. To effect this sharing, an instance of the ontology is shared between IDS's in the form of a set of DAML (or RDF) statements. If the recipient does not understand some aspect of the data, it obtains the ontology in order to interpret and use the data as intended by its originator.

Ontologies therefore, unlike taxonomies, provide powerful constructs that include machine interpretable definitions of the concepts within a specific domain and the relations between them. In our case the domain is that of a particular computer or a software system acting on the computer's behalf in order to detect attacks and intrusions. Ontologies may be utilized to not only provide an IDS with the ability to share a common understanding of the information at issue but also further enable the IDS with improved capacity to reason over and analyze instances of data representing an intrusion. Moreover, within an ontology, characteristics such as cardinality, range and exclusion may be specified and the notion of inheritance is supported.

5 Target-Centric Ontology

We have constructed our ontology in accordance with the results of a detailed analysis of the *CERT/CC Advisories* maintained by the "Computer Emergency Response Team/Coordination Center" of Carnegie Mellon University's Software Engineering Institute and the *"Internet Catalog of Assailable Technologies"* (ICAT) maintained by the National Institute of Standards [13, 14]. Accordingly, the attributes of the Class *Intrusion* consist of:

- System Component Most Often Targeted. This includes the Protocol Stack, Operating System and Applications.
- Means of Attack. Consisting of: Input Validation Errors, Buffer Overflows, Boundary Condition Errors and other Malformed Input.
- Consequences of Attack. Where the result of the attack is manifested as a Denial of Service, Unothrozed Access (user or root), Loss of Confidentiality and Information Leakage resulting from a probe.
- Location of Attack. We categorize "Location of Attack" as Remote, Local, or Remote/Local.

Our ontology modeling the domain of computer attacks, is specified in the DARPA Agent Markup Language (DAML). DAML is a description logic language (DLL), which as a knowledge representation language is tailored for expressing knowledge about concepts and concept hierarchies and is well grounded in axiomatic and model semantics. DAML defines a number of constructs, such as intersection, union, and quantification, which can be used to define concepts and roles. DAML, like all description logics, supports classification and satisfiability, subsumption and instance checking.

At the top most level of the ontology, we define the classes *Host*, *State* and *Intrusion*. The relationship between Host and State is *Current State* while the relationship between Host and Intrusion is *Victim of*. It is important to note that while the ontology defines classes, properties, and their relationships, instances of the class *Intrusion* will only be instantiated if certain properties of the class *Current State* are found to exist.

Figure 1 presents a high level graphical illustration of our target-centric ontology. In the illustration, an ellipse denotes a subject and object while an arc represents the predicate (relationship).

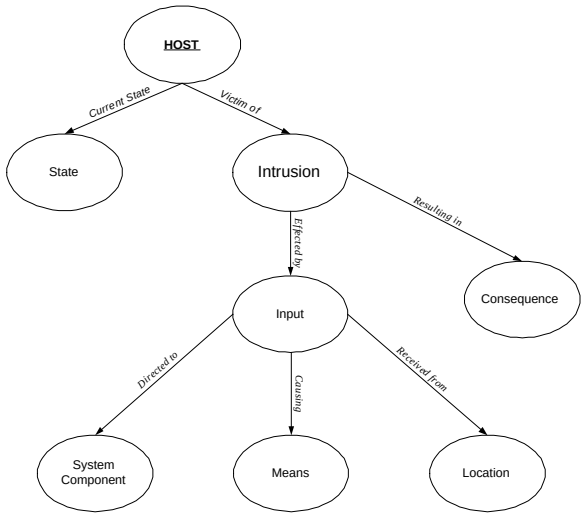

Figure 1. High Level Illustration of the Target-Centric Attack Ontology

Figure 2 depicts the class *Host* and its subclasses. The classes *Network*, *System* and *Process* are comprised of properties that reflect the current state of the particular host. For example, and as will be later demonstrated, the network class includes the properties *TCP_MAX* which defines the maximum number of TCP connections, *WAIT_STATE* defining the number of connections waiting on the final *ack* of the three-way

handshake required to establish a TCP connection, *THRESHOLD* which specifies the allowable ratio between maximum connections, and partially established connections and *EXCEED_T* a boolean value indicating that the allowable ratio has been exceeded. It should be noted that these are only four of several network properties.

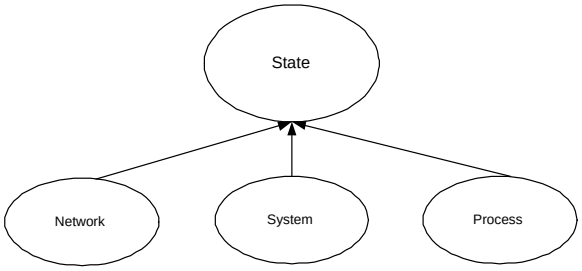

Figure 2. The State Class

Likewise, the class *Process* contains properties that measure conformance of that specific process (e.g.: sendmail, sshd, httpd, etc.) to its baselined profile. The class *System* is comprised of properties such as memory usage, processor load, etc., that describe the overall system state.

The class *Intrusion* is comprised of the classes *Consequence* and *Input* where the relationship (predicate) between Intrusion and Consequence is *Resulting in*. Likewise, the predicate connecting Intrusion and Input is *Effected by*. In turn, the class Input is comprised of the classes *Component*, *Means* and *Location*. Similarly, the predicates defining the relationship between Input and Component is *Directed to*, Input and Means is *Causing*, and Input and Location is *Received From*.

Figure 3 illustrates the class *System Component* and all of its subclasses. The class *Protocol Stack* has the subclasses *IP, TCP* and *UDP* which inherit their properties directly from the class *Network* that is defined under *State*. Accordingly, the class *TCP* inherits the properties *WAIT_STATE*, *TCP_MAX*, *THRESHOLD*, and *EXCEED_T* with restrictions. (We address DAML restrictions shortly).

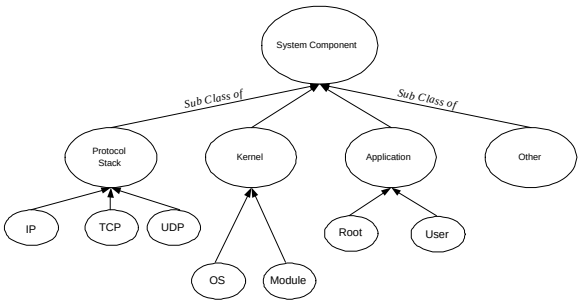

Figure 3. The System Component Class and its Subclasses

The class *Means* is comprised of the subclasses *Input Validation* and *Logic Exploit* both of which have subclasses associated with them. The subclass *Buffer Overflow* has a boolean attribute, inherited from the class *State*, providing it is true, that indicates if the instruction pointer points to an address within the process' stack frame. Figure 4 illustrates the class hierarchy rooted at the class *Means*.

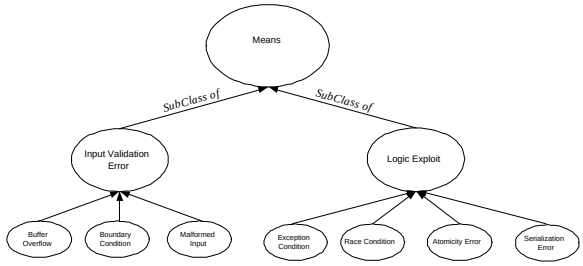

Figure 4. The Means Class and its Subclasses

Figure 5 illustrates the class Consequence and its subclasses. Although it is not illustrated, *Denial of Service* has a subclass called *Syn_Flood* which has a property that is a boolean value restricted by the property *EXCEED_T*. In DAML, a restriction means that a property of one class will be inherited by another class providing the property is of a specific value. We provide an example of this in the DAML listing in Section 6, where the class *Syn_Flood* inherits the property *Exceed_T* from the class *Network* if the value of *Exceed_T* is true.

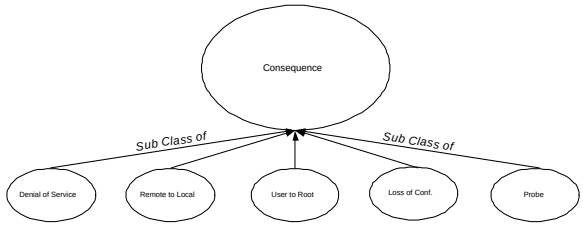

Figure 5. The Consequence Class and its Subclasses

Currently, we are focusing on TCP/IP traffic, hence the class *Location* is comprised of the subclasses *Local* and *Remote*. Figure 6 illustrates the class *Location* and its children. As was the case with the class *Protocol Stack* the class *TCP/IP* inherits some of its properties, with restrictions, from the class *Network*. Moreover, in the event that additional networking protocols were to be supported by the specific host, the ontology is easily extensible and additional protocols may added as subclasses.

6 Implementation

We have prototyped our ontology using *DAMLJessKB* [19], an extension to the *Java Expert System Shell* (JESS) [6]. JESS is a Java implementation of the *C Language Integrated Production System* (CLIPS) [7]. Accordingly, we use DAMLJessKB to reason over instances of our target centric ontology. Upon initialization we convert the DAML statements

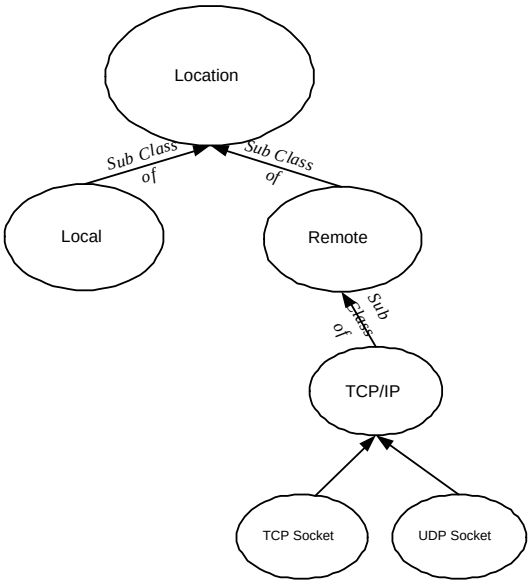

Figure 6. The Location Class and its Subclasses

representing the ontology into *N-Triples* and assert them into a knowledge base as rules. The assertions are of the form:

```
(assert
(PropertyValue (predicate) (subject) (object)))
```

Once asserted, DAMLJessKB it evaluates them and produces additional statement to include all of the chains of implication derived from the ontology.

The following is the DAML representation of a subset of our ontology. It consists of the classes Host, State, Network, Intrusion, Consequence, DoS, and Syn Flood and some of their properties:

```
<?xml version='1.0' encoding='ISO-8859-1'?>
   <!DOCTYPE rdf:RDF [
   <!ENTITY rdf "http://www.w3.org/1999/02/22-rdf-syntax-ns#">
   <!ENTITY IntrOnt "http://security.umbc.edu/IntrOnt#">
   <!ENTITY rdfs "http://www.w3.org/2000/01/rdf-schema#">
   <!ENTITY xsd "http://www.w3.org/2000/10/XMLSchema#">
   <!ENTITY daml "http://www.daml.org/2001/03/daml+oil#">
   <!ENTITY dex "http://www.daml.org/2001/03/daml+oil-ex#">
   <!ENTITY exd "http://www.daml.org/2001/03/daml+oil-ex-dt#">]>

<rdf:RDF xmlns:rdf="&rdf;"
   xmlns:IntrOnt="&IntrOnt;"
   xmlns:rdfs="&rdfs;"
   xmlns:daml="&daml;"
   xmlns:xsd="&xsd;"
   xmlns:dex="&dex;"
   xmlns:exd="&exd;"
   xmlns     ="http://www.daml.org/2001/03/daml+oil-ex#">


<daml:Class rdf:ID="BooleanValue">
  <daml:oneOf rdf:parseType="daml:collection">
    <BooleanValue rdf:ID="true"/>
    <BooleanValue rdf:ID="false"/>
   </daml:oneOf>
 </daml:Class>

<daml:Class rdf:about="&IntrOnt;Host"
     rdfs:label="Host">
    <rdfs:subClassOf rdf:resource="&rdfs;Resource"/>
</daml:Class>

<rdf:Property rdf:about="&IntrOnt;IP_Address"
     rdfs:label="IP_Address">
    <rdfs:domain rdf:resource="&IntrOnt;Host"/>
   <rdfs:range rdf:resource="&rdfs;Literal"/>
</rdf:Property>

<rdf:Property rdf:about="&IntrOnt;Current_state"
    rdfs:label="Current_state">
   <rdfs:domain rdf:resource="&IntrOnt;Host"/>
   <rdfs:range rdf:resource="&IntrOnt;State"/>
</rdf:Property>

<rdf:Property rdf:about="&IntrOnt;Target_of"
    rdfs:label="Target_of">
   <rdfs:domain rdf:resource="&IntrOnt;Host"/>
   <rdfs:range rdf:resource="&IntrOnt;Intrusion"/>
</rdf:Property>

<daml:Class rdf:about="&IntrOnt;State"
    rdfs:label="State">
   <rdfs:subClassOf rdf:resource="&rdfs;Resource"/>
</daml:Class>

<daml:Class rdf:about="&IntrOnt;Network"
   rdfs:label="Network">
  <rdfs:subClassOf rdf:resource="&IntrOnt;State"/>
</daml:Class>

rdf:Property rdf:about="&IntrOnt;TCP_Max"
   rdfs:label="TCP_Max">
  <rdfs:domain rdf:resource="&IntrOnt;Network"/>
  <rdfs:range rdf:resource="&rdfs;nonNegativeInteger"/>
</rdf:Property>

<rdf:Property rdf:about="&IntrOnt;Wait_State"
   rdfs:label="Wait_State">
  <rdfs:domain rdf:resource="&IntrOnt;Network"/>
  <rdfs:range rdf:resource="&rdfs;nonNegativeInteger"/>
</rdf:Property>

<rdf:Property rdf:about="&IntrOnt;Threshold"
   rdfs:label="Threshold">
  <rdfs:domain rdf:resource="&IntrOnt;Network"/>
  <rdfs:range rdf:resource="&rdfs;nonNegativeInteger"/>
</rdf:Property>

<rdf:Property rdf:about="&IntrOnt;Exceed_T"
 rdfs:label="Exceed_T">
<rdfs:domain rdf:resource="&IntrOnt;Network"/>
<rdfs:range rdf:resource="&IntrOnt;BooleanValue"/>
</rdf:Property>

<rdf:Property rdf:about="&IntrOnt;result_in"
   rdfs:label="restult_in">
  <rdfs:range rdf:resource="&IntrOnt;Consequence"/>
  <rdfs:domain rdf:resource="&IntrOnt;Intrusion"/>
</rdf:Property>

<rdfs:Class rdf:about="&IntrOnt;Consequence"
   rdfs:label="Consequence">
  <rdfs:subClassOf rdf:resource="&rdfs;Resource"/>
</rdfs:Class>

<rdfs:Class rdf:about="&IntrOnt;DoS"
   rdfs:label="DoS">
  <rdfs:subClassOf rdf:resource="&IntrOnt;Consequence"/>
</rdfs:Class>

<daml:Class rdf:about="&IntrOnt;Syn_Flood"
   rdfs:label="Syn_Flood">
  <rdfs:subClassOf rdf:resource="&IntrOnt;DoS"/>
     <rdfs:subClassOf rdf:resource="&IntrOnt;Network">
       <daml:Restriction>
          <daml:onProperty rdf:resource="&IntrOnt;Exceed_T"/>
          <daml:hasValue rdf:resource="&IntrOnt;true"/>
        </daml:Restriction>
       </rdfs:subClassOf>
</daml:Class>
```

As previously stated, our IDS model is distributed intrusion detection system that includes the use of data mining techniques to profile the quiescent state of a host [13] and [14]. Once a host has been profiled our IDS model calls for the continual sampling of network, process and system data, sub-

sequently comparing it to the profiled baseline. If we detect more than a marginal deviation between the sample and the baseline, the data, consisting of hundreds of low-level system parameters, are asserted into the knowledge base and reasoned over.

The following rule, when applied to the knowledge base, will indicate the occurrence of a denial of service attack if an instance of it or any of its subclasses exist.

```
(defrule isDOS

(PropertyValue
(p http://www.w3.org/1999/02/22-rdf-syntax-ns#type)
(s ?attack-id)
(o http://security.umbc.edu/Intrusion#DoS))

=>

(printout t ‘‘A DoS attack has occured.’’  crlf
            ‘‘with ID number: ‘‘ ?var))
```

Because of DAML's notion of classes and subclasses, a denial of service attack is inclusive of Syn Floods, Mailstorms, Pings of Death, or any other denials of service.

A DAML representation of an instance of a *Syn_Flood* attack is illustrated below. The first statement indicates that an event numbered 00035 has occurred which has the *resulting_in* property instantiated to an instance of a Syn Flood attack and is uniquely identified as 00038. Note: the following instance was produced by querying the knowledge base for instances of a denial of service attack, not for instances of a Syn Flood attack.

```
<Intrusion:Intrusion rdf:about="&IntrOnt;00035"
 Intrusion:IP_Address="130.85.112.231"
<Intrusion:resulting_in rdf:resource="&IntrOnt;00038"/>
</Intrusion:Intrusion>

<Intrusion:Syn_Flood rdf:about="&IntrOnt;00038"
 Intrusion:Exceed_T="true"
 Intrusion:time="20021212 154312">
</Intrusion:Syn_Flood>
```

Because our IDS model is distributed each individual IDS shares information with the other IDS's in its coalition. To that effect, the instance of the Syn Flood attack is sent to each IDS in the coalition where it is asserted into the corresponding knowledge bases.

7 Using the Ontology to Detect a Distributed Attack

The following example of a distributed attack illustrates the utility of our ontology.

The Mitnick attack is multi-phased; consisting of a Denial of Service attack, TCP sequence number prediction and IP spoofing. When this attack first occurred a Syn Flood was used to effect the denial of service, however any denial of service attack would have sufficed.

In the following example, which is illustrated in figure 7, **Host B** is the ultimate target and **Host A** is trusted by **Host B**.

The attack is structured as follows:

1. The attacker initiates a Syn/Flood attack against **Host A** to prevent **Host A** from responding to **Host B**.
2. The attacker sends multiple TCP packets to the target, **Host B** in order to be able to predict the values of TCP sequence numbers generated by **Host B**.
3. The attacker then pretends to be **Host A**, by spoofing **Host A**'s IP address, and sends a Syn packet to **Host B** in order to establish a TCP session between **Host A** and **Host B**.
4. **Host B** responds with a SYN/ACK to **Host A**. The attacker does not see this packet. **Host A**, since its input queue is full due to number of half open connections caused by the Syn/Flood attack, cannot send a *RST* message to **Host B** in response to the spurious Syn message.
5. Using the calculated TCP sequence number of **Host B** (recall that the attacker did not see the Syn/ACK message sent from **Host B** to **Host A**) the attacker sends an *Ack* with the predicted TCP sequence number packet in response to the *Syn/Ack* packet sent to **Host A**.
6. **Host B** is now in a state where it believes that a TCP session has been established with a trusted host **Host A**. The attacker now has a one way session with the target, **Host B**, and can issue commands to the target.

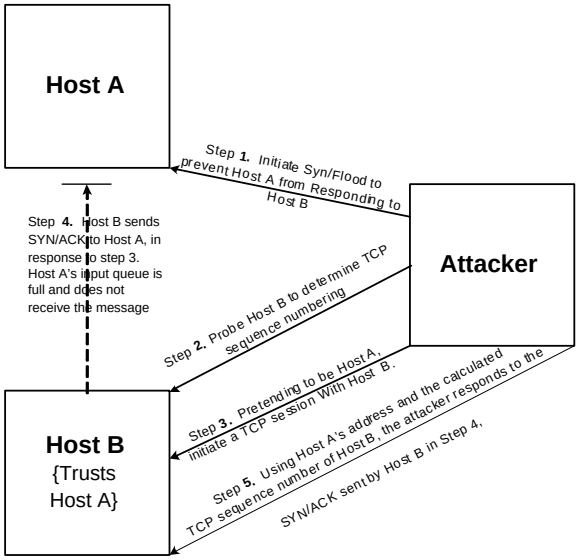

Figure 7. Illustration of the Mitnick Attack

It should be noted that an intrusion detection system running exclusively at either host will not detect this multi-phased and distributed attack. At best, Host A's IDS would see a relatively short lived Syn Flood attack, and Host B's IDS might observe an attempt to infer TCP sequence numbers, although this may not stand out from other non-intrusive but ill-formed TCP connection attempts.

The following explains the utility of our ontology, as well as the importance of forming coalitions of IDSs. In our IDS

model, we form coalitions of IDS services each of which is responsible for specific parts of an enterprise or domain. For example, one IDS service may be responsible for a specific host, while another is responsible for a group of hosts, while yet still another is responsible for monitoring network traffic. The IDS's all share a common ontology and utilize a secure communications infrastructure that has been optimized for IDS's. We present such a infrastructure in [15, 30, 31].

Consider the case of the instance of the Syn Flood attack presented in Section 6 and that it was directed against **Host A** in our example scenario. As the IDS responsible for **Host A** is continually monitoring for anomalous behavior, asserting and de-asserting data as necessary, it will detect the occurrence of an inordinate number of partially established TCP connections, and transmit the instance of the Syn Flood to the other IDS's in its coalition.

That instance is repeated below:

```
<IntrOnt:Intrusion rdf:about="&IntrOnt;00035"
     IntrOnt:IP_Address="130.85.112.231">
     <IntrOnt:resulting_in
               rdf:resource="&IntrOnt;00038"/>
</IntrOnt:Intrusion>

<IntrOnt:Syn_Flood rdf:about="&IntrOnt;00038"
     IntrOnt:Exceed_T="true"
     IntrOnt:int_time="20021212 154312"/>
```

This instance is converted into a set of *N-Triples* and asserted into the knowledge base of each IDS in the coalition. Those same*N-Triples* will be de-asserted when the responsible IDS transmits a message stating that the particular host is no longer the victim of a Syn Flood attack. This situation, especially in conjunction with **Host B** being subjected to a series of probes meant to determine its TCP sequencing, could be the prelude to a distributed attack the current connections and pending connections are also asserted into the knowledge base.

The following is a set DAML statements describing connections:

```
<IntrOnt:Connection rdf:about="&IntrOnt;00038"
 IntrOnt:IP_Address="130.85.112.231"
 IntrOnt:conn_time="20021212 154417"/>

<IntrOnt:Connection rdf:about="&IntrOnt;00101"
 IntrOnt:IP_Address="202.85.191.121"
 IntrOnt:conn_time="20021212 151221"/>

<IntrOnt:Connection rdf:about="&IntrOnt;00102"
 IntrOnt:IP_Address="68.54.101.78"
 IntrOnt:conn_time="20021212 150152"/>
```

In order to detect an Mitnick type attack, we include the following DAML statements that partially specify an ontology of the Mitnick attack (the class is identified as P_Mitnick for partial):

```
<daml:Class rdf:about="&Intrusion;P_Mitnick"
   rdfs:label="P_Mitnick">
    <daml:intersectionOf rdf:parseType=''daml:collection''>
      <daml:Class rdf:about="&IntrOnt;DoS"/>
      <daml:Class rdf:about="&IntrOnt;Connection"/>
    </daml:intersectionOf>
</daml:Class>
```

The ontology is partial because the Mitnick attack has the additional property that the connection time with the victim must be greater than or equal to the time of the denial of service attack. An instance of this ontology will be instantiated provided that there exists an instance of a denial of service attack that has the same unique identifier as that of an established connection. In fact there will be an instance created in each case where this condition holds. In our prototype, we check each instance to determine if the time of the connection is greater than or equal the time of the attack.

The following rules are used to check each instance:

```
(defrule isMitnick

(PropertyValue
(p http://security.umbc.edu/IntrOnt#P\_Mitnick )
(s ?eventNumber) (o "true"))

(PropertyValue
(p http://security.umbc.edu/IntrOnt#Int_time)
(s ?eventNumber) (o ?Int_Time))

(PropertyValue
(p http://security.umbc.edu/IntrOnt#Conn_time)
(s ?eventNumber) (o ?Conn_Time))

=>
(if (>=  ?Conn_Time ?Int_Time) then
(printout t ''event number: ''
?eventnumber '' is a Mitnick Attack: crlf)))
```

this rule will fire and event number 00038, the instance of the intersection of the connection and the denial of service attack, will be displayed.

At this point it is important to review the sequence of events leading up to the discovery of the Mitnick attack. Recall, that the IDS responsible for the victim of the Syn Flood attack queried its knowledge base for an instance of a *DoS* denial of service attack. The query returned an instance of a Syn Flood which was instantiated solely on the condition that the *Exced_T* property of the *Network* class was true.

The instance (its properties) of the Syn Flood attack was transmitted in the form of a set of DAML statements to the other IDS's in the coalition. In turn, these IDS's converted the DAML to a set of *N-Triples* and asserted them into their respective knowledge bases. As a Syn Flood is a precursor to a more insidious attack, instances of established and pending connections were asserted into the knowledge base. As the state of the knowledge base is dynamic due to the assertions and de-assertions, the rule set of each IDS is continually applied to the knowledge base.

The ontology specifying the Mitnick class states that it is the intersection of both the *DoS* and *Connection* classes. Because each IDS instantiates an instance when this constraints imposed by intersection is true, we need to examine each instance to ensure that $Connection\ Time \geq Intrusion\ Time$.

8 Conclusion and Future Work

We have analyzed vulnerability and intrusion data derived from CERT advisories and NIST's ICAT meta-base resulting in the identification of the components (network, kernel, application and other) most frequently attacked. We have also identified the most common means and consequences of the attack

as well as the location of the attacker. Our analysis shows that non-kernel space (non operating system) applications, running as either root or user, are the most frequently attacked and are attacked remotely. The most common means of attack are exploits. According to the CERT advisories issued in response to severe vulnerabilities, *root* access is the most common consequence of an exploit whereas the ICAT data shows *denial of service* to be the most common consequence.

Our analysis was conducted in order to identify the observable and measurable properties of computer attacks and intrusions. Accordingly, we have developed a target-centric ontology characterized by *System Component*, *Means of Attack*, *Consequences of Attack* and *Location of Attacker*. We have stated the case for replacing simple taxonomies with ontologies for use in IDS's and have presented an initial ontology specifying the class *Intrusion*. Our ontology is available at: http://security.cs.umbc.edu/Intrusion.

We have prototyped our ontology using the DAMLJessKB, which has some limitations. We intend to either modify DAMLJessKB in order to make it a full and complete reasoner or use Stanford's *Java Theorem Prover* [5] or *Rename ABox and Concept Expression Reasoner* [10].

References

[1] J. Allen, A. Christie, W. Fithen, J. McHugh, J. Pickel, and E. Stoner. State of the Practice of Intrusion Detection Technologies. Technical Report 99tr028, Carnegie Mellon - Software Engineering Institute, 2000.

[2] E. G. Amoroso. *Fundamentals of Computer Security Technology*. Prentice-Hall PTR, 1994.

[3] D. Curry and H. Debar. Intrusion detection message exchange format data model and extensible markup language (xml)document type definition. draft-ietf-idwg-idmef-xml-07.txt, June 2002. expires December 19, 2002.

[4] R. Davis, H. Shrobe, and P. Szolovits. What is knowledge representation? *AI Magazine*, 14(1):17 – 33, 1993.

[5] G. Frank, J. Jenkins, and R. Fikes. Jtp: An object oriented modular reasoning system. http://kst.stanford.edu/software/jtp.

[6] E. J. Friedman-Hill. Jess, the java expert system shell. http://herzberg.ca.sandia.gov/jess/docs/52/, November 1977.

[7] J. Giarratano and G. Riley. *Expert Systems Principles and Programming*. PWS Publishing Company, third edition, 1998.

[8] R. L. Glass and I. Vessey. Contemporary application-domain taxonomies. *IEEE Software*, pages 63 – 76, July 1995.

[9] B. Guha and B. Mukherjee. Network Security via Reverse Engineering of TCP Code: Vulnerability Analysis and Proposed Solutions. In *IEEE Networks*, pages 40 – 48. IEEE, July/August 1997.

[10] V. Haarslev and R. Moller. Racer. http://www.cs.concordia.ca/ faculty/haarslev/racer.

[11] J. W. Haines, L. M. Rossey, R. P. Lippman, and R. K. Cunningham. Extending the darpa off-line intrusion detection evaluations. In *DARPA Information Survivability Conference and Exposition II*, volume 1, pages 77 – 88. IEEE, 2001.

[12] J. Hendler. DARPA Agent Markup Language. http://www.daml.org, 2000.

[13] A. Joshi and J. Undercoffer. On web semantics and data mining: Intrusion detection as a case study. In *Proceedings of the National Science Foundation Workshop on Next Generation Data Mining*, 2002.

[14] A. Joshi and J. Undercoffer. *Data Mining, Semantics and Intrusion Detection: What to dig for and Where to find it*, chapter tbd. MIT Press, 2003.

[15] L. Kagal, J. Undercoffer, A. Joshi, and T. Finin. Vigil: Enforcing Security in Ubiquitous Environments . In *Grace Hooper Celebration of Women in Computing 2002*, 2002.

[16] R. A. Kemmerer and G. Vigna. Intrusion detection: A brief history and overview. *Security and Privacy a Supplement to IEEE Computer Magazine*, pages 27 – 30, April 2002.

[17] K. Kendall. A database of computer attacks for the evaluation of intrusion detection systems. Master's thesis, MIT, 1999.

[18] D. Koller and A. Pfeffer. Probabilistic Frame-Based Systems. In *Proceedings of the Fifteenth National Conference on Artifical Intelligence*, pages 580 – 587, Madison, Wisconsin, July 1998. AAAI.

[19] J. Kopena. Damljesskb. http://edge.mcs.drexel.edu/assemblies/software/damljesskb/articles/DAMLJessKB-2002.pdf, October 2002.

[20] I. Krusl. *Software Vulnerability Analysis*. PhD thesis, Purdue, 1998.

[21] C. E. Landwehr, A. R. Bull, J. P. McDermott, and W. S. Choi. A taxonomy of computer program security flaws. *ACM Computing Surveys*, 26(3):211 – 254, September 1994.

[22] U. Lindqvist and E. Jonsson. How to systematically classify computer security intrusions. In *Proceedings of the 1997 IEEE Symposium on Security and Privacy*, pages 154 – 163. IEEE, May 1997.

[23] R. Lippmann, D. Fried, I. Graf, J. Haines, K. Kendall, D. McClung, D. Weber, S. Webster, D. Wyschogrod, R. Cunningham, and M. Zissman. Evaluating intrusion detection systems: The 1998 darpa off-line intrusion detection evaluation. In *Proceedings of the DARPA Information Survivability Conference and Exposition,2000*, pages 12 – 26, January 2000.

[24] J. McHugh. Testing Intrusion Detection Systems: A Critique of the 1998 and 1999 DARPA Intrusion Detection System Evaluations as Performed by Lincoln Laboratory. *ACM Transactions on Information and System Security*, November 2000.

[25] P. Ning, S. Jajodia, and X. S. Wang. Abstraction-based intrusion in distributed environments. *ACM Transactions on Information and Systems Security*, 4(4):407 – 452, November 2001.

[26] N. F. Noy and D. L. McGuinnes. Ontology development 101: A guide to creating your fisrt ontology. Stanford University.

[27] V. Raskin, C. F. Hempelmann, K. E. Triezenberg, and S. Nirenburg. Ontology in information security: A useful theoretical foundation and methodological tool. In *Proceedings of NSPW-2001*, pages 53 – 59. ACM, ACM, September 2001.

[28] RDF. Resource description framework (rdf) schema specification, 1999.

[29] G. G. Sumpson. *Principals of Animal Taxonomy*. Columbia University Press, 1961.

[30] J. Undercoffer, F. Perich, A. Cedilnik, L. Kagal, and A. Joshi. Centarusu2: A Secure Infrastructure for Service Discovery and Delivery in Pervasive Computing. *Mobile Networks and Applications: Special Issue on Security*, 2003.

[31] J. Undercoffer, F. Perich, and C. Nicholas. Shomar: An architecure for distributed intrusion detection services. University of Maryland Baltimore County, Department of Computer Science and Electrical Engineering, 2002.

[32] i. WEBSTERS, editor. *Merriam-Webster's Collegiate Dictionary*. Merriam-Webster, Inc., tenth edition, 1993.

[33] C. Welty. Towards a semantics for the web. Vassar College, 2000.

[34] M. Wood and M. Erlinger. Intrusion detection message exchange requirements. draft-ietf-idwg-requirements-08, August 2002.