

Investigating the Dark Cyberspace: Profiling, Threat-Based Analysis and Correlation

Claude Fachkha, Elias Bou-Harb, Amine Boukhtouta, Son Dinh, Farkhund Iqbal, Mourad Debbabi

NCFTA Canada and Concordia Institute for Information Systems Engineering
Concordia University, Montreal, Quebec, Canada
{c_fachkh, e_bouh, a_boukh, so_din, iqbal_f, debbabi}@encs.concordia.ca

Abstract—An effective approach to gather cyber threat intelligence is to collect and analyze traffic destined to unused Internet addresses known as darknets. In this paper, we elaborate on such capability by profiling darknet data. Such information could generate indicators of cyber threat activity as well as providing in-depth understanding of the nature of its traffic. Particularly, we analyze darknet packets distribution, its used transport, network and application layer protocols and pinpoint its resolved domain names. Furthermore, we identify its IP classes and destination ports as well as geo-locate its source countries. We further investigate darknet-triggered threats. The aim is to explore darknet embedded threats and categorize their severities. Finally, we contribute by exploring the inter-correlation of such threats, by applying association rule mining techniques, to build threat association rules. Specifically, we generate clusters of threats that co-occur targeting a specific victim. Such work proves that specific darknet threats are correlated. Moreover, it provides insights about threat patterns and allows the interpretation of threat scenarios.

I. INTRODUCTION

Today, the safety and security of our society is entirely dependent on having a secure infrastructure. This infrastructure is largely controlled and operated using cyberspace; the electronic world created by interconnected networks of information technology and the information on those networks. Recent events demonstrate that cyberspace could be subjected, at the speed of light and in full anonymity, to severe attacks with drastic consequences. One particular research revealed that 90% of United States companies have been the target of a cyber attack, with 80% suffering a significant financial loss [1]. In addition, the Canadian's cyber security strategy report [2] highlighted the fact that Canadian households and businesses have completely embraced the cyberspace. The report as well elaborated that in a recent one year period, 86% of large Canadian organizations had suffered a cyber attack where the loss of intellectual property as a result of these attacks doubled between 2008 and 2010. Moreover, the report alarmed that more than 60% of all the malicious code ever detected, originating from more than 190 countries, was introduced into cyberspace solely in 2010. In this context, cyber threat intelligence generation is of vital importance. A promising approach to gather cyber threat intelligence is to collect and analyze darknet traffic. The term *darknet* could be defined as one of the following:

- A set of unallocated network addresses and communication ports that belong either to the public cyberspace or to a specific organization. Such unallocated space could be maliciously utilized to launch cyber attacks.
- Untraceable and inaccessible cyberspace material that is concealed from web services and search engines.
- Concealed communication platforms (e.g., P2P) including social networks, chat channels and file-sharing environments.

In this paper, we refer to darknets as per the first definition; since these network addresses correspond to illegitimate hosts or devices, any observed traffic destined to them may be suspicious and hence need to be investigated. Therefore, there is a requirement to answer the following set of high-level questions:

- 1) What is the nature of darknet traffic and its underlying content?
- 2) Who contributes to darknet traffic?
- 3) Are there any embedded darknet threats?
- 4) Can we show that such threats are correlated and hence provide their real world interpretation and impact?

In this work, we answer such questions. Specifically, we profile darknet traffic (e.g., protocols distributions, type of traffic, IP classes, sources, etc.). Such information generates indicators of cyber threat activity as well as it provides an in-depth understanding of the nature of such traffic. We, as well, explore darknet-triggered infections/intrusions and geo-locate their corresponding sources. Moreover, we generate, analyze and interpret threat association rules by applying association rule mining techniques. Such work demonstrates that certain darknet threats are correlated when targeting specific victims. Moreover, it provides insights about threat patterns and allows the interpretation of threat scenarios. The rest of the paper is organized as follows. Section II discusses the related work. Section III presents our profiling results while Section IV demonstrates the results of darknet embedded threats. Moreover, Section V elaborates on the association rule mining approach for the purpose of generating and interpreting clusters of threats that co-occur targeting a specific victim (network destination). Finally, Section VI summarizes our contributions and discusses future work.

II. RELATED WORK

Several studies explored darknet traffic analysis. We can classify these proposals into two main categories. The first category is based on designing, implementing and managing darknet platforms, while the second focuses on the analysis of darknet traffic feeds.

In the following, we describe some of the projects in the area of darknet monitoring systems. In [3], the author presented Honeyd as a framework for the deployment of honeypots using virtual machines. This project runs on unallocated addresses within various operating systems. Such environments provide numerous services which aid in detecting and mitigating worms, preventing spam distribution and alerting about suspicious attacks. Another project is the network telescope which was proposed in [4] to monitor cyber incidents through the dark address space. Moreover, the Internet Motion Sensor (IMS) system, a distributed system, described in [5], reports the network behavior originating from different monitored IP blocks. Furthermore, Yegneswaran et al. [6] developed Internet Sink (iSink) to monitor unused IP address space. The iSink approach was conceived to address the scalability issue that is related to large address spaces. It incorporates passive detection and monitor sensors as well as honeynet components.

In the other category, namely darknet analysis, the research in [7] elaborated on a detailed analysis of the darknet data. Their active and passive analyses assessed darknet samples from different networks and over a long time period. Another study [8] has reviewed the last mentioned work to render the state of this Internet background radiation at that current year. The authors observed significant changes and pinpointed several factors that are behind these measures. Moreover, Fukuda et al. [9] studied correlations among darknet traffic for estimating their behaviors through small address blocks by analyzing a specific type of traffic packets (i.e., TCP SYN). There are other research proposals that investigated threats triggered through darknets such as in [10] where the authors were able to study the Slammer worm. Moreover, denial of service (DoS) attacks were as well addressed in [11] by analyzing the replies of DoS attacks from spoofed sources in darknet feeds. Other studies such as [12] elaborated on scanning events, misconfiguration and other suspicious activities.

The work presented in this paper, which belongs to the second category, contributes in the following three aspects:

- Analysis accuracy: The analyzed darknet data includes packet types that were omitted by other research works (e.g., ICMP in [13]). As such, the data set is rich which contributes to a better accuracy of the analysis.
- Threat analysis: By adopting an analysis methodology based on the use of network intrusion detection systems (NIDSs), our approach yields real world threats that are embedded in darknet traffic. Such results will be presented in Section IV.
- Association rule mining approach: By applying association rule mining and correlation techniques on the threat

data, we investigate clusters of threats that co-occur. Such cyber threat intelligence proves that specific threats are correlated in addition to providing better understating by interpreting the attack scenarios targeting specific network destinations. To the best of our knowledge, such work has never been previously investigated.

III. DARKNET PROFILING

To achieve a better understanding of the nature of darknet traffic and its underlying content and threats, we have performed darknet traffic profiling. To accomplish this task, we analyzed some darknet data collected in the period between September 16th, 2011 and May 9th, 2012. The analyzed data feeds are retrieved in real-time from a trusted third party framework. The data consists of pure darknet traffic collected from many /16 address blocks. For more information regarding the source of data, please contact the authors.

We initiated our analysis by differentiating darknet packets according to their types following the method in [14]:

- Scanning traffic; TCP SYN packets
- Backscattering traffic, which commonly refers to unsolicited traffic that is the result of responses to attacks with spoofed source IP address; TCP SYN+ACK, RST, RST+ACK, and ACK packets
- The remaining traffic packets are classified as misconfiguration

Table I depicts the outcome distribution. These results reveal

Scanning Traffic	Backscattering	Misconfiguration
68.02%	2.00%	29.98%

TABLE I: Packets Distribution - Nature of Traffic

that scanning or network probing constitutes the majority of darknet traffic. Note that, such traffic could be interpreted as an indication of port scanning and/or vulnerability probing. Such attacks, in general, are preliminary triggered before launching a targeted attack towards a specific system. We next aimed to identify the major protocols that are used in darknet traffic. Table II provides the percentages of darknet transport and network layer protocols. It is observed that TCP plays the major role.

TCP	UDP	ICMP	Others
91.9%	5.5%	2.9%	0.3%

TABLE II: Protocols Distribution

Figure 1 corroborates this fact by plotting the protocols distribution in a day sample, which is the average of daily samples collected over a month's period. TCP dominance can be explained by two facts: First, is that the majority of scanning attacks use TCP and second is that there exist known attacks that specifically target TCP ports as noted in [7]. TCP increase in Figure 1, especially after the 12th hour, indicates




Fig. 1: Darknet Network and Transport Layer Protocols

that the darknet sensors record an increasing number of TCP packets after that period.

Such information pinpoints the need for a thorough temporal analysis and comparison of that phenomena which may uncover and explain the occurrence of certain attacks at specific time periods and their absence during other periods at any given day. Next we profiled darknet application protocols. Figure 2 illustrates the top 16 application protocols that have been found. The results demonstrate that the Session Initiation Protocol (SIP) is leading while the Domain Name Service is ranked second and NetBIOS is ranked third. It is worthy to note, that the SIP protocol is excessively used in DoS attacks, specifically against voice over IP (VoIP) servers [15], and thus its appearance as a top darknet application protocol is significant and maybe alarming. We further studied the source

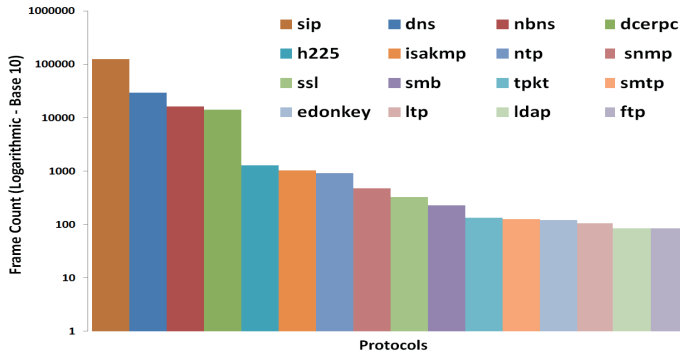


Fig. 2: Darknet Application Layer Protocols

and destination distributions of IP classes in the darknet traffic. Table III depicts the results.

Class	Usage (%)	
	Source	Destination
A	62.529	0.017
B	18.529	7.138
C	18.942	92.845

TABLE III: IP Class Distribution

It is revealed that the majority of source IPs belong to class 'A', whereas in the case of destination IPs, class 'C'

plays the major role. Furthermore, Class 'A' proportion in the destination IPs is almost negligible, i.e., 0.017% whereas class 'B' contributes relatively more. It is substantial to mention that class 'C', being the most destined and smallest class, could be an indication that it is as well the most targeted class by cyber attacks and hence further investigation in it could yield relevant cyber intelligence. Moreover, we were interested in identifying the resolved domain names in darknets. After performing this task, we identified that the top-most darknet resolved domain belongs to a .cc Internet country code top-level domain for Cocos (Keeling) Islands. Note that, this domain, according to the anti-phishing working group, constituted a significant 7.3% of all phishing attacks detected in 2010 [16]. Similar results could feed us, in general, with relevant information about unsolicited/malicious domains that could be used by attackers. Another analysis has been performed on the TCP and UDP ports that are used in the collected darknet traffic. Specifically, we aimed to pinpoint the destination ports. Such insights could reveal the targeted ports used in cyber attacks. Figures 3 and 4 illustrate such results.




Fig. 3: Darknet TCP Targeted Ports




Fig. 4: Darknet UDP Targeted Ports

The top three destination darknet TCP ports, namely, ports 445, 80, and 3389 are the Microsoft active directory service, the hypertext transfer protocol, and the Microsoft terminal server respectively. These service ports have previously suffered from security issues and vulnerabilities. A sample of the threats targeting such services are pinpointed in [17], [18] and [19] respectively. Hence, it is alarming that such ports




Fig. 5: Darknet Sources - Heat Map

appear as the top darknet destination TCP ports. On the other hand, the top three destination darknet UDP ports, namely, ports 5060, 397, 1280 are the SIP, the multi-protocol transport network (mptn) service, and the pictography protocol respectively. The SIP, as mentioned in Section III, is a significant target of attack. This result further validates the integrity of our insights. Moreover, the mptn and the pictography services are known to suffer from denial of service attacks when a malformed request is destined to them. For the purpose of pinpointing the sources that contribute to the darknet traffic, we perform darknet geo-localization. Figure 5 depicts the heat map. According to our analysis, the source countries reached 196 countries where the majority of source IPs are located in USA. It is as well noticeable that Brazil, China, and Russia represent the major portion of source IPs compared to other countries. Note that, in Section IV, when we reveal the darknet threat analysis and geo-locate the sources behind those threats, the three aforementioned countries as well appear amongst the top contributed threat countries.

IV. THREAT ANALYSIS

In this section, we extend our profiling task. The aim is to uncover real world threats that are embedded in darknet traffic in addition to categorizing their severities and geo-locating their sources. For that purpose, we executed threat-based severity analysis. To accomplish this task, Snort [20] and Bro [21], two open source NIDSs, combining the benefits of signature, protocol and anomaly-based inspection, were implemented and utilized. Part of their content signature detection, Snort and Bro implement the Boyer-Moore exact string matching detection algorithm in addition to a non-deterministic finite automata regular expression detection algorithm. To perform the threat analysis, we configured the NIDSs with rule sets from the Sourcefire Vulnerability Research Team and The Bro Network Security Monitor. Consequently, we fed the darknet data to the NIDSs. A partial outcome of this procedure is summarized in Table IV. The results reveal 30 distinct threats. According to the NIDSs, three threats are of high priority, two are of medium severity and the rest are of low priority. The first high priority threat (t_1) is in fact an attempt to possibly overflow a buffer. Specifically, a series of NOOP (no operation instructions) were found in the data stream. Typically, most buffer overflow exploits use NOOPs sleds to pad the code [22]. Hence, this threat might indicate an attempt to use a buffer overflow exploit. Thus, a full compromise of

a system is possible if the exploit is successful. Another high priority threat (t_2) is rendered as an attempt to cause a DoS. Particularly, a heap-based buffer overflow in Microsoft MSN Messenger [23] was found that allows user-assisted remote attackers to execute arbitrary code via unspecified vectors involving video conversation handling in Web Cam and video chat sessions. As a result, DoS and complete administrator access to a targeted system is possible. The last high priority threat (t_3) is in reality a detected virtual private network (VPN) remote attempt on a set of darknet addresses. Although, in general, VPN is not considered a threat, however an attempt to gain VPN access on a specific system can be alarming. On the other hand, threats t_4 and t_5 , and according to the NIDSs, are of medium severity. Threat (t_4) represents an attempt to use a traceroute software where an attacker can discover live hosts and routers on a target network in preparation for an attack. Moreover, (t_5) is a portmap GETPORT request to discover the port where the Remote Procedure Call (RPC) `statd` is listening [24]. An attacker can query the portmapper to discover the port where `statd` runs. Consequently, this may be a precursor to accessing `statd`. The remaining of the incidents are mainly scanning attempts and are considered of low severities. Although their techniques may vary, their end goal is to either perform port scanning or vulnerability probing in preparation to a possible targeted attack. It is very significant to note, for the purpose of results integrity, that such scanning attempts, that constitute the majority of the threats, are in accordance with our darknet profiling results, specifically the packets distribution - nature of traffic percentages (68.02%) that was demonstrated in Table I in Section III. For the purpose of accomplishing high level attribution, we perform geo-location of the threats sources. Figure 6 depicts the heat map. Note that, the threat count metric is of the order of thousands. The results reveal that Russia and China lead in terms of number of darknet threats.

Threat	Type	Priority
t_1	Buffer Overflow Exploit	High
t_2	Denial of Service	
t_3	VPN Attempt	
t_4	Traceroute Utilization	Medium
t_5	Service Port Discovery	
t_{6-30}	Scanning Attempts	Low

TABLE IV: Darknet Threats and Corresponding Severities




Fig. 6: Threats Sources - Heat Map (in thousands)

V. THREATS CORRELATION

There is a crucial need to further analyze the threats that have been previously detected and discussed. This section explores the inter-correlation of such threats, by applying association rule mining techniques, to build threat association rules. Such work demonstrates that specific darknet threats are in fact correlated or co-occur when targeting specific victims. Moreover, it provides insights about threat patterns and allows the interpretation of threat scenarios.

A. Approach

The goal is to investigate the interdependence and inter-correlation of darknet threats. Particularly, we aim to answer the following questions: Are there any threats targeting a specific victim that follow a certain pattern? Moreover, if some of the co-occurring threats appear in a darknet traffic, how confidently one can predict the existence of other threats? To investigate this, we employed the technique of frequent pattern mining (frequent item-set) and association rule mining [25]. Another outcome of this approach besides the ones mentioned above, is the generation of threat association rules that could be used as an input to a classification model that is able to predict and hence mitigate future threat occurrences. Frequent pattern and association rule mining techniques have been proven to be very successful for identifying hidden patterns in DNA sequences, customer purchasing habits, text categorization, and many other applications of pattern recognition. The proposed threat correlation approach is a three-step process, namely, frequent pattern mining, association rule generation from each frequent threat-set, and rule analysis by applying various correlation techniques. Each of these steps is detailed below.

B. Frequent Pattern Mining

An item-set or a pattern is a group of two or more objects that appear together. An item-set is a *frequent* pattern if its members appear together for some minimum number of times. In the context of threat analysis, an item or an object is a threat and an item-set is the threat-set. Table V, which is used for illustration and explanation purposes, depicts 10 threat-sets, one threat-set per row. Let $T = \{t_1, \dots, t_m\}$ denote the universe of all threats detected from the given darknet feeds F . Suppose a threat-set $T_i \subseteq T$ detected at a time interval τ_i represents a row or an instance in the threat Table V.

Time Intervals	Identified Threats
τ_1	$\{t_2, t_5, t_7, t_9\}$
τ_2	$\{t_2, t_5, t_7\}$
τ_3	$\{t_2, t_5\}$
τ_4	$\{t_1, t_5, t_7\}$
τ_5	$\{t_4, t_5, t_7\}$
τ_6	$\{t_3, t_6, t_8\}$
τ_7	$\{t_4, t_5, t_8\}$
τ_8	$\{t_3, t_6, t_8\}$
τ_9	$\{t_2, t_5, t_8\}$
τ_{10}	$\{t_1, t_5, t_7, t_8, t_9\}$

TABLE V: Vectors of Darknet Threats

This table shows ten threat-sets captured at time intervals $\{\tau_1, \dots, \tau_{10}\}$. Let $T_i \subseteq T$ be a threat-set or a pattern in the threat table. A pattern that contains k threats is a k -*pattern*. For instance, $\tau_1 = \{t_2, t_5, t_7, t_9\}$ is a 4-*pattern*. Similarly, the support of a pattern T_i is the percentage of all the instances T in the threat table containing T_i , denoted by $support(T_i|T)$. Note that, the probability $P(t_a \cup t_b)$, where $t_a \cup t_b$ indicates that a pattern contains both t_a and t_b , is the union of itemsets t_a and t_b . The support is defined in equation 1:

$$support(t_a \Rightarrow t_b) = P(t_a \cup t_b) \quad (1)$$

A pattern T_i is a *frequent pattern* if the support of T_i is greater than or equal to some user specified minimum support threshold, which is a real number in an interval of $[0, 1]$. Further explanation of these terms is given in Example 5.1.

Example 5.1: Consider Table V. Suppose the user-specified threshold $min_sup = 0.3$, which means that a pattern $T_i = \{t_1, \dots, t_k\}$ is frequent if at least 3 out of the 10 rows contain all threat-items in T_i . For instance, $\{t_2, t_5, t_7, t_9\}$ is not a frequent pattern because it has support $1/10 = 0.1$. Similarly, $\{t_2, t_5\}$ is a frequent 2-pattern because it has support $4/10 = 0.4$ and contains two threats. Likewise, $\{t_5, t_8\}$ is a frequent 2-pattern with support $3/10 = 0.3$.

There are various data mining algorithms for extracting frequent patterns, such as the Apriori [25], FP-growth [26], and ECLAT [27]. In this work, we employ the Apriori algorithm since it is easy to comprehend and it has been validated in several text mining studies [28]. Below, we provide an overview of the Apriori algorithm. Apriori is a level-wise iterative search algorithm that uses frequent k -patterns to explore the frequent

$(k + 1)$ -patterns. First, the set of frequent 1-patterns is found by scanning the threat table, accumulating the support count of each threat-set, and collecting the threat patterns containing T that also contains T_i with $support(T_i|T) \geq min_sup$. The resulting frequent 1-patterns are then used to find frequent 2-patterns, which are then used to find frequent 3-patterns, and so on, until no more frequent k -patterns can be found. The generation of frequent $(k+1)$ -pattern from frequent k -patterns is based on the following Apriori property.

Property 5.1 (Apriori property): All nonempty subsets of a frequent pattern must be frequent.

By definition, a pattern T'_i is not frequent if $support(T'_i|T) < min_sup$. The above property implies that adding a threat t to a non-frequent pattern T'_i will not make it frequent. Thus, if a k -pattern T'_i is not frequent, then there is no need to generate $(k + 1)$ -pattern $T'_i \cup T$ because $T'_i \cup T$ is also not frequent. The following example shows how the Apriori algorithm exploits this property to efficiently extract all frequent patterns or threat-sets. For a formal description, we refer the reader to [25].

Example 5.2: Consider Table V with $min_sup = 0.3$. First, identify all frequent 1-patterns by scanning the threat table once to obtain the support of every threat-set. The items having support ≥ 0.3 are frequent 1-patterns, denoted by $L_1 = \{\{t_2\}, \{t_5\}, \{t_7\}, \{t_8\}\}$. Then, join L_1 with itself, i.e., $L_1 \bowtie L_1$, to generate the candidate set $C_2 = \{\{t_2, t_5\}, \{t_2, t_7\}, \{t_2, t_8\}, \{t_5, t_7\}, \{t_5, t_8\}, \{t_7, t_8\}\}$ and scan the threat table once to obtain the support of every pattern in C_2 . Identify the frequent 2-patterns, denoted by $L_2 = \{\{t_2, t_5\}, \{t_5, t_7\}, \{t_5, t_8\}\}$. Similarly, perform $L_2 \bowtie L_2$ to generate $C_3 = \{t_5, t_7, t_8\}$. By scanning the threat table once, we found that $\{t_5, t_7, t_8\}$ is not frequent, i.e., 3-pattern L_3 is empty. The finding of each set of frequent k -patterns requires one full scan of the rows in Table V.

C. Association Rule Mining

The selected frequent patterns or frequent threat-sets are used to investigate the correlation and interdependence of the subsets of each frequent threat-set. This can be achieved by applying association rule mining techniques [29]. For this, all 1-patterns are deleted as they contain only one threat and thus can not be associated with any other threat. The 2-patterns threat-sets are used to extract single-dimensional association rules while the 3-patterns and higher patterns are used to construct multi-dimensional association rules. To construct an association rule of threats, we need to calculate the confidence for each frequent threat-set. The confidence is the percentage of threat-sets containing threat Y in addition to threat X with regard to the overall number of threat-sets containing X . Assume we have a threat-set $\{t_a, t_b\}$ for which the association rule would be $\{t_a\} \Rightarrow t_b$. Hence, the association rule has a confidence c in the threat table T , where P is the probability and c is the percentage of threat-sets in T containing t_a that

also contains t_b .

This statement is mathematically expressed in Equation 2.

$$confidence(t_a \Rightarrow t_b) = P(t_b|t_a) = \frac{support\{t_a \cup t_b\}}{support\{t_a\}} \quad (2)$$

Having support-count of $(t_a \cup t_b)$ and t_a , we can calculate $confidence(t_a \Rightarrow t_b)$ using Equation 2. Once the frequent threat-sets are extracted, the related association rule of a frequent threat-set T_i can be constructed as follows:

- Generate all non-empty subsets of T_i
- For every non-empty subset S , construct a rule $(S \Rightarrow (T_i - S))$, provided that the $\frac{support(T_i)}{support(S)} \geq min_conf$

D. Correlation Analysis

In order to investigate the interdependency of the threats, various correlation techniques including χ^2 , cosine measure, and lift [29] can be used. In the current study, we use lift, which is easy to understand since it is based on probabilities and its results are interpretable by non-technical domain experts without the help of data mining experts. The correlation technique lift measures how many times more often threats t_a and t_b occur together than expected if they are statistically independent. The lift indicates whether the identified threat patterns are correlated together. It is mathematically expressed as follows:

$$lift(t_a, t_b) = \frac{P(t_a \cup t_b)}{P(t_a)P(t_b)} \quad (3)$$

If the value of Equation 3 is equal to 1 then threats t_a and t_b are independent and therefore have no correlation; otherwise they are either negatively correlated (i.e., $lift < 1$) or positively correlated (i.e., $lift > 1$).

E. Experimental Results

We used Weka [30] to run the Apriori algorithm. In summary, the Apriori takes the threat table in ARFF file type as input along with the user-defined parameters including minimum support min_sup and confidence c , and generates association rules. To assess our approach, we experimented with different threats that were detected and mentioned in Table IV. The experimental results are achieved by employing sequential rule mining techniques for correlating same set of threats. Consequently, the generated rules can be used to build an associative classification model for predicting the occurrences of specific threats in real-time darknet traffic. In general, the threat rules generated by the Apriori, provided the threshold is kept low, is usually very large. However, we can tune and filter the results to bring the rules to a manageable level by applying the following steps:

- Choosing a suitable value for the minimum support based on the occurrence count of the targeted threat. Note that, the choice of selecting a minimum support threshold is inversely proportional to the number of generated threat-sets.
- Taking into consideration the size of the association rules by specifying the number of items per threat-set as input to the algorithm.

Darknet Feed Providers	Analyzed Address Blocks	Association Rules	Confidence	Lift	Count
Destination Network 1	w1.x1.y1.z1/24	1. $\{t_7, t_8, t_9\} \Rightarrow t_{10}$	0.63	3.64	282
		2. $\{t_{10}, t_{14}, t_{13}\} \Rightarrow t_{11}$	0.56	7.06	306
Destination Network 2	w2.x2.y2.z2/24	3. $\{t_{10}, t_{15}, t_4\} \Rightarrow t_1$	0.76	1.54	193
		4. $\{t_{12}, t_{11}, t_{13}\} \Rightarrow t_{10}$	0.92	3.81	359
Destination Network 3	w3.x3.y3.z3/24	5. $\{t_{10}, t_7, t_8, t_9, t_{13}\} \Rightarrow t_4$	0.55	10.75	218
		6. $\{t_{10}, t_8, t_9, t_{13}\} \Rightarrow t_{12}$	0.26	3.68	348
Destination Network 4	w4.x4.y4.z4/24	7. $\{t_7, t_8, t_9\} \Rightarrow t_{10}$	0.43	4.12	113
		8. $\{t_4, t_8, t_9\} \Rightarrow t_{10}$	0.98	6.6	102
Destination Network 5	w5.x5.y5.z5/24	9. $\{t_{10}, t_7, t_8, t_9, t_{13}\} \Rightarrow t_{11}$	0.41	3.56	260
		10. $\{t_7, t_8, t_9, t_{11}, t_{13}\} \Rightarrow t_{10}$	0.82	3.65	131

TABLE VI: Darknet Threat Patterns

- Removing threats, prior to the analysis, that do not contribute in information gain (i.e., a threat that is absent during the analyzed period).

In the current work, we selected a portion of darknet providers network blocks as the target of attacks. Specifically, we restricted the target of the attacks to five /24 network blocks. Table VI represents our frequent pattern and association rule mining results. For confidentiality and privacy matters, we anonymized some sensitive information. This table discloses the analyzed IP blocks, their corresponding identified threat patterns or association rules, coupled with their lift, their confidence and their number of occurrences per day. The latter metric is an indication that the identified threat pattern is valid since it frequently occurs per unit of time (a day in our current analysis). Up to this point, we have demonstrated that certain darknet threats are in fact correlated or co-occur when targeting specific network destinations. For example, consider association rule 1 in Table VI. This rule discloses that if we detect threats $\{t_7, t_8, t_9\}$ in some order in the live darknet data stream, then with 63% confidence we can as well expect to predict that threat t_{10} will follow or occur. Note that these threats are correlated since the value of the *lift* is > 1 . In the sequel, we attempt to provide an interpretation to the identified threat patterns. Please refer to the numbered association rules in Table VI as a reference to the below interpretations. It is worthy to note, that such interpretations are solely derived from the threat patterns and the NIDSs threat descriptions. Hence, we aimed to provide the most logical and best fit scenario considering the threat association rules. We believe that one interesting outcome of this work is the ability to provide insights about threat patterns and interpret real world threat scenarios. Future work in this area could provide more elaborative interpretations.

The first association rule discloses the following information. A Unix host, running FreeBSD, attempts to fingerprint a target Voice over IP (VoIP) Session Initiation Protocol (SIP) server on port 5060. By fingerprinting, the attacker hopes to retrieve the servers identification information such as operating system and installed services. Finally, he leverages his attack by sending an enormous number of malformed ICMP packets directed towards the SIP server. The latter can be interpreted as a denial of service attempt. The second association rule reveals the subsequent information. An exploited Windows

host first attempts to ping a target to check if it is alive. To retrieve more information, he initiates various traceroute commands. Moreover, he attempts to connect to a certain undisclosed port. However, he is faced with an unable to connect error message. The latter effort can be explained by an attempt to gain system access. The third association rule can be interpreted as the following. A typical attacker first performs port and host scanning to identify security vulnerabilities and possible ways to get system access. Sequentially, he can trigger various traceroute commands to retrieve more information on how to reach his target. Finally, he will attempt to execute a high priority threat (a buffer overflow exploit) to gain elevated privilege on the victim’s system. The fourth association rule presents a scanning attack targeting IP version 6. Specifically, it discloses that an attacker first attempts to fingerprint a server running IPv6. After receiving a request timed out reply, he launches a traceroute command to further explore his target’s path. Finally, he extends his attack by sending a series of ICMP packets. The latter can be interpreted as a denial of service attempt against the IPv6 server. The fifth association rule discloses the following information. A Unix host, running FreeBSD as an operating system, attempts to fingerprint a target server on TCP port 80. By fingerprinting, the attacker hopes to retrieve the server’s (possibly the web server’s) identification information such as operating system and installed services. This can be a prelude to discovering vulnerabilities and sequentially instrumenting a targeted attack. His scanning request is made from a Flowpoint 2200 DSL router. However, the reply is a message indicating that such port is unreachable. In an attempt to gather more information about the target, the attacker consequently launches various traceroute commands. The sixth association rule can be interpreted as the following. An attacker aims to target a Microsoft server running as a domain controller. The server, running Windows 2000 Server, has the Microsoft directory services installed and running. The attacker first tries pinging the server to see if it is operational. After receiving a positive confirmation, he elevates his attack by tracing the path to reach the server. Finally, he leverages his attack by sending an enormous number of malformed ICMP packets directed towards the domain controller. The seventh association rule is a series of scanning attempts on UDP port 53, a port normally dedicated for the domain name service (DNS). A host running Windows 9x generated a significant

number of ICMP echo requests directed towards the server. In an attempt to gather more information about the target, the attacker consequently launches traceroute commands. The eighth association rule unveils the following information. An attacker launches various traceroute commands from a Unix host. He leverages his scanning attempts by sequentially targeting TCP port 3389, the Windows Remote Desktop Protocol (RDP). This event is alarming since it can be interpreted as an attempt to gain system access especially if the mentioned service is vulnerable or if its authentication is inadequately configured. The ninth and tenth association rules are syntactically different, however contextually, they can be interpreted similarly. They disclose that an exploited host is generating enormous malformed ICMP packets towards a certain target. This is an indication of an attempt to launch a denial of service attack against the target victims.

VI. CONCLUSION AND FUTURE WORK

In this paper, we investigated darknets by performing darknet characterization and traffic profiling. We interpreted the output of this step by providing insights as indicators for cyber threat activity. Particularly, the results can be summarized in the following: scanning traffic constitute the majority of darknet traffic; TCP leads the darknet protocol distribution; SIP contributes as the major darknet application layer protocol; IP Class 'C' is the most destined class of darknet traffic; TCP port 445, pertaining to Microsoft active directory service, is the most targeted port. We presented and discussed darknet-triggered threats. Distinctively, we highlighted various threats as well as their severities and elaborated on their nature and consequences. This analysis step revealed three high severity threats, namely, denial of service attempts, buffer overflow exploits and unsolicited VPN access. Furthermore, we explored the inter-correlation of such threats, by applying association rule mining techniques, to build threat association rules. Such work demonstrated that in fact certain darknet threats are correlated when targeting specific network destinations. Moreover, it provided insights about threat patterns and allowed the interpretation of threat scenarios. Among the identified threat clusters, was one leading to a high priority buffer overflow exploit. For future work, we intend to provide more cyber threat insights and build a classification model from the threat association rules to experiment its predictability features with near real time darknet traffic.

REFERENCES

- [1] Stephen Hinde. The law, cybercrime, risk assessment and cyber protection. *Computers & Security*, pages 90–95, 2003.
- [2] Public Safety Canada. Canada's Cyber Security Strategy, 2009. Available at: http://www.publicsafety.gc.ca/prg/ns/cbr/_fl/ccss-scc-eng.pdf.
- [3] Niels Provos. A virtual honeypot framework. In *Proceedings of the 13th conference on USENIX Security Symposium - Volume 13*, SSYM'04, pages 1–1, Berkeley, CA, USA, 2004. USENIX Association.
- [4] David Moore, Colleen Shannon, Geoffrey M. Voelker, and Stefan Savage. Network telescopes: Technical report. pages 1–14, 2004.
- [5] Michael Bailey, Evan Cooke, Farnam Jahanian, Jose Nazario, and David Watson. The Internet motion sensor: A distributed blackhole monitoring system. In *Proceedings of Network and Distributed System Security Symposium (NDSS'05)*, pages 1–13, San Diego, CA, February 2005.
- [6] Vinod Yegneswaran, Paul Barford, and Dave Plonka. On the design and use of internet sinks for network abuse monitoring. In *Proceedings of the 7th International Symposium on Recent Advances in Intrusion Detection (RAID)*, pages 146–165, 2004.
- [7] R. Pang, V. Yegneswaran, P. Barford, V. Paxson, and L. Peterson. Characteristics of internet background radiation. In *Proceedings of the 4th ACM SIGCOMM conference on Internet measurement*, IMC '04, pages 27–40, New York, NY, USA, 2004. ACM.
- [8] Eric Wustrow, Manish Karir, Michael Bailey, Farnam Jahanian, and Geoff Huston. Internet background radiation revisited. In *Proceedings of the 10th annual conference on Internet measurement*, IMC '10, pages 62–74, New York, NY, USA, 2010. ACM.
- [9] K. Fukuda, L.A.N. Amaral, and H.E. Stanley. Dynamics of temporal correlation in daily internet traffic. In *Global Telecommunications Conference, GLOBECOM. IEEE*, December 2003.
- [10] D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford, and N. Weaver. Inside the slammer worm. *Security Privacy, IEEE*, 1(4):33–39, July-August 2003.
- [11] David Moore, Colleen Shannon, Douglas J. Brown, Geoffrey M. Voelker, and Stefan Savage. Inferring internet denial-of-service activity. *ACM Transactions on Computer Systems*, 24:115–139, May 2006.
- [12] Mark Allman, Vern Paxson, and Jeff Terrell. A brief history of scanning. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, IMC '07, pages 77–82, New York, NY, USA, 2007. ACM.
- [13] K. Fukuda, T. Hirotsu, O. Akashi, and T. Sugawara. Correlation among piecewise unwanted traffic time series. In *Global Telecommunications Conference, 2008. IEEE GLOBECOM 2008. IEEE*, pages 1–5, December 2008.
- [14] Eric Wustrow, Manish Karir, Michael Bailey, Farnam Jahanian, and Geoff Huston. Internet background radiation revisited. In *Proceedings of the 10th annual conference on Internet measurement*, IMC '10, pages 62–74, New York, NY, USA, 2010. ACM.
- [15] D. Sisalem, J. Kuthan, and S. Ehlert. Denial of service attacks targeting a SIP VoIP infrastructure: attack scenarios and prevention mechanisms. *Network, IEEE*, 20(5):26–31, September-October 2006.
- [16] APWG. Global phishing survey: Trends and domain name use in 2h2011.
- [17] Microsoft Security TechCenter. Microsoft security bulletin ms09-018 - critical. Available at: <http://technet.microsoft.com/en-us/security/bulletin/MS09-018>.
- [18] Saamil Shah. Top ten web attacks. Available at: <http://www.blackhat.com/presentations/bh-asia-02/bh-asia-02-shah.pdf>.
- [19] Learn Security Online. MS terminal service cracking. Available at: http://www.carnal0wnage.com/papers/iso_ms_terminal_server_cracking.pdf.
- [20] Snort. Available at: <http://www.snort.org>.
- [21] The bro network security monitor. <http://bro-ids.org/>, last accessed on December 2011.
- [22] A. Thomas. Rapid: Reputation based approach for improving intrusion detection effectiveness. In *Information Assurance and Security (IAS), 2010 Sixth International Conference on*, pages 118–124, August 2010.
- [23] National Institute of Standards and Technology (NIST)-National Cyber-Alert System. Vulnerability summary for cve-2007-2931, 2011. Available at: <http://web.nvd.nist.gov/view/vuln/detail?vulnId=CVE-2007-2931>.
- [24] Sourcefire-Snort, 2011. Available at: <http://www.snort.org/search/sid/587>.
- [25] R. Agrawal, T. Imielinski, and A. Swami. Mining association rules between sets of items in large databases. *ACM SIGMOD Record*, 22(2):207–216, June 1993.
- [26] J. Han and J. Pei. Mining frequent patterns by pattern-growth: methodology and implications. *ACM SIGKDD Explorations Newsletter*, pages 14–20, December 2000.
- [27] M. J. Zaki. Scalable algorithms for association mining. *IEEE Transactions of Knowledge and Data Engineering (TKDE)*, 12:372–390, 2000.
- [28] John D. Holt and Soon M. Chung. Efficient mining of association rules in text databases. In *Proc. of the 8th ACM International Conference on Information and Knowledge Management (CIKM)*, pages 234–242, Kansas City, Missouri, United States, 1999.
- [29] Jiawei Han, Micheline Kamber, and Jian Pei. *Data Mining: Concepts and Techniques (The Morgan Kaufmann Series in Data Management Systems)*. Morgan Kaufmann, 2nd edition, January 2006.
- [30] Mark A. Hall Ian H. Witten, Eibe Frank. *Data Mining: Practical Machine Learning Tools and Techniques*. Morgan Kaufmann, Third edition, January 2011.