

Intelligent Defense using Pretense against Targeted Attacks in Cloud Platforms

Roshan Lal Neupane^{1a}, Travis Neely^{1b}, Prasad Calyam^b, Nishant Chettri^a, Mark Vassell^a, Ramakrishnan Durairajan^c

^a{rlnzq8, nc5ff, mdvy96}@mail.missouri.edu

^b{neelyt, calyamp}@missouri.edu

^cram@cs.uoregon.edu

Abstract

Cloud-hosted services are being increasingly used in online businesses in e.g., retail, healthcare, manufacturing, entertainment due to benefits such as scalability and reliability. These benefits are fueled by innovations in orchestration of cloud platforms that make them programmable as Software Defined everything Infrastructures (SDxI). At the same time, sophisticated targeted attacks such as Distributed Denial-of-Service (DDoS) and Advanced Persistent Threats (APTs) are growing on an unprecedented scale threatening the availability of online businesses. In this paper, we present a novel defense system called *Dolus* to mitigate the impact of targeted attacks launched against high-value services hosted in SDxI-based cloud platforms. Our Dolus system is able to initiate a ‘pretense’ in a scalable and collaborative manner to deter the attacker based on threat intelligence obtained from attack feature analysis. Using foundations from *pretense theory in child play*, Dolus takes advantage of elastic capacity provisioning via ‘quarantine virtual machines’ and SDxI policy co-ordination across multiple network domains to deceive the attacker by creating a false sense of success. We evaluate the efficacy of Dolus using a GENI Cloud testbed and demonstrate its real-time capabilities to: (a) detect DDoS and APT attacks and redirect attack traffic to quarantine resources to engage the attacker under pretense, (b) coordinate SDxI policies to possibly block attacks closer to the attack source(s).

Keywords: Software-defined Infrastructure, DDoS Attacks, Advanced Persistent Threats, Pretense Theory, Network Analytics for Targeted Attack Defense

¹ These authors contributed equally to this work

1 **1. Introduction**

2 Cloud computing has become an essential aspect of online services available
3 to customers in major consumer fields such as e.g., retail, healthcare, manufactur-
4 ing, and entertainment. On-demand elasticity, and other benefits including diver-
5 sity of resources, reliability and cost flexibility have led enterprises to pursue the
6 development and operations of their applications in a “cloud-first” fashion [1].

7 Technological trends indicate that the aforementioned benefits typically rely
8 on software-centric innovations in the orchestration of cloud resources. These
9 innovations include cloud platforms based on Software Defined everything In-
10 frastructures (SDxI) that allow programmability to achieve capabilities such as
11 speed and agility [2] in elastic capacity provisioning. Additionally, they provide
12 opportunities to create Software-Defined Internet Exchange Points (SDXs) be-
13 tween multiple Software-Defined Network (SDN) domains (or Autonomous Sys-
14 tems (ASes)) that can enable application-specific peering, knowledge sharing of
15 cyber threats, and other cross-domain collaborations [3].

16 While the adoption of SDxI-based clouds is starting to mature, sophisticated
17 targeted attacks such as Distributed Denial-of-Service (DDoS) attacks and Ad-
18 vanced Persistent Threats (APTs) are simultaneously growing on an unprece-
19 dented scale. DDoS attacks can have significant effects on cloud-hosted ser-
20 vices (i.e., attack “targets”) and are continual threats on the availability of online
21 businesses to customers. If successful, they also cause significant loss of rev-
22 enue/reputation for a large number of enterprises for extended periods of time.
23 From the customers’ perspective, application consumption interruptions due to
24 cyber attacks can lower their overall Quality of Experience (QoE) and can lead
25 to loss of trust, or in worst cases, the termination of cloud-hosted application
26 provider services.

27 Different from DDoS attacks, APTs are a form of attacks that are character-
28 ized by computer viruses/trojans/worms, which hide on network devices (personal
29 computers, servers, mobile devices). The nature of APT attack behavior is to ex-
30 filtrate data from within the network, to devices outside the network. While DDoS
31 attacks are large scale and forthright with a goal of obvious disruption, APTs are
32 quite the opposite, subtle and secretive while also ranging from small to large
33 scale attacks. The aim of the long-term attack is to go unnoticed for as long as
34 possible so that maximum exfiltration can occur. Many APTs will attempt to ex-
35 ploit both Zero-Day attacks (faults in software which have not been discovered

36 by the application developers or hardware vendors and can be exploited) as well
37 as human error (e.g., the curiosity of finding a flash drive in a parking lot, tak-
38 ing it, and attempting to use it). A combination of these methods are also used
39 for initially breaking into an application system as well as spreading through an
40 enterprise infrastructure [4].

41 Given the benefits of SDxI-based cloud platforms, the traditional Intrusion
42 Prevention Systems (IPS) and Intrusion Detection Systems (IDS) solutions are
43 undergoing major transformations. Recently, defense strategies such as SDN-
44 based “moving target defense” [5] [6] have been proposed to protect networks
45 and users against DDoS attacks by migrating networks and users from targeted
46 virtual machines (VMs) to other healthy/safe VMs in a cloud platform. However,
47 such strategies may cause the application response behavior to change to an extent
48 that alerts the attacker that a high-value target has been hit. Given such a discovery
49 that a service provider is moving a target in order to shelter from the attack impact,
50 the attacker may then deflect more resources to seek ransom demands in order to
51 stop the DDoS on the target.

52 Moreover, if the DDoS attack flows are blacklisted, traditional approaches al-
53 low defense only at the attack destination side i.e., any related traffic is dropped
54 at the target-end. In such cases, the attacker still can escalate the DDoS attacks
55 by crossing many other neighboring domain paths, who may not be inclined to
56 drop the attack flow traffic assuming it may be legitimate traffic of a peer net-
57 work. We suppose that SDxI-based cloud platforms can facilitate capabilities for
58 coordination of policies and creation of incentives to block such targeted attack.
59 Threat intelligence collection and corresponding analytics can be developed to
60 block malicious flows closer to the attack source side, which can then mitigate the
61 impact on resource flooding for all the providers involved. However, this might
62 require the target service provider to buy some time in order to bring ‘humans into
63 the loop’ to actually enforce attack traffic blocking measures closer to the attack
64 source side.

65 The above defense strategies in SDxI-based cloud platforms could also be ap-
66 plied to defend against APTs, however they pose a different set of challenges.
67 Since the APTs attempt to be stealthy and commonly use Zero-Day attacks, it is
68 difficult to detect them with existing IDS solutions. Many of these attacks go un-
69 noticed for years, such as Red October, which was active for over five years [7].
70 With such long lasting and subtle attacks, new threat intelligence collection meth-
71 ods and corresponding analytics technologies are needed to detect APT related
72 attacks quickly and defend against them before any further long term damage or
73 exfiltration can be accomplished.

74 In this paper, we address the above challenges and present a novel defense
75 system called *Dolus* (named after the spirit of trickery in Greek Mythology) to
76 mitigate the impact of targeted attacks such as DDoS attacks and APTs launched
77 against high-value services hosted in SDxI-based cloud platforms. Our Dolus
78 approach is novel owing to a scalable and collaborative defense strategy which
79 use foundations from *pretense theory in child play* [8] [9] along with SDxI-based
80 cloud platform capabilities for: (a) elastic capacity provisioning via ‘quarantine
81 VMs’, and (b) SDxI policy coordination across multiple network domains. Such a
82 strategy is aimed at preventing the disruption of cloud-hosted services (i.e., Loss
83 of Availability) and/or the exfiltration of data (i.e., Loss of Confidentiality) by
84 deceiving the attacker through creation of a false sense of success, and by allowing
85 the attacker to believe that a high-value target has been impacted or that high value
86 data has been accessed or obtained.

87 DDoS attack detection is performed in the Dolus system using the threat in-
88 telligence obtained from attack feature analysis in a two-stage ensemble learning
89 scheme that we developed. The first stage focuses on anomaly detection to iden-
90 tify salient events of interest (e.g., connection exhaustion), and the second stage
91 is invoked to distinguish the DDoS attack event type amongst the 5 common at-
92 tack vectors: DNS (Domain Name System), UDP (User Datagram Protocol) frag-
93 mentation, NTP (Network Time Protocol), SYN (short for synchronize), SSDP
94 (simple service discovery protocol).

95 Dolus uses an automated defense strategy that we developed to mitigate APT
96 attacks, which we refer to as Automated Defense against Advanced Persistent
97 Threats (ADAPTs). Our goal in ADAPTs design is to detect which devices may
98 be infected by an APT, by pursuing tracking for data exfiltration outside of an
99 enterprise network. Once a device is suspected of being infected by an APT,
100 the device’s traffic can be rerouted so that it does not leave the enterprise net-
101 work, but can instead be analyzed to determine what is being exfiltrated or what
102 has been compromised. In order to detect possible APTs and identify systems,
103 which have been compromised by an APT, we use a concept called *Suspicious-*
104 *ness Scores* [10]. A Suspiciousness Score is assigned to each device on or off the
105 network. Each device will be assigned a score which is calculated based upon its
106 total number of unique destinations contacted, total number of connections, and
107 total number of bytes transmitted. Using these scores we create a baseline for the
108 entire network. Consequently, devices which are ‘suspicious’ will stand out with
109 higher scores. Suspiciousness Scores are calculated for internal devices, external
110 devices and domains. Consequently, an external device or domain, which we find
111 to be suspicious can later be blocked from devices on the internal network.

112 In addition to the baseline scoring for individual devices and the overall net-
113 work, we also introduce a novel concept of *Targeted Suspiciousness Scores*. Such
114 a scoring overcomes limitations in effectiveness of tracking suspiciousness while
115 accurately detecting a variety of attack traffic with unique characteristics. E.g.,
116 we can differentiate APT attacks involving malicious data exfiltration from their
117 variants such as Advanced Persistent Miner (APM) attacks (also referred to as
118 cryptojacking attacks), which involve malicious resource exfiltration by infecting
119 user devices with miners such as CoinHive and geth for cryptocurrency mining
120 purposes [11].

121 We evaluate the efficacy of our Dolus using two GENI Cloud [12] testbeds,
122 one for DDoS detection and the other for APT attacks detection. The DDoS de-
123 tection testbed contains three SDN switches, two slave switches and a single root
124 switch. The slave switches are each attached to users and attackers, a quarantine
125 VM, and a connection to the root switch. Likewise, the root switch is connected
126 to elastic VMs, each of which could serve as a candidate for the target application
127 (i.e., a video gaming portal) hosting that could be compromised by the attackers.
128 All switches are connected to a unified SDN controller located in the cloud ser-
129 vice provider domain, which directs the policy updates. Our experiment results
130 demonstrate the real-time capabilities of our Dolus system to: (a) detect DDoS
131 attacks and redirect attack traffic to quarantine resources to engage the attacker
132 under pretense, and (b) coordinate SDxI policies to possibly block DDoS attacks
133 closer to the attack source(s) without affecting the (benign) cloud users/customers.
134 The APT detection testbed with ADAPTs is similar to the DDoS detection testbed
135 but features dynamic traffic manipulation, monitoring, and analysis to calculate
136 Suspiciousness Scores for each device on the network.

137 We also present experiment results for targeted suspiciousness based detec-
138 tion of APM attacks before and after whitelisting in order to evaluate the detec-
139 tion accuracy effectiveness. Lastly, we demonstrate ‘defense by pretense’ traffic
140 characteristics in comparison to normal traffic characteristics to show our pre-
141 tense novelty. Specifically, we show a resource allocation ‘fading pretense’ de-
142 fense mechanism, where resource limitations are emulated in order to diminish
143 the value of the compromised resources in the middle of an APM attack. Our
144 findings show how the defense mechanism eventually influences an attacker’s de-
145 cision to abandon the resource being exfiltrated when its money-making potential
146 drops notably.

147 Another testbed development contribution that is used in our evaluation exper-
148 iments is an Administrative User Interface (Admin UI) which we developed for a
149 network administrator to defend against targeted attacks. The Admin UI acts as a

150 central analytics hub for defense against both types of targeted attacks considered
151 in this paper. The Admin UI informs the administrator of total bytes being trans-
152 mitted through each switch connected to the controller. The administrator can
153 also update policies dynamically and on-the-fly, thus allowing for customization
154 of the network data flows allowing for human control of any data flowing through
155 the network. The Admin UI also displays suspiciousness scores for each device
156 on the network, as well as overall network suspiciousness. The administrator can
157 then use this information to determine if a device has cause for closer investigation
158 to determine if an APT exists on the device or if the device has been compromised
159 by other means.

160 The remainder of this paper is organized as follows: In Section 2, we discuss
161 related work. Section 3 provides an overview of the Dolus System design. In
162 Section 4, we provide detailed description of Dolus defense methodology against
163 DDoS attacks. Section 5 details our Dolus strategy for defense against APT
164 attacks. Section 6 evaluates the performance of Dolus system in GENI Cloud
165 testbeds. Section 7 concludes this paper.

166 **2. Related Work**

167 *2.1. Attack Defense using Trickery*

168 There have been efforts that seek to implement defense mechanisms using
169 some form of ‘trickery’ to engage an attacker. For example, authors in [13] in-
170 troduce the notion of tricking the attackers through IP randomization methods
171 in decoy-based MTD efforts. In contrast, the notion of pretense in our Dolus
172 approach is akin to Honeypots and Honeynets which are effective in gaining in-
173 formation about possible attacks based on minimal active interactions with attack-
174 ers [14]. Primarily they are used in a setting to either gain more information about
175 potential attacks or the behavior of attackers.

176 Our work is complementary to Honeypots/Honeynets: we employ pretense
177 to deceive attackers by rerouting and responding to attack traffic using quaran-
178 tine VMs. Dolus system’s pretense theory is mainly built upon the work in [8]
179 and [9] belonging to the field of child pretend play psychology. Our novel *defense*
180 *by pretense* mechanism for effective mitigation of targeted attacks is inspired by
181 the authors’ experiments where they show children (analogous to our attackers)
182 various pictures of the animals along with a mismatch of the sounds made by the
183 associated animals. Observations are made on how a pretense is effective based on
184 how long it takes for a child to understand/protest that the information portrayed
185 is actually false. In our case, the longer an attacker is tricked by our pretense, the

186 more time a cloud service provider has to perform MTD mechanisms, strategize
187 on patching identified vulnerabilities, as well as implement a SDxI-based infras-
188 tructure policy coordination for mitigation of the impact of a targeted attack.

189 *2.2. Defense against DDoS Attacks*

190 Defense against flooding attacks such as DDoS typically involves attack traf-
191 fic feature learning that provides intelligence on where the attack is coming from,
192 and the specific attack type(s) [15] [16] [17]. Analysis of features such as source
193 IP, destination IP, source port, destination port, size of packets, packet identifiers
194 commonly help in subsequent filtering of flooding attacks. Authors in [18] show
195 that the Internet traffic patterns are distinguishable, which can help filter and iso-
196 late attack traffic flows. Once attack flows are filtered, blacklists are created [19],
197 which can then be used to “scrub” the flows through scrubbing SaaS services as a
198 low-cost solution [20].

199 A number of other network-based defense strategies have been proposed in
200 efforts that involve analysis of traffic and dynamic updation of rules to effectively
201 reroute malicious traffic. Such efforts include [21], where a network reacts to tar-
202 getted attacks using accountability and content-aware supervision concepts. Simi-
203 larly, using volume counting, authors in [22] provide a DDoS defense mechanism
204 that involves monitoring SDN traffic flows in OpenFlow-enabled switches. In
205 the context of programmability of SDN switches to mitigate targeted attacks, au-
206 thors in [23] present a programming framework. In another similar effort, authors
207 in [24] propose a memory-efficient system that uses Bloom filter and monitoring
208 tools to dynamically update SDN rules to mitigate DDoS attacks. Also leveraging
209 the dynamic rule update feature of SDN, authors in [25] analyze the probability
210 that a flow is traced back across multiple ASes’ hops by sampling the probability
211 and the analyzing signatures of attack traffic flows.

212 Alternately, cloud service providers allow mitigation of DDoS attacks by uti-
213 lizing the elastic capacity provisioning capabilities in the cloud platforms that
214 allow “moving target defense” (MTD) techniques to be implemented. MTD ba-
215 sically involves replication and live migration [26] of compromised application
216 services (with pre-attack state information) in new VM(s) to redirect legitimate
217 users, and keep attackers in a quarantine VM(s) [6]. As an added defense strat-
218 egy, authors in works such as [27] present a survey of SDN-based mechanisms to
219 detect attacks closer to the attackers/attack sources.

220 2.3. Defense against APT Attacks

221 Techniques to detect APTs have been of interest to the community [28, 29,
222 30, 30, 10, 31, 32, 33, 34, 35, 36, 37]. This includes finite angular state velocity
223 machines and vector mathematics to model benign versus attack traffic, allowing
224 a network operator to easily view the differences [34], assessing the outbound net-
225 work traffic [10, 31], using honeypots [38] and using distributed computing [37].
226 Another APT detection technique is based on a ranking system where all inter-
227 nal hosts are ranked based on number of bytes sent outside the network, number
228 of data transfers initiated to an entity outside the network, and number of dis-
229 tinct destinations contacted outside the network per host [10]. Yet, another APT
230 detection technique is to monitor attack traffic using a detector in an enterprise
231 network [35].

232 Potential countermeasures against APTs are discussed in [4, 39, 40, 41, 42,
233 43, 44, 45]. Defense strategies include: (a) running routine software updates to
234 avoid backdoors, bugs and vulnerabilities; (b) strengthening network access con-
235 trol and monitoring services; (c) enabling strict Internet access policies; and (d)
236 dropping encrypted traffic from unknown hosts. Similarly, authors in [4] discuss
237 a number of counter measures against APTs including training users about social
238 engineering attacks, blacklisting hosts, dropping packets, etc. Furthermore, SDN-
239 based defense [43] involves: (i) defining and maintaining a network baseline to
240 identify any deviation from the baseline through analytical tools, and (ii) updat-
241 ion of flow policies for (re)directing and blocking traffic in any of the network seg-
242 ments. A framework that realizes such an SDN-based defense is discussed in [44].
243 In a similar vein, authors in [45] provide a sandbox environment using which a
244 security professional can emulate the propagation of APTs across an enterprise
245 network environment.

246 3. Dolus Defense Methodology

247 In this section, we first present an overview of pretense theory. Following this,
248 we describe how the pretense theory is used in our Dolus system design.

249 3.1. Pretense Theory

250 The pretense in the Dolus system is designed to create stimulus from the target
251 side that matches the initial expectation of an attacker that a high-value target has
252 not yet been compromised through an automated bot activity. Pretense theory
253 concepts from [8] motivate us to address the issue of how a cognitive agent can

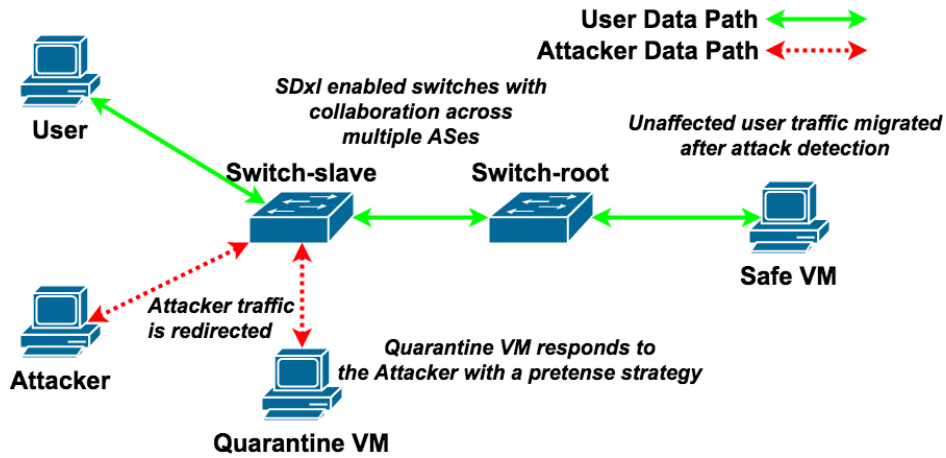


Figure 1: Illustration of the proposed Dolus system scheme wherein the attacker is *tricked* by redirection of the attack traffic to a quarantine VM for pretense initiation, while the providers work collaboratively to block the attack traffic closer to the source side.

254 present a pretense world, which is different from the real world using the following
 255 four steps:

- 256 (a) The basic assumption(s) or premise(s) that is used by a pretender on *what*
 257 is being pretended.
- 258 (b) Inferential elaboration which details of what goes into or what actually hap-
 259 pens in the process of pretense.
- 260 (c) Appropriate behavior production which answers the question of whether the
 261 pretender was successful on the audience being tricked.
- 262 (d) Balancing and steering the effects of pretense.

263 For use cases to guide our design, we borrow ideas from an example experi-
 264 ment from [9], where a child (i.e., the attacker in our case) is shown the image of
 265 a dog that makes the sound of a duck. In this situation, the child protests saying
 266 that it is not the sound that a dog makes. However, if the same child is shown a
 267 image that seemingly looks like a duck (in reality, it is not) and makes the sound
 268 of a duck, then there is no protest and the child falls for the pretense. However,
 269 given additional observation time, the child realizes he/she has been tricked and
 270 protests.

271 *3.2. Pretense in Dolus System Design*

272 In our Dolus system, effective pretense design methodology is illustrated in
 273 Figure 1. We create pertinent stimulus from the target side i.e., redirecting attack
 274 traffic to a quarantine VM that mimics original target behavior, when our two-
 275 stage ensemble learning algorithm (explained in Section 4.2) can identify and then
 276 blacklist the attacker flows while allowing benign user flows to continue unim-
 277 peded. This in turn could help in keeping an attacker distracted for a brief period
 278 of time while the pretense is in effect.

279 From the time gained through such a pretense initiation, Dolus enables cloud
 280 service providers to decide on a variety of policies by dynamically generating net-
 281 work policies using Frenetic [46] to mitigate the attack impact, without disrupting
 282 the cloud services experience for legitimate users. In the worst case, destination-
 283 side blocking can be enforced. Alternately, if the cloud service provider uses the
 284 attack intelligence information and successful pretense time to coordinate the ‘hu-
 285 mans in the loop’ of neighboring SDN-enabled domains, together they can direct
 286 a unified SDN controller that directs SDN-enabled switches to actually enforce
 287 attack traffic blocking measures closer to the attack source side.

Table 1: Objectives for the pretense zero-sum game considered in the Dolus System design.

Objective	Attacker	Defender
Goal	Evade Defender’s detection	Protect attacker’s target(s)
Trick	Defender to provide access	Attacker to reveal presence
Time	Mislead defender to spend time on false positives	Mislead attacker to spend time on true negatives
Outcome	Make defender believe that an attack is simple	Make the attacker believe that the target attack is successful
Attribution	Hide attackers’ identity	Induce attackers to believe that their identities are unknown

288 The goal of our Dolus approach is to model the notion of pretense as a zero-
 289 sum game. Specifically, a zero-sum game is one in which the sum of the individual
 290 payoffs for each outcome is zero. That is, (1) loss to an attacker is gain for a de-
 291 fender and vice versa, and (2) total sum of gain and loss is (roughly) zero. There
 292 are two strategies to play a zero-sum game, one from an attacker’s perspective
 293 and the other from the defender’s perspective. Specifically, we plan to employ
 294 the following two strategies: (a) `minimax` strategy: minimizing defenders own

295 maximum loss (from defenders perspective), and (b) `maximin` strategy: max-
296 imize attacker’s own minimum gain (from attacker’s perspective). We explore
297 these two strategies in our Dolus system on a number of objectives, which are
298 summarized in Table 1.

299 Our guiding strategy for targeted attack defense using pretense is to use a form
300 of *pretense machine learning* which we propose to be understood as - *If you don’t*
301 *know the enemy and don’t know yourself, then you will succumb in every battle.*
302 *If the attacker does not know you but you know the attacker, for some victories*
303 *gained you will suffer some defeat. If the enemy knows you and you know yourself,*
304 *you need not fear the result of a hundred battles.*²

305 For our defense by pretense strategy, we consider multiple vectors: (1) aware-
306 ness of the attack surface, i.e., cloud/physical topology aware; (2) behavior and/or
307 psychological aspects of the attacker, i.e., data science and pretense theory to
308 understand an attacker’s desire and to identify an effective countermeasure; (3)
309 development of theory and algorithms to deceive the attacker into a false sense
310 of success *without* affecting the network resources; and (4) sharing multi-domain
311 threat intelligence across SDxI entities.

312 4. DDoS Attack Defense with Dolus

313 In this section, we first describe the DDoS attack model that we assume to
314 design our Dolus defense. Subsequently, we detail our DDoS defense solution
315 that uses a ‘defense by pretense’ scheme in the Dolus system.

316 4.1. Attack Model

317 DDoS attacks aim to overwhelm network-accessible devices such as networks,
318 firewalls and end-systems in enterprises by sending packets at excessively high
319 rates from multiple attack points. With cloud-hosted applications with large mon-
320 etary value becoming highly common, DDoS attacks can cause LoA for users and
321 customers, and can be used for extortion from vulnerable online businesses. Com-
322 mon DDoS attack event types are amongst the 5 common attack vectors: DNS
323 (Domain Name System), UDP (User Datagram Protocol) fragmentation, NTP
324 (Network Time Protocol), SYN (short for synchronize), SSDP (simple service
325 discovery protocol). For the purposes of our work, we assume the DDoS attacker
326 uses SYN [47] and ICMP/Ping [48] flooding. Such attacks typically inundate a

²A quote modified from “The Art of War” by Sun Tzu.

327 networks' resources with Echo Request packets. We also assume that the attack-
 328 ers' traffic is sent constantly and may or may not solicit a response in return. Such
 329 attacks can bring the network to a standstill due to the high volume of both incom-
 330 ing and outgoing traffic. To effectively capture the semantics of this attack model
 331 and to exhaust the target application services, we generate and emit synthetic ping
 332 and HTTP traffic using `hping3` [49] and `SlowHTTPTest` [50] tools, respec-
 333 tively. Furthermore, to capture the dynamics of an attacker, we randomly change
 334 the number of attack packets emitted by these tools.




Figure 2: Cross-domain physical setup in a Dolus system deployment to share threat intelligence for a unified controller to coordinate policy management with a federation of ASes to block attack traffic closer to the source side.

335 4.2. Defense by Pretense Scheme

336 Figure 2 depicts the cross-domain setup in a Dolus system deployment to im-
 337 plement a defense by pretense scheme. To complement Figure 2, interactions
 338 between different phases of a Dolus system configured for spoofing pretense are
 339 shown in Figure 3 and Algorithm 1, respectively.

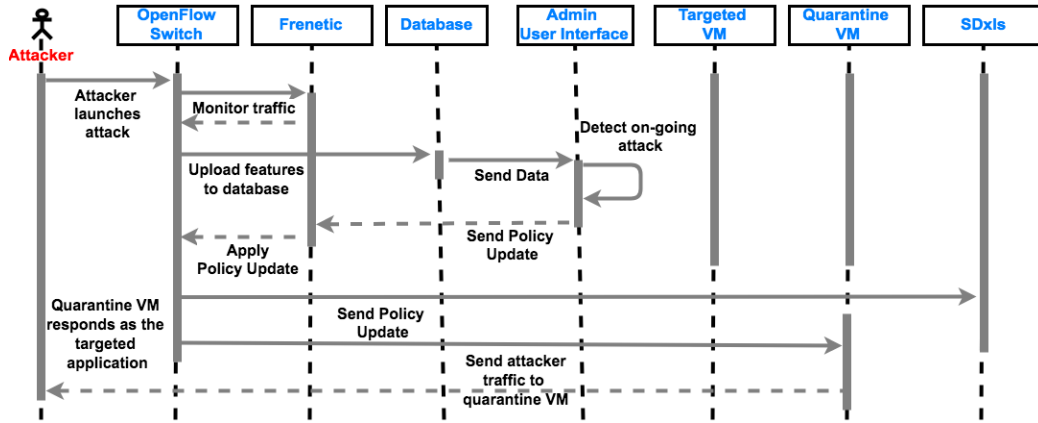


Figure 3: Sequence diagram of the Dolus system interactions for attack detection, quarantine setup, pretense initiation/maintenance and DDoS attack impact mitigation.

340 **Attack Detection.** First, traffic within a cloud provider’s network (which is gen-
 341 erated by the SDN switches) or across multiple transit provider ASes (which are
 342 composed of SDX plus SDN switch substrates) is monitored using a Frenetic run-
 343 time [46]-enabled monitoring subcomponent (line 24 of Algorithm 1). Next, in
 344 order to learn and classify the attacks (line 25 of Algorithm 1), we employ a two-
 345 stage ensemble learning scheme on the incoming traffic, both from the attackers
 346 and from the benign users. In order to differentiate attackers from benign users,
 347 the first stage handles outlier detection to identify salient events of interest (e.g.,
 348 connection exhaustion), whereas the second stage handles outlier classification to
 349 distinguish different event types (e.g., DDoS attack).

350 *Outlier Detection.* We use basic/static methods such as multivariate Gaussian to
 351 detect outliers and build upon our prior work on detecting network-wide correlated
 352 anomaly events [51, 52] that are typical of the traffic from multiple attack sources.
 353 Specifically, the outlier detection is a composition of many efficient, multivariate
 354 outlier detectors or hypotheses functions: $\mathcal{H} = \{h_1, h_2, \dots, h_n\}$ and the result,
 355 \mathcal{F} , is an ensemble of the different hypotheses. Furthermore, we note that the
 356 traditional methods for ensemble learning use averaging or majority voting [53].

357 In our case, to achieve higher accuracy with a minimum size of the training dataset
 358 D , we use the Bayesian voting scheme [54] as the ensemble method to predict the
 359 result for new data x , which can be represented as Equation 1.

$$\mathcal{F} = \sum_{h \in \mathcal{H}} h(x)P(h|D) \quad (1)$$

360 Final ensemble result \mathcal{F} consists of all of the hypotheses in \mathcal{H} , and each hy-
 361 pothesis h weighted by its posterior probability $P(h|D)$. The posterior probability
 362 is proportional to the likelihood of the training data D times the prior probability
 363 of h (2).

$$P(h|D) \propto P(h)P(D|h) \quad (2)$$

364 *Outlier Classification.* The outliers detected are classified into either interesting
 365 events (e.g., attacks) or erroneous conditions (e.g., router failure). We use a simple
 366 classifier to this end: if the final ensemble results of consecutive events (detected
 367 in the first stage) fall in the same range, we classify them as an attack; otherwise,
 368 we ignore those events. We remark that the above two-stage ensemble learning
 369 scheme requires a sizable amount of data to classify the attacks effectively. To
 370 overcome this challenge, we initially let the attacker(s) to attack the cloud ser-
 371 vices. However, we also monitor the incoming traffic carefully and make sure that
 372 the attack does not disrupt the network resources. Once an attack is classified,
 373 which are shown separately in Figure 2, we reroute the attack traffic using Fre-
 374 netic runtime to quarantine VM (QVM) along with sample server responses (see
 375 lines 1 through 22 of Algorithm 1).

376 **Quarantine Setup for Pretense.** Dolus calls the quarantine setup procedure
 377 (lines 1 to 9) where a new QVM is instantiated using a cloud platform’s elas-
 378 tic provisioning capability and the update policy routine is invoked (line 3). In
 379 the update policy routine (lines 10 to 14), we log the attack traffic to prevent fu-
 380 ture attack events as well as invoke the Frenetic runtime to generate new policies
 381 (line 12). Frenetic executes Python scripts to identify suspicious packets, learn
 382 from patterns and directs switches to redirect packets to QVMs. We then adver-
 383 tise this information (attack intelligence) to the neighboring switches (line 13),
 384 where, apart from the policy updates, the IP addresses of the attackers are black-
 385 listed. Following this, based on the stored attack traffic logs, the QVM uses Scapy
 386 libraries [55] to generate responses with spoofed IP addresses and pretends as the
 387 targeted VM under attack from the perspective of the attacker(s) (lines 20 to 22).

388 Subsequently, depending on the nature and volume of the incoming data, we

389 decide either to move forward with the pretense or drop the traffic—which is the
390 third step of production of appropriate behavior in pretense theory (lines 28 to
391 30). In order to gain more information about the attackers/attacks, we typically
392 continue the process of pretense. While we continue the pretense, we routinely
393 update threat intelligence such as the attacker’s IP, targeted VM’s IP where ser-
394 vice(s) under attack is hosted, type of attacks, etc. Furthermore, we assume that
395 an attacker has enough knowledge on how a successful attack should affect our
396 system, which is another reason why we keep the attacker involved in the system
397 as long as is usually expected. If we drop the attack traffic too early or maintain it
398 for too long, attacker might potentially infer our pretense.

399 Finally, we redirect the flow of the attack traffic by pushing a new policy from
400 the unified controller running in the cloud to the switch(es). This will redirect
401 the attacker’s traffic that is intended for the targeted VM towards the QVM. The
402 QVM then responds to the attacker’s traffic as though it is the targeted VM/server
403 under attack with spoofed IP address and hostname of the target, which creates
404 the pretense effect, from an attacker’s perspective, that the targeted DDoS attack
405 is successful. Depending on the nature of the attack, we want the attacker to
406 believe that services are no longer up/available on the targeted VM. We therefore
407 allow the QVM to continue to respond to the attacker for a limited amount of time
408 t . We tune t based on the type of attack traffic and how the targeted VM would
409 respond if it was under attack. For example, if the targeted VM went down after
410 10 seconds of attack, the QVM would do the same by not responding at the same
411 time with a variable random delay factor of $[-1,1]$ seconds added. This allows the
412 attacker to see that the services are available until, suddenly, they no longer are.

413 **Policy Decision Making.** In this sense, our defense maintains the pretense: gives
414 the attacker the confirmation of a successful attack, when in reality the service has
415 not been affected at all considering the scenario that the user is running a video
416 gaming portal application. This also gives us sufficient time to collect information
417 about the attackers and their attack patterns. We use the collected information to
418 create a blacklist of attacker information. To help network administrators effec-
419 tively manage the network in the face of attacks, our system also consists of a
420 Administrator User Interface (Admin UI) module and a unified controller module
421 that can be customized in a Dolus system instance deployment depicted in Fig-
422 ure 2. The Admin UI shown in Figure 4 can be used for e.g., to enforce users
423 to adhere to the policies generated by Frenetic runtime when they connect to the
424 cloud. Policies generated by Frenetic internally are updated through the User
425 Interface using JSON arrays. These policies (e.g., open/block flows) could be in-
426 stalled in the switches using the unified controller module, which is also linked

427 with a back-end database that logs traffic characteristics and user profiles.

428 The after effects of our pretense only lasts for as long as they are needed. Dur-
429 ing the pretense, the attackers' traffic continues to be redirected a QVM near the
430 attacker. However, this process need not continue indefinitely i.e., once if it has
431 been determined that the attack traffic is no longer impacting the network, the poli-
432 cies can be updated to redirect the attacker traffic back to where it was prior to the
433 start of pretense. There are several reasons to do this: (i) changes in the dynamics
434 of the attack (e.g., bandwidth usage dropping back down to normal, absence of
435 SYN packets in a SYN flooding attack, fixing of malware in an affected machine
436 and hence it is no longer an attacker, etc.) calls for network policy changes so
437 that the network resources can be effectively used, (ii) changes in traffic e.g., IP
438 address change in incoming service requests sent from a benign user must be ser-
439 viced to meet the service level agreement (SLA), and (iii) to save the operational
440 cost of QVMs by reusing them for a different purpose e.g., periodic backups.

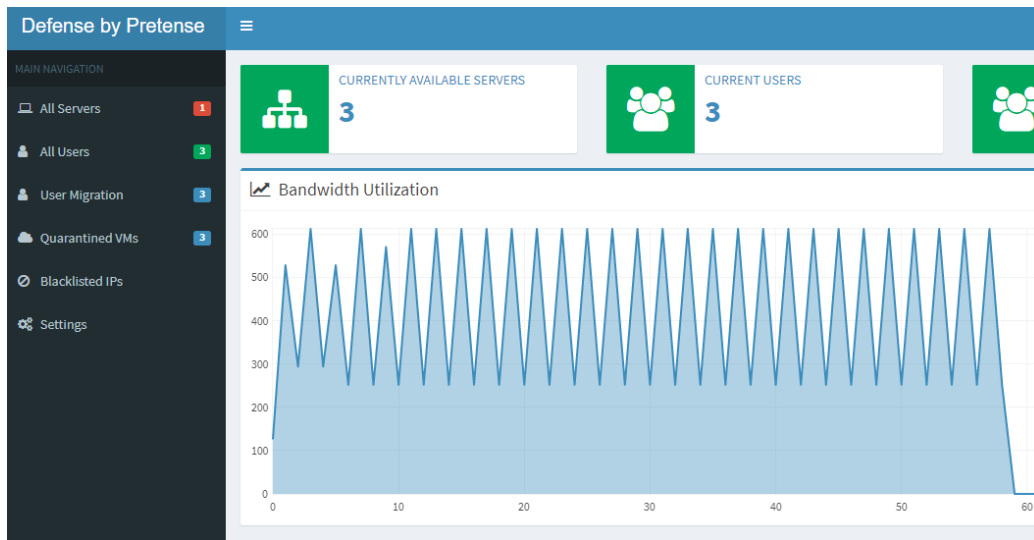


Figure 4: Administrator User Interface of an Dolus system instance.

441 **Threat Intelligence Sharing.** Algorithm 1 runs in the monitor component and
442 coordinates/shares intelligence with the switches deployed in the network and
443 across different providers. This in turn enables a collaborative environment among
444 providers such that the targeted attacks can be detected closer to the source *with-*
445 *out* affecting the cloud infrastructure. A natural question is why would a provider
446 share the attack intelligence, especially in a business that is driven by competi-

447 tion? We posit that the coordination among different ASes/providers is mutually
448 beneficial for all the entities involved. Of course, a particular AS/provider can de-
449 cide not to share the attack intelligence to others. However, if an AS experiences
450 an attack and if it shares the intelligence with other ASes, a global and unified
451 hardening of infrastructure against such targeted attacks can be achieved. In addi-
452 tion, any downtime is money lost in a business; sharing the attack intelligence in
453 turn provides a cheaper alternative to lost downtime and business.

Algorithm 1: Dolus system defense algorithm against DDoS attacks

Input: *attacker_ID* = attacker ID,

src_ip = source IP,

dst_ip = destination IP,

no_of_packets = number of packets,

spoof_dst_ip = spoofed IP,

black_ip blacklisted IP list

Result: Attack traffic will be redirected to the quarantine VM and DDoS blocking policy will be generated

```
1 function initQuarantine()
2   | createVM();
3   | updatePolicy(src_ip);
4   | do
5   |   | redirectTraffic();
6   |   | pretense_data = generateUsingScapy();
7   |   | vmResponse(spoof_dest_ip, src_ip, dst_ip, pretense_data);
8   | while timeout == false;
9 end
10 function updatePolicy(src_ip)
11   | logAttackTraffic();
12   | new_policy = generateNewPolicy();
13   | collaborate(new_policy);
14 end
15 function collaborate(new_policy)
16   | advertisePoliciesToNeighbors(new_policy);
17   | black_ip = updateList(src_ip);
18   | redirectTraffic();
19 end
20 function redirectTraffic()
21   | sendTrafficToQuarantineVM();
22 end
23 function main()
24   | /* Receive incoming data from external machine */
24   | data = monitorPackets(attacker_ID, src_ip, no_of_packets, start_time,
24   | end_time);
25   | attack = twoStageEnsembleLearning(data);
26   | /* Update policy in case of attack detected */
26   | if attack == true then
27   |   | initQuarantine(src_ip);
27   |   |
28   | end
28   | decideToStopOrContinue();
29 end
30 end
```

454 **5. APT Attack Defense with Dolus**

455 *5.1. Attack Model*

456 APTs are long-term attacks and affect a target in four stages: *preparation*,
457 *access*, *resident*, and *harvest* [7, 56, 57, 58, 59, 60, 61, 39]. In the preparation
458 stage, attackers apply a reconnaissance tactic through social engineering (e.g.,
459 via social networks) to bootstrap the attack [4]. Once the attack is bootstrapped,
460 attackers identify a vulnerability, and/or a vulnerable target and send malwares
461 either through email (e.g., spear phishing) or through third-party software/service
462 (e.g., watering-hole attack) in the access stage. Subsequently, the malwares estab-
463 lish external communication paths with attackers' Command and Control (C&C)
464 server(s), and spread across other targets in the resident stage; which is a slow and
465 a stealthy phenomenon. Finally, in the harvest stage, attackers extract any vital
466 information in an on-going fashion for extended periods of time.

467 *5.2. Defense by Pretense Scheme*

468 Our novel Dolus system with ADAPTs is designed to automatically defend
469 against APT attacks. Its design is similar to the original Dolus system algorithm
470 (i.e., Algorithm 1) for DDoS attack defense described in Section 5, however the
471 threat intelligence collection and defense are adapted towards mitigation of APT
472 attacks. More specifically, ADAPTs consists of: (1) a Suspiciousness Score-based
473 detection mechanism, which is robust against the threshold evasion problem; (2)
474 internal quarantine VMs (iQVMs), which are a minimal version of honeypots to
475 mimic hosts internal to an organization, along with performance/topology views
476 to aid network administrators; (3) a coordination mechanism driven by enterprise
477 defense policies to share threat intelligence about APTs among hosts; and (4) net-
478 work policy update mechanism to mitigate attack spreading based on coordinated
479 intelligence using iQVMs. We outline each one these mechanisms/components in
480 the remainder of this sub-section.

481 **Attack Detection.** Inspired by the work of authors in [10] to identify hosts ex-
482 hibiting suspiciousness in a network, we propose a Suspiciousness Score (SS) in
483 a similar vein. We calculate SS based on captured network traces (.pcap) using
484 three main features: *destinations* (dst), *flows*, and *bytes*.

Table 2: Features captured from a network trace for APT attack defense analytics.

Value	Description
switch_id	ID of the switch which received the frame
trace_id	ID for the trace under consideration
frame_number	Order in which the frame was received
frame_time	Unix timestamp at which the frame was received
frame_time_relative	Unix timestamp for frame receipt relative to last frame received
frame_protocols	Protocols used in the frame
frame_len	Size of the frame in bytes
ip_src	Source IP of the frame
ip_dst	Destination IP of the frame

485 Table 2 shows the list of values/features captured in network traces for APT
 486 attack defense analytics. For each packet trace, a trace_id t is assigned. For each
 487 t , we perform the following: the features are normalized and their combined Root
 488 Mean Square Error (RMSE) values are calculated. Using the RMSE values, we
 489 calculate the Suspiciousness Scores of each device as follows. The *Min* and *Max*
 490 values (below) are assumptions made per device type regarding what one may
 491 expect the minimum and maximum values to be on the type of device, network
 492 and traffic expectations. These values are determined by the system/network ad-
 493 ministrators and could vary vastly depending on each ASeS or domain’s threat
 494 monitoring objectives.

495 Destination suspiciousness for trace t :

$$dst_i = w_{dst} * \frac{numDst_i - numDistMin_i}{numDstMax_i - numDistMin_i}; w_{dst} \in [0.0, 1.0] \quad (3)$$

496 Flow suspiciousness for trace t :

$$flows_i = w_{flows} * \frac{numFlows_i - numFlowsMin_i}{numFlowsMax_i - numFlowsMin_i}; w_{flows} \in [0.0, 1.0] \quad (4)$$

497 Bytes suspiciousness for trace t :

$$bytes_i = w_{bytes} * \frac{numBytes_i - numBytesMin_i}{numBytesMax_i - numBytesMin_i}; w_{bytes} \in [0.0, 1.0] \quad (5)$$

498 Device suspiciousness for trace t is based on equations 3, 4 and 5 as shown
 499 below.

$$ss_i = \sqrt{\frac{dst_i^2 + flows_i^2 + bytes_i^2}{3}} \quad (6)$$

500 Note that for each device on the network i we calculate a Suspiciousness Score
 501 and the overall network suspiciousness for trace t is calculated based on ss for
 502 each individual device (equation 6) that is connected. That is, the sum of all ss
 503 for each devices on the network n is the *overall network suspiciousness* SS for that
 504 particular t .

$$SS_t = \sqrt{\frac{(ss_1^2 + ss_2^2 + ss_3^2 + \dots + ss_n^2)}{n}} \quad (7)$$

505 Relative change in device i 's suspiciousness score on new traffic t is simply
 506 given by:

$$\Delta ss_{it} = \frac{ss_{it} - \sqrt{\frac{ss_{i_1}^2 + ss_{i_2}^2 + \dots + ss_{i_{t-1}}^2}{t-1}}}{\sqrt{\frac{ss_{i_1}^2 + ss_{i_2}^2 + \dots + ss_{i_{t-1}}^2}{t-1}}} \quad (8)$$

507 In equations 3, 4 and 5, we assume the weight parameters i.e., w_{dst} , w_{flows}
 508 and w_{bytes} to be equal to 1 in a general case of SS calculations. As shown later
 509 through experiment findings in Section 6, assigning suitable weights for a variety
 510 of suspicious traffic can minimize the RMSE in the attack detection accuracy, as
 511 opposed to the general case. Consequently, we extend SS calculations as detailed
 512 in Algorithm 2 by introducing a novel concept of Targeted Suspiciousness Scores
 513 for specific traffic types that a system/network administrator would like to clas-
 514 sify as suspicious. We remark that system/network administrators could whitelist
 515 certain traffic types, however none of the devices on the network will be entirely
 516 whitelisted. Thus, by using a targeted suspiciousness scoring for certain traffic
 517 types (e.g., for suspected APM-like traffic) can still be effective to detect mali-
 518 cious activities, even when whitelisting is performed for legitimate user traffic.

519 **Quarantine Setup for Pretense.** VMs which are internal to an organization and
 520 which implement minimal versions of honeypot-like hosts are internal quarantine
 521 VMs (iQVMs), whose Suspiciousness Scores are monitored continuously. These
 522 are also the hosts that play the game of pretense i.e., they create a false notion of
 523 *high-value targets within an organization with sensitive data* to the external world.
 524 An attacker is lured to attack iQVMs first; they maintain pretense by sending data
 525 similar to what a host with sensitive data would send. Apart from monitoring the

Algorithm 2: Targeted Suspiciousness Score Calculations

Input: $devices \in [1 \dots n]$ = array of all devices on the network,
 $maxTrace$ = Maximum number of packet traces to be evaluated,
 t = current trace being evaluated

Result: Targeted Suspiciousness Scores calculated for each network device
after traffic analysis

```
1 function calcDst()
2 function calcFlows()
3 function calcBytes()
4 function calculateTargetedSuspiciousness(devicei)
5   | calcDst(devicei, wdst);
6   | calcFlows(devicei, wflows);
7   | calcBytes(devicei, wbytes);
8 end function main()
9   | do
10  |   | for each devicei;
11  |   |   calculateTargetedSuspiciousness(devicei);
12  |   | while  $t \leq maxTrace$ ;
13 end
```

526 data sent out of iQVMs, they also add weights to the calculated Suspiciousness
527 Scores, overcoming the threshold evasion problem.

528 To simplify the process of monitoring iQVMs and other hosts effectively, we
529 also extend our Dolus related Admin UI for use with ADAPTs. This allows the
530 administrator a more robust monitoring of the network with views separated based
531 on the various requirements: devices connected to the network, blacklisted IPs,
532 metrics, as well as any other the requirements of administrator. The user interface
533 is developed using the traditional LAMP stack (Linux OS, Apache Web Server,
534 MySQL, PHP), with views specifically built for ADAPTs including the following:

- 535 1. SS view: `Flot.js`-based bar and line graphs as depicted in Figure 5 ex-
536 cerpted from the Admin UI. SS per device or for the overall network can be
537 viewed in temporal fashion as shown in Figure 6 extracted from the Admin
538 UI. Moreover, when a suspiciousness score of a blacklisted device is shown
539 to be above a certain threshold, an administrator can block all traffic from
540 that device to the network or take an appropriate action.




Figure 5: Suspiciousness score per device over time.




Figure 6: Overall network suspiciousness score over time.

- 541 2. Upload policy view: This view on the user interface enables administra-
 542 tors to push NetKAT-based policies [62] to a centralized database, which
 543 stores device configurations, thresholds, policies maintained by the organi-
 544 zation. Interfaces are provided to select a specific device and a correspond-
 545 ing NetKAT policy to affect that device as shown in Figure 7.

Device	Filter Policies	Loaded	Remove
server1 (10.0.0.1)	Filter(SwitchEq(51570677359425) & IP4DstEq("10.0.0.4")) >> SetPort(4)	1	
attacker1 (10.0.0.7)	Filter(SwitchEq(51570677359425) & IP4DstEq("10.0.0.7")) >> SetPort(6)	1	
attacker1 (10.0.0.7)	Filter(SwitchEq(51570677359425) & IP4DstEq("10.0.0.5")) >> SetPort(8)	1	

[Add New Policy](#)

[Update Policy](#)

Figure 7: Policy table view.

546 3. Network view: a `vis.js`-based view to monitor the network as a graph of
 547 connected devices as depicted in Figure 8.

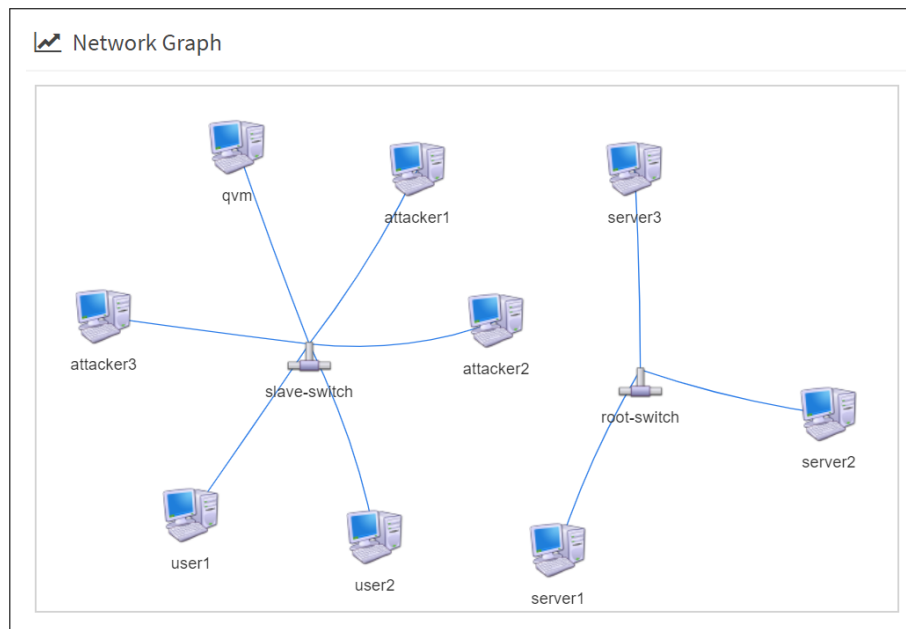


Figure 8: Network graph of all the connected devices.

548 **Policy Decision Making.** In ADAPTs, each device has a corresponding access
 549 control policy to control/configure it remotely. We call this a configuration policy,

550 which determines the virtual structure of the network and decides how traffic flows
551 traverse through the network in normal versus attack conditions.

552 Similarly, ADAPTs also features a *defense policy* for the enterprise network.
553 The defense policy is reactive i.e., it will take effect when the original configura-
554 tion policy has failed to communicate erratic host behaviors such as *SS* threshold
555 changes, jump in the number of external hosts contacted, etc., or if an attack has
556 been detected and communication privileges need revocation. The interface can fa-
557 cilitate administrators to update policies directly in the event of an attack.

558 Both these policies and the revocation/enabling functionalities are instantiated
559 based on the policy updater mechanism, whose main objective is to simplify the
560 learning curve for users/administrators to get proficient at writing policies (e.g.,
561 using network programming languages such as Frenetic [46])—a daunting and te-
562 dious task. With this in mind, the updater component can auto-generate policies
563 based on simplified inputs that are provided via the user interface. For example, to
564 minimize the process, the policy updater takes a generic command such as “user1
565 to server1” and all possible configuration policies would be generated by the up-
566 dater. The updater works with the centralized database and is pre-programmed
567 with the network architecture.

568 One of the advanced defense policies set by a system/network administrator
569 could be a ‘defense by pretense’ policy. Such a policy can adopt our novel pretense
570 concept that is adapted to a particular attack threat. Herein, we provide an example
571 of one such policy viz., ‘fading pretense’ that can be used to defend against APM
572 attacks. Invoking the defense policy on a compromised device essentially results
573 in a situation where the defense mechanisms simulates resource limitations. Such
574 a resource limitation diminishes the value of the compromised resources in the
575 middle of an APM attack. This is because, by limiting the resource allocation
576 through a hypervisor on a device that is part of a mining pool, the ability in terms
577 of the computation speed of a miner software on that device to mine a single
578 block (within a blockchain) can be slowed down. When the computation speed
579 goes below a threshold, the overhead of communicating with this compromised
580 miner device in a mining pool becomes excessively large, which in turn impacts
581 the revenue for the attacker [63]. Eventually, such a scenario will influence an
582 attacker to abandon the resources exfiltration on that device, and move onto other
583 more-potent compromised devices with higher resource allocations. A realistic
584 demonstration of the implementation of the ‘fading pretense’ policy is provided
585 later in Section 6.

586 **Threat Intelligence Sharing.** Our iQVM monitors also coordinate and share the
587 APT threat intelligence such as *SS* thresholds, policy updates, etc. with other

588 hosts in the network. Apart from providing a collaborative environment amongst
 589 pertinent hosts to effectively counter APTs, the mechanism also provides a way
 590 to drill down on specific segments of the network with suspicious hosts. Further-
 591 more, we believe that the coordination mechanism will pave the way to achieve
 592 a global and unified hardening of the enterprise network against APTs. In addi-
 593 tion, any sensitive data sent out is money lost in a business; sharing the threat
 594 intelligence in turn provides a cheaper alternative to lost data and host/business
 595 downtimes.

596 6. Performance Evaluation

597 In this section, we describe the evaluation of our Dolus methodology in GENI
 598 Cloud testbeds for DDoS and APT attacks. For showing effectiveness of Dolus
 599 for each targeted attack type, we start by describing our testbed, followed by the
 600 experiments and results discussion. The source code and instructions to replicate
 601 below experiments are openly available at [64] [65].

602 6.1. Dolus Experiments for DDoS Attack Defense

603 6.1.1. Testbed Setup

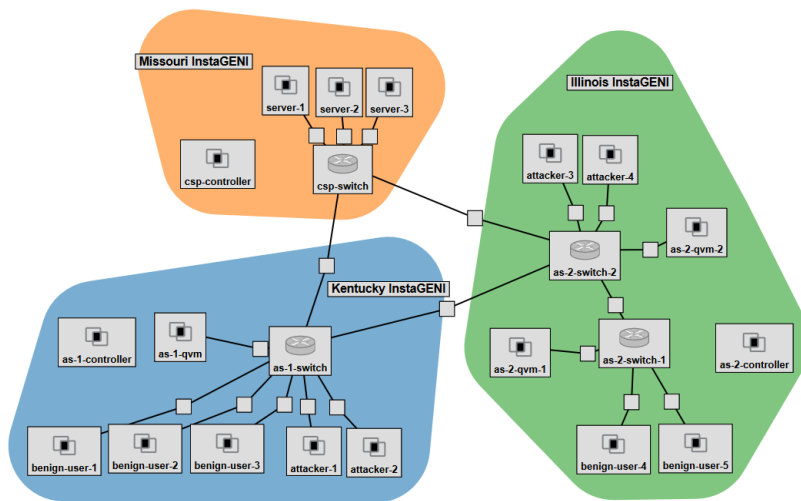


Figure 9: GENI Cloud testbed used to evaluate Dolus for DDoS attack defense.

604 We evaluate the efficacy of our Dolus system using a realistic, GENI Cloud [12]
 605 testbed as shown in Figure 9. The testbed contains three SDN switches, two slave

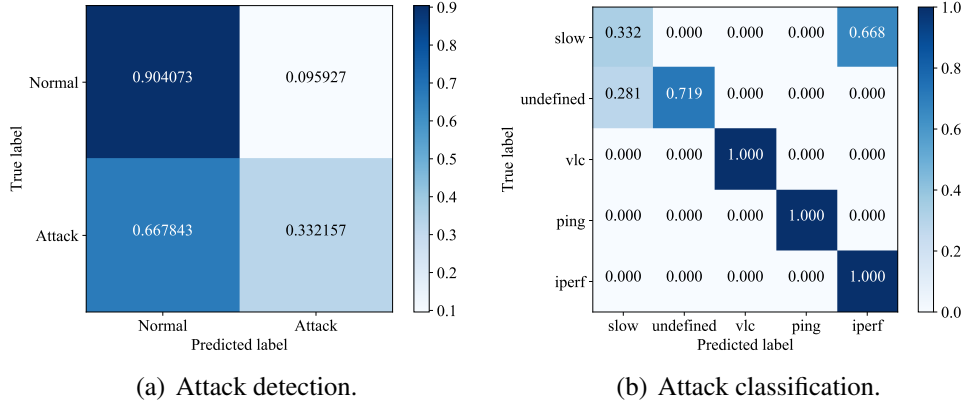


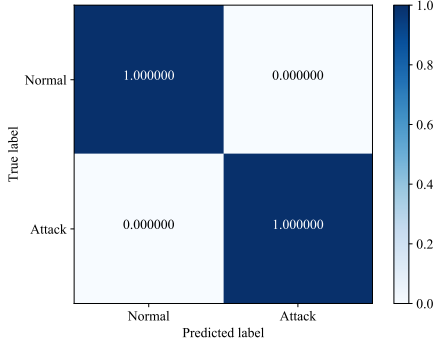
Figure 10: Confusion matrices for attack detection and classification for multiple traffic flows sent to multiple hosts.

606 switches and a single root switch. Such a system could also be extended to host
 607 many more switches and devices. The slave switches are each attached to users
 608 and attackers, a quarantine VM, and a connection to the root switch. Likewise, the
 609 root switch is connected to elastic VMs, each of which could serve as a candidate
 610 for the target application (i.e., a video gaming portal) hosting that could be com-
 611 promised by the attackers. All switches are connected to a unified SDN controller
 612 located in the cloud service provider domain, which directs the policy updates.
 613 In the following, we show the attack detection and classification accuracy using
 614 our two-stage ensemble learning scheme and then present results from two sets of
 615 experiments that were run for a maximum of 28 *seconds* to show how our Dolus
 616 implementation can be used in real-time to restore cloud services under DDoS
 617 attack situations.

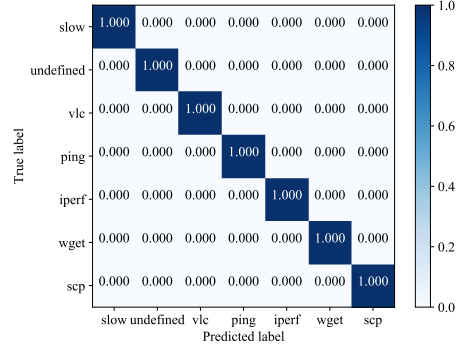
6.1.2. Attack Detection and Classification Results

619 Using the Dolus system, we monitor different types of data that are permitted
 620 to enter the GENI Cloud testbed depicted in Figure 9. We send both normal and
 621 attack traffic (i.e., our datasets) to the targeted server to test the efficacy of our two-
 622 stage ensemble learning scheme. Our evaluation results span over two instances of
 623 learning of datasets as explained in the following.

624 The first instance shows multiple traffic types from a single attacker VM to
 625 a single target node. For this instance, we divide $\sim 180,000$ lines of data into
 626 two sets, one for training and the other to test the accuracy of our scheme. Fur-
 627 thermore, the types of traffic used to create these instances are composed of
 628 SlowHTTPTest, iperf, VLC and ICMP ping. Figure 15 shows the two confu-



(a) Attack detection.



(b) Attack classification.

Figure 11: Confusion matrices for outlier detection and classification for multiple traffic flows comprising of familiar attack flows.

Table 3: Overall Attack Detection and Classification Time and Accuracy

Tests	Time (in Seconds)	Accuracy (in %)
Single server stage 1	<1	99.99
Single server stage 2	<1	99.98
Multiple hosts stage 1	7	89.12
Multiple hosts stage 2	13	98.49

629 sion matrices for attack detection and classification in a normalized fashion. We
 630 note that both the detection and the classification of attack took less than a sec-
 631 ond. In addition to the rapid detection and classification, our approach is highly
 632 accurate as shown in Table 3, where stage 1 is the detection stage and stage 2 is
 633 the classification stage.

634 In the second instance, we consider multiple traffic types to multiple hosts.
 635 This instance is composed of 2.5 million rows per test, totaling 5 million rows of
 636 data. The types of traffic that we use to create this dataset include SlowHTTPTest,
 637 iperf, VLC, scp, wget, and ICMP ping. This dataset also contains some unlabeled/undefined data for the scheme to assess and classify the training data to evaluate the effectiveness of our two-stage ensemble learning scheme. Figure 10 shows the two confusion matrices in normalized form for attack detection and classification. Detection and classification of attack took ~ 7 and ~ 13 seconds, respectively. Despite the slowdown in attack detection/classification in comparison with the first instance, the accuracy of our approach is still high as shown in
 643

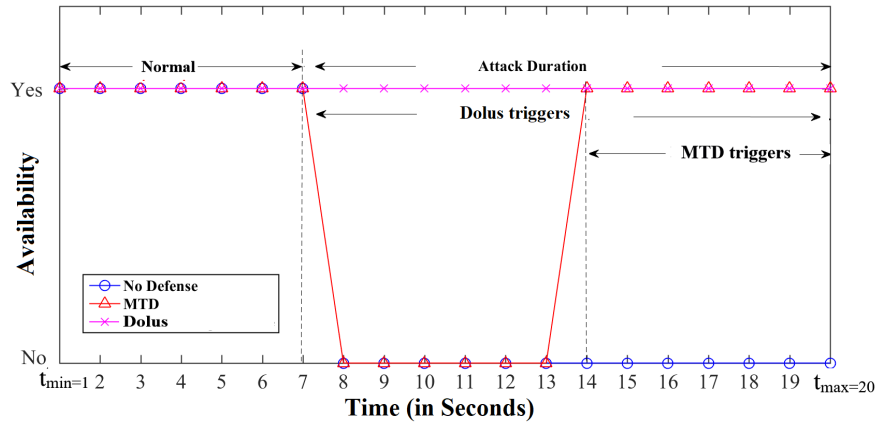


Figure 12: Comparison of the *cloud service restoration time* metric with cases of: no Defense, with MTD and with Dolus.




Figure 13: Traffic processed (in Bytes) in one of the slave switches.

644 Table 3.

645 While the two-stage ensemble learning scheme is effective in detecting test
 646 data, a new attack that has not been used in training could initially go undetected
 647 and impact services. However, with pertinent labeling of attack traffic flows dur-
 648 ing training, the accuracy of the ensemble learning scheme can be improved sig-
 649 nificantly. We depict the outlier detection and classification for a trained case
 650 in Figure 11, where we make use of 60% of the data as training data and 40%




Figure 14: Traffic Processed at the root switch only shows user traffic proving that the attack traffic is redirected to quarantine VM.

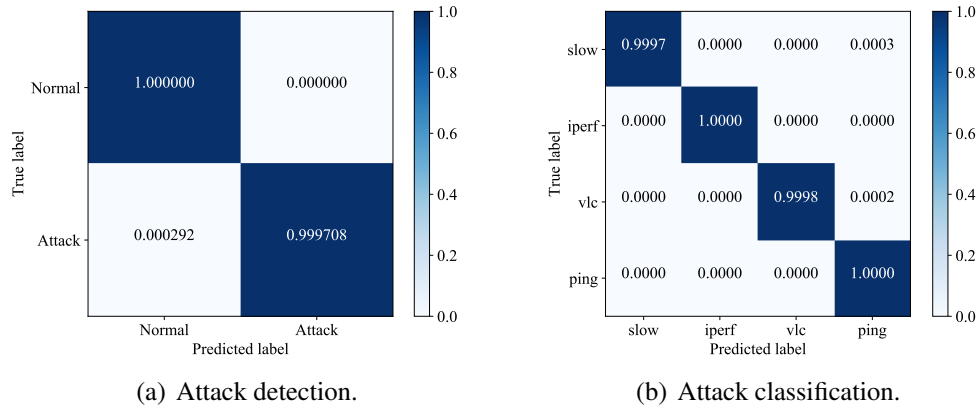


Figure 15: Confusion matrices for attack detection and classification for multiple traffic flows sent to a single server.

651 as test data for the same dataset used in the 2nd instance. For the purpose of our
 652 evaluation, the sorted dataset has randomized time stamps.

653 Though the dataset that we use is discrete with differences in traffic such as
 654 protocol, bytes transmitted, number of packets, source and destination addresses,
 655 our two-stage ensemble learning scheme is effective in detecting the attacks with
 656 good accuracy and efficiency. The ensemble learning scheme can further be mod-
 657 ified based on other characteristics of network traffic, and such modifications are

658 beyond the scope of the work in this paper.

659 6.1.3. Time to Restore a Cloud-hosted Application Service

660 Figure 12 compares the time taken by our Dolus system to stop a DDoS at-
661 tack versus MTD-based and no defense strategies. After a warm-up period of
662 6 seconds, we start the SlowHTTPTest and hping3 at the 7th second from the at-
663 tackers. In a SDxI-based cloud network with no defense strategy, the services are
664 immediately affected by the attack traffic. Consequently, an absence of service
665 availability after the 7th second as shown in the graph results in a situation where
666 cloud service restoration does not occur. MTD-based defense strategy is able to
667 restore the service after taking ~ 6 seconds to mitigate the attack traffic impact.
668 However, our Dolus system supported service on the other hand, does not suffer
669 from any loss of availability in comparison with the other two strategies. This is
670 due to the sharing of attack intelligence between the slave switches and redirec-
671 tion of attack traffic to quarantine VMs closer to the attackers, making the cloud
672 network completely oblivious to the attackers.

673 6.1.4. Amount of Traffic Processed at the Root Switch

674 Figures 13 and 14 depict the amount of traffic processed (in Bytes) at one of
675 the slave switches and the root switch. From Figure 14, it is evident that the SDxI-
676 based cloud network is oblivious to the attack traffic impact, complementing the
677 result in Figure 12. Since the slave switch represented in Figure 14 redirects attack
678 traffic to the quarantine VMs, we observe a 5X increase in the amount of traffic
679 processed in comparison with the root switch.

680 Overall, we find that our Dolus can effectively detect DDoS attack and redirect
681 traffic in real-time i.e., on the order of seconds depending on the knowledge of
682 the DDoS attack pattern, and block it closer to the attack source in 1-2 seconds
683 if automated policy updates are possible in the cross-domain setting. However,
684 if humans need to be brought into the loop, the time to block the attack can be
685 adjusted so that there is enough time for cross-domain manual coordination during
686 which an effective pretense of the quarantine VM is deceiving the attacker with a
687 false sense of success.

688 6.2. Dolus Experiments for APT Attack Defense

689 6.2.1. Testbed Setup

690 For the purposes of APT attack detection and defense, a modified GENI Cloud
691 testbed was setup as shown in Figure 16. The purpose of the GENI Cloud Testbed
692 is to simulate a collaborative cross-domain SDxI architecture with the core servers

693 and services located at the switch-root in the Clemson InstaGENI (blue) domain.
 694 Correspondingly, the user traffic originates from three other separate domains
 695 with two distinct paths to where the services are located. Since an APT is not
 696 a distributed attack, there was no need to consider multiple attack vectors from
 697 many directions. However, due to the nature of an APT attack being secretive and
 698 stealthy, we assume that an APT can be hiding anywhere in an SDxI. Our testbed
 699 is comprised of multiple open vSwitches (a slaves and a single root), numerous
 700 nodes (which are hosts), and a controller. The slave switches connect all the user
 701 nodes, and the root switch connects all the servers hosting the application system
 702 and related services to the slave switches. The controller in the setup is a stand-
 703 alone node, running the monitor and policy updaters, calculating SS thresholds
 704 for nodes and the overall network, managing all the traffic and defense by pretense
 705 mechanisms of the Dolus system.

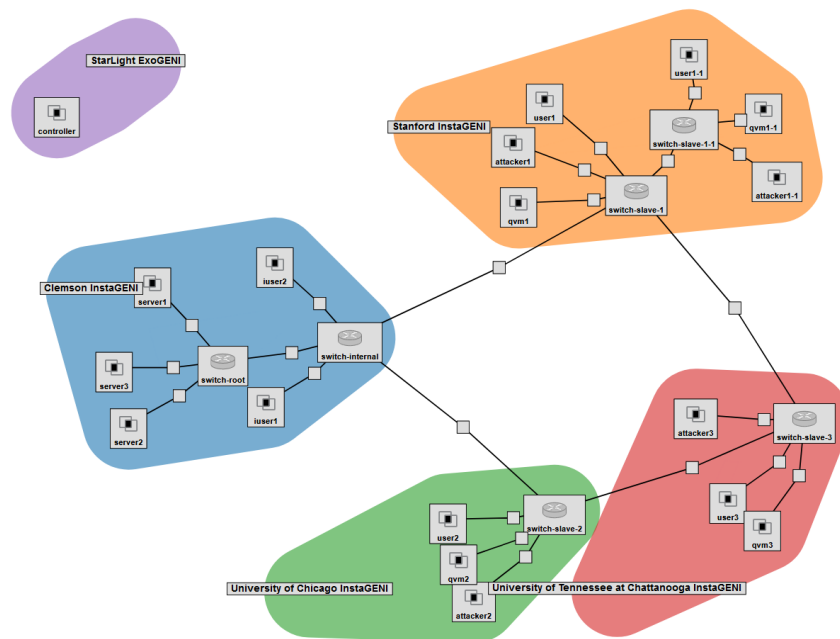


Figure 16: GENI Cloud testbed used to evaluate Dolus for APT attack defense.

706 6.2.2. Suspiciousness Score Calculation Results

707 In the first experiment, we randomly selected three hosts, and compromised
 708 them by running `slowhttp` attacks from attacker 1 and attacker 3, and a secure

Table 4: Suspiciousness Scores before Whitelisting

Node	Command	Score
Attacker1	slowhttp	8.8
Attacker2	scp	215.5
Attacker3	slowhttp	18.0
Server1	ping	17.3
Server2	Traffic Response	16.4
Server3	iperf -s	9.0
User1	iperf -c	5645.7
User2	wget	200.7

709 copy (`scp`) from attacker 2. This configuration allows us to compare the suspi-
710 ciousness between a DDoS attack, and a file exfiltration attack. Before running
711 the experiment, we specify minimum and maximum values for flows, connections,
712 and bytes: the user and attacker nodes are each set to a minimum of 1 and a max-
713 imum of 10 connections, a minimum of 100 and a maximum of 1,000 flows, and
714 a minimum of 10 and a maximum of 100,000 bytes. The servers had a minimum
715 of 10 and a maximum of 1000 connections, a minimum of 1,000 and maximum
716 of 10,000 flows, and minimum of 100,000 and a maximum of 100000000 bytes.

717 From the controller, we obtain the SS for these three attackers before (see
718 Table 4) and after (see Table 5) whitelisting. Note that all devices have SS calcu-
719 lated for them, as we don't initially whitelist any devices or traffic on our testbed
720 network. Attacker 2 exhibited the highest SS out of three, due to data exfiltra-
721 tion [10]. The traffic that is being exfiltrated generates a much higher score than
722 the regular traffic in the network.

723 The purpose of whitelisting is to allow administrators to ignore traffic, which
724 is not going outside of the network. For example, if we consider that both server
725 1 and user 1 are within our own network, then any data transmitted between those
726 two machines would not be data being exfiltrated from the enterprise network.
727 Therefore, we can consider such traffic as benign. However, whenever we con-
728 sider attacker 1 and server 1, since attacker 1 is compromised, we consider all
729 traffic from attacker 1 to be possible data exfiltrated from the enterprise network.

Table 5: Suspiciousness Scores after Whitelisting

Node	Command	Score
Attacker1	slowhttp	8.8
Attacker2	scp	215.5
Attacker3	slowhttp	18.0

730 Furthermore, we consider a case where - if attacker 1 compromised user 1 within
 731 our network and then used user 1 to exfiltrate data from server 1 to user 1 then
 732 from user 1 to attacker 1. In such a case, we are able to detect the suspicious-
 733 ness between user 1 and attacker 1 since that is where the actual data exfiltration
 734 is taking place. As you can see in Table 5, we ignore the traffic between users
 735 and servers, even though there was data moving between them (as seen in Table
 736 4). Moreover, by considering the whitelisting prior to the Suspiciousness Score
 737 calculations, we decrease the overall time spent on speed of the calculations since
 738 we will need to calculate scores for *only* a portion of the network.

739 6.2.3. Targeted Suspiciousness Score Effectiveness

740 As an extension of our overall Suspiciousness Scores calculation, we also per-
 741 formed experiments using our novel Targeted Suspiciousness Scores detailed in
 742 Section 5.2. Similar to the previous experiment, we tested a variety of different
 743 types of network traffic with the addition of a two new types i.e., BitTorrent and
 744 cryptocurrency mining traffic. For generating the cryptocurrency mining traffic,
 745 we use a miner software called ‘geth’ which could be used for CPU resource ex-
 746 filtration. We also use variants of geth, where we limit the geth mining using
 747 various tools for both network rate limiting as well as CPU limiting. Both these
 748 variants could mimic the tools used by an attacker attempting to stealthily exfil-
 749 trate resources from an enterprise system without being detected. In addition to
 750 the cryptocurrency mining attack traffic, we also generated several types of be-
 751 nign, other attack, and suspicious traffic. We expected the BitTorrent traffic in
 752 particular to have similar characteristics and behave similar to the cryptocurrency
 753 mining traffic, since both use peer-to-peer protocols on distributed systems.

754 Figure 17 shows the Overall Suspiciousness Score results from the different
 755 tests we ran over a ten minute period. We surprisingly found that the Overall
 756 Suspiciousness Scores indicate that the cryptocurrency mining traffic is far less

Table 6: Test traffic types for Targeted Suspiciousness Score effectiveness evaluation experiments.

Test Traffic Types	
Devices	Number of traces
geth	Cryptocurrency mining
geth	Cryptocurrency mining with trickle (network rate limited to 10 kbps download and upload)
geth	Cryptocurrency mining with cpulimit (limited to 10% of total cpu)
ping	Ping traffic on one of the IP addresses contacted during crypto mining
wget	Traffic of a 121 MB video file download
scp	Traffic of the same 121 MB video file upload
DDoS	slowhttptest traffic against the same IP address used for scp
BitTorrent	Up to 200 connections relating to downloading/seeding Ubuntu 17.10 iso

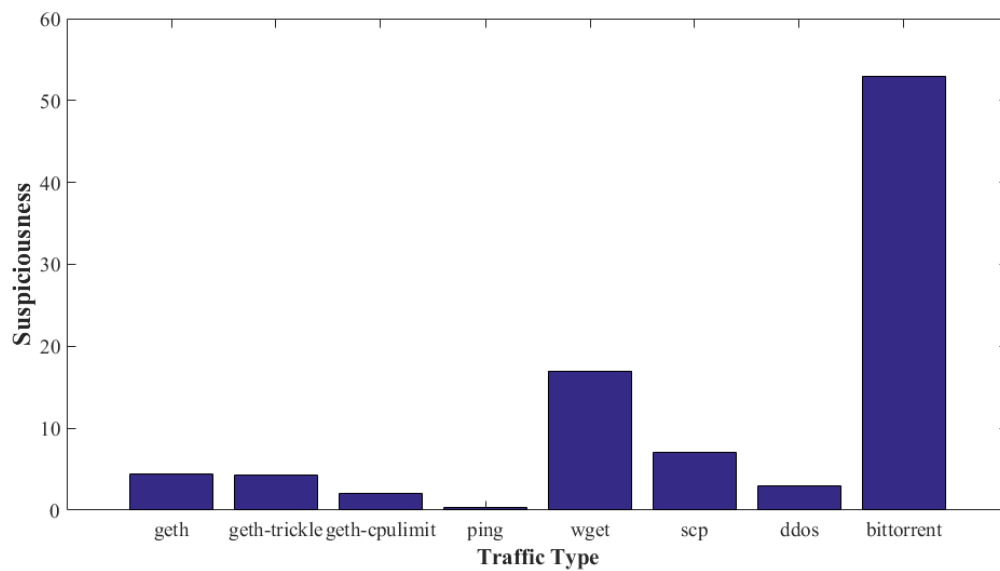


Figure 17: Overall Suspiciousness Score results that motivate the need for using targeted suspiciousness to improve APT attack detection accuracy.

757 suspicious than the BitTorrent traffic, and even the wget and scp traffic. This
 758 is mainly due to the fact that the suspiciousness calculations in the general case
 759 were created to detect data exfiltration attacks, and do not account for the resource
 760 exfiltration characteristics of a cryptocurrency mining attack. This motivated us to
 761 reconsider the use of Overall Suspiciousness Scores, and introduce weights for a
 762 variety of suspicious traffic in order to minimize the RMSE in the attack detection
 763 accuracy.




Figure 18: Distinct IP Addresses contacted per every 15 seconds.

764 Digging deeper as shown in Figure 18, we observed that - even though the
 765 Overall Suspiciousness Score for the benign bittorrent traffic was higher than
 766 cryptocurrency mining traffic, there is a distinct difference in the traffic varia-
 767 tion over time. The cryptocurrency mining traffic was far more varied in total
 768 distinct IP addresses contacted at any given period of time. We also saw highly
 769 similar results with the total flows over a given time period, and also for the total
 770 bytes transmitted over the same time period. This led us to conclude that there is
 771 a distinct difference in the variation of traffic over time when comparing BitTor-
 772 rent and cryptocurrency mining traffic. Specifically, BitTorrent traffic will feature
 773 connections with a large number of IP addresses but will continue to maintain con-
 774 nections with the same number of IP addresses every few seconds with little or no
 775 change. In contrast, the cryptocurrency mining traffic will connect with many IP
 776 addresses, and the total number of distinct connections at any given period of time

777 will fluctuate over a wide range. With the knowledge of such attack traffic fea-
 778 tures, suitable weights can be assigned in Targeted Suspiciousness Scores that will
 779 not trigger BitTorrent traffic on a network as suspiciousness, but will accurately
 780 detect unauthorized cryptocurrency mining traffic that are part of APM attacks.




Figure 19: Demonstration of a Fading Pretense policy implementation to deter an APM attack after its detection on a compromised device.

781 6.2.4. Fading Pretense Policy Demonstration

782 The Fading Pretense policy can be implemented as an effective defense mech-
 783 anism against APM attacks as described earlier in Section 5.2. Herein, we de-
 784 scribe a demonstration of an experiment that illustrates how the fading pretense
 785 policy could deter an APM attacker in practice. Figure 19 shows the experiment
 786 conducted using a safe system zone, and an attacked system zone. The x-axis is an
 787 arbitrary unit of time progression that can be configured (on the order of several
 788 minutes, hours or even days) by the system/network administrator depending on
 789 the duration of the pretense policy being in effect. Behavioral psychology con-
 790 siderations also could be factored into determining the rate of progressive decline
 791 of the resource allocation on a compromised device. In any case, we can observe
 792 that three distinct phases of pretense should occur:

793 *(Phase-1): Fading pretense begins.* Assuming at time $x=1$, a resource exfiltration
 794 attack is detected to be occurring in an attacked system zone, at which point the
 795 attacker will have 100% availability to the system resources, and the fading pre-
 796 tense policy is initiated at time $x=2$.

797 *(Phase-2): Attacker is deterred.* In the time after the fading pretense policy is in
 798 effect, the availability of the system resources is reduced progressively to 90%,
 799 60%, 30% and 10 % to ultimately cause the attacker to consider a redirection of
 800 the APM attack to a more-potent device with higher resource allocations.

801 *(Phase-3) Safe system zone restoration.* Once the attacker is found to have been
 802 deterred at time $x=6$, the previously compromised device can be added to into the
 803 safe system zone with 100% availability of system resources. We remark that the
 804 safe system zone restoration should be performed only after suitable patching or
 805 system re-imaging in order to ensure that there is no re-occurrence of the APM
 806 attack on that particular device in the future.

807 6.2.5. Time Overhead for Suspiciousness Score Calculation

808 Table 7 shows the time taken by ADAPTs to calculate the ss for devices, each
 809 running on a single core. It also shows the number of traces, and their correspond-
 810 ing processing times. As high as 1.8 million packets for 8 devices can be processed
 811 under 100 seconds, which demonstrates the efficacy of ADAPTs. However, there
 812 is a linear increase in time as the number of traces grow. If such a linear increase
 813 does not meet the threat monitoring objectives of a domain, a parallel implementa-
 814 tion of ADAPTs can be extended and used on nodes with multi-core functionality
 815 to reduce the computation times in the Suspiciousness Score calculations.

Table 7: Processing time taken by ADAPT with single threaded processing

Single Threaded		
Devices	Number of traces	Time (in seconds)
3	590,492	50
6	1,249,490	77
8	1,839,982	94

816 7. Conclusion

817 Recent innovations in the orchestration of cloud resources are fueled by emer-
818 gence of the Software-Defined everything Infrastructure (SDxI) paradigm. At the
819 same time, the sophistication of targeted attacks such as Distributed Denial-of-
820 Service (DDoS) attacks and Advanced Persistent Threat (APT) attacks are grow-
821 ing on an unprecedented scale. Consequently, online businesses in retail, health-
822 care and other fields are under constant threat of targeted attacks. In this paper,
823 we presented a novel defense system called *Dolus* to mitigate the impact of DDoS
824 and APT attacks against high-value services hosted in SDxI-based cloud plat-
825 forms. We proposed a *defense by pretense* mechanism that can be used during
826 defense against targeted attacks, which involves threat detection algorithms based
827 on a number of attack vector features. Using blacklisting information, our pre-
828 tense initiation builds upon pretense theory concepts in child play psychology to
829 trick an attacker through creation of a false sense of success.

830 Our above approach for DDoS and APT attacks defense takes advantage of
831 elastic capacity provisioning in cloud platforms to implement moving target de-
832 fense techniques that does not affect the cloud-hosted application users, and con-
833 tains the attack traffic in a quarantine VM(s). With the time gained through ef-
834 fective pretense initiation in the case of DDoS attacks, cloud service providers
835 could coordinate across a unified SDxI infrastructure involving multiple ASes to
836 decide on policies that help in blocking the attack flows closer to the source side.
837 Performance evaluation results of our Dolus system in a GENI cloud testbed for
838 DDoS attacks show that our approach can be effective in filtering, detection and
839 implementation of SDxI-based infrastructure policy coordination for mitigation
840 of the impact of the DDoS attacks. In addition, we also showed how the Do-
841 lus system can be an effective defense using pretense against APTs and APMs.
842 Using the general Suspiciousness Scores and a novel Targeted Suspiciousness
843 Score concept, we proposed novel threat intelligence collection and accurate at-
844 tack detection of subtle and secretive targeted attacks at a device level and also
845 at a network-wide level. Further, we found that our Admin UI capability can
846 greatly help network operators and cloud service providers to overcome their dif-
847 ficulty in determining which devices on an enterprise network or a cloud service
848 deployment may be compromised. Lastly, we demonstrated how a pertinent de-
849 fense strategy such as a fading pretense policy can be effective in the mitigation of
850 APM attacks that target resource exfiltration within an SDxI-based infrastructure.

851 Future work can be pursued to investigate more sophisticated pretense schemes
852 that use threat intelligence collection on effectiveness of a working pretense, and

853 initiate more involved adaptations. In addition, data analytics extensions can be
854 pursued for more sophisticated targeted attacks with significantly larger number
855 of features that need to be involved in effective detection and defense schemes.

856 **Acknowledgement**

857 This material is based upon work supported by the National Science Founda-
858 tion under Award Number: CNS-1205658. Any opinions, findings, and conclu-
859 sions or recommendations expressed in this publication are those of the author(s)
860 and do not necessarily reflect the views of the National Science Foundation.

861 **References**

- 862 [1] Verizon. State of the Market: Enterprise Cloud,
863 [http://www.verizonenterprise.com/resources/reports/rp_state-of-the-market-](http://www.verizonenterprise.com/resources/reports/rp_state-of-the-market-enterprise-cloud-2016_en_xg.pdf)
864 [enterprise-cloud-2016_en_xg.pdf](http://www.verizonenterprise.com/resources/reports/rp_state-of-the-market-enterprise-cloud-2016_en_xg.pdf), 2017.
- 865 [2] SDxCentral, SDxCentral. Software Defined Everything Part 3: SDx
866 infrastructure, [https://www.sdxcentral.com/cloud/definitions/software-](https://www.sdxcentral.com/cloud/definitions/software-defined-everything-part-3-sdx-infrastructure)
867 [defined-everything-part-3-sdx-infrastructure](https://www.sdxcentral.com/cloud/definitions/software-defined-everything-part-3-sdx-infrastructure), 2017.
- 868 [3] A. Gupta, N. Feamster, L. Vanbever, Authorizing network control at soft-
869 ware defined internet exchange points, in: Proceedings of the Symposium
870 on SDN Research, SOSR '16, ACM, New York, NY, USA, 2016, pp. 16:1–
871 16:6.
- 872 [4] P. Chen, L. Desmet, C. Huygens, A study on advanced persistent threats, in:
873 IFIP International Conference on Communications and Multimedia Security,
874 Springer, pp. 63–72.
- 875 [5] P. Kampanakis, H. G. Perros, T. Beyene, Sdn-based solutions for moving
876 target defense network protection., in: WoWMoM, IEEE Computer Society,
877 2014, pp. 1–6.
- 878 [6] S. Debroy, P. Calyam, M. Nguyen, A. Stage, V. Georgiev, Frequency-
879 minimal moving target defense using software-defined networking, 2016
880 International Conference on Computing, Networking and Communications
881 (ICNC) 00 (2016) 1–6.

- 882 [7] N. Virvilis, D. Gritzalis, T. Apostolopoulos, Trusted computing vs. advanced
883 persistent threats: Can a defender win this game?, in: Ubiquitous intel-
884 ligence and computing, 2013 IEEE 10th international conference on and
885 10th international conference on autonomic and trusted computing (uic/atc),
886 IEEE, pp. 396–403.
- 887 [8] S. Nichols, S. Stich, A cognitive theory of pretense, *Cognition* 74 (2000)
888 115–147.
- 889 [9] J. W. V. de Vondervoort, O. Friedman, Young children protest and correct
890 pretense that contradicts their general knowledge, *Cognitive Development*
891 43 (2017) 182–189.
- 892 [10] M. Marchetti, F. Pierazzi, M. Colajanni, A. Guido, Analysis of high vol-
893 umes of network traffic for advanced persistent threat detection, *Computer*
894 *Networks* 109 (2016) 127–141.
- 895 [11] J. Leyden, CryptoLurker hacker crew skulk about like cyberspies, earn \$\$\$,
896 https://www.theregister.co.uk/2018/03/06/apt_cryptomining, 2018.
- 897 [12] M. Berman, J. S. Chase, L. Landweber, A. Nakao, M. Ott, D. Raychaud-
898 huri, R. Ricci, I. Seskar, Geni: A federated testbed for innovative network
899 experiments, *Comput. Netw.* 61 (2014) 5–23.
- 900 [13] A. Clark, K. Sun, R. Poovendran, Effectiveness of IP address randomization
901 in decoy-based moving target defense, in: *Proceedings of the 52nd IEEE*
902 *Conference on Decision and Control, CDC 2013, December 10-13, 2013,*
903 *Firenze, Italy*, pp. 678–685.
- 904 [14] T. Sochor, M. Zuzcak, Study of internet threats and attack methods using
905 honeypots and honeynets, in: *International Conference on Computer Net-*
906 *works*, Springer, pp. 118–127.
- 907 [15] S. Seufert, D. O’Brien, Machine learning for automatic defence against
908 distributed denial of service attacks, in: *2007 IEEE International Conference*
909 *on Communications*, pp. 1217–1222.
- 910 [16] J. Choi, C. Choi, B. Ko, P. Kim, A method of ddos attack detection using
911 http packet pattern and rule engine in cloud computing environment, *Soft*
912 *Computing* 18 (2014) 1697–1703.

- 913 [17] T. Thapngam, S. Yu, W. Zhou, S. K. Makki, Distributed denial of service
914 (ddos) detection by traffic pattern analysis, *Peer-to-Peer Networking and*
915 *Applications* 7 (2014) 346–358.
- 916 [18] Y. Chen, S. Jain, V. K. Adhikari, Z. L. Zhang, K. Xu, A first look at inter-data
917 center traffic characteristics via Yahoo! datasets, pp. 1620–1628.
- 918 [19] L. Yang, T. Zhang, J. Song, J. S. Wang, P. Chen, Defense of ddos attack
919 for cloud computing, in: *2012 IEEE International Conference on Computer*
920 *Science and Automation Engineering (CSAE)*, volume 2, pp. 626–629.
- 921 [20] Internet2’s ddos mitigation strategy, <https://www.internet2.edu/blogs/detail/12234>,
922 2016.
- 923 [21] Y. Choi, Implementation of content-oriented networking architecture (cona):
924 a focus on ddos countermeasure.
- 925 [22] C. YuHunag, T. MinChi, C. YaoTing, C. YuChieh, C. YanRen, A novel
926 design for future on-demand service and security, in: *2010 IEEE 12th Inter-*
927 *national Conference on Communication Technology*, pp. 385–388.
- 928 [23] S. Shin, P. A. Porras, V. Yegneswaran, M. W. Fong, G. Gu, M. Tyson, Fresco:
929 Modular composable security services for software-defined networks., in:
930 *NDSS*, The Internet Society, 2013.
- 931 [24] Y. Yu, C. Qian, X. Li, Distributed and collaborative traffic monitoring in
932 software defined networks, in: *Proceedings of the Third Workshop on Hot*
933 *Topics in Software Defined Networking, HotSDN ’14*, ACM, New York,
934 NY, USA, 2014, pp. 85–90.
- 935 [25] H. Tian, J. Bi, An incrementally deployable flow-based scheme for ip trace-
936 back., *IEEE Communications Letters* 16 (2012) 1140–1143.
- 937 [26] C. Clark, K. Fraser, S. Hand, J. G. Hansen, E. Jul, C. Limpach, I. Pratt,
938 A. Warfield, Live migration of virtual machines, in: *Proceedings of the 2nd*
939 *Conference on Symposium on Networked Systems Design & Implementa-*
940 *tion - Volume 2, NSDI’05*, USENIX Association, Berkeley, CA, USA, 2005,
941 pp. 273–286.
- 942 [27] Q. Yan, F. R. Yu, Q. Gong, J. Li, Software-defined networking (sdn) and dis-
943 tributed denial of service (ddos) attacks in cloud computing environments:

- 944 A survey, some research issues, and challenges, *IEEE Communications Sur-*
945 *veys Tutorials* 18 (2016) 602–622.
- 946 [28] J. Kim, T. Lee, H.-g. Kim, H. Park, Detection of advanced persistent threat
947 by analyzing the big data log, *Cloud Security Alliance* 13 (2013) 101–107.
- 948 [29] P. K. Sharma, S. Y. Moon, D. Moon, J. H. Park, Dfa-ad: a distributed frame-
949 work architecture for the detection of advanced persistent threats, *Cluster*
950 *Computing* 20 (2017) 597–609.
- 951 [30] P. Bhatt, E. T. Yano, P. Gustavsson, Towards a framework to detect multi-
952 stage advanced persistent threats attacks, in: *Service Oriented System En-*
953 *gineering (SOSE), 2014 IEEE 8th International Symposium on*, IEEE, pp.
954 390–395.
- 955 [31] B. Binde, R. McRee, T. J. OConnor, Assessing outbound traffic to uncover
956 advanced persistent threat, SANS Institute. Whitepaper (2011).
- 957 [32] J. de Vries, H. Hoogstraaten, J. van den Berg, S. Daskapan, Systems for
958 detecting advanced persistent threats: A development roadmap using intel-
959 ligent data analysis, in: *Cyber Security (CyberSecurity), 2012 International*
960 *Conference on*, IEEE, pp. 54–61.
- 961 [33] J. Vukalović, D. Delija, Advanced persistent threats-detection and defense,
962 in: *Information and Communication Technology, Electronics and Micro-*
963 *electronics (MIPRO), 2015 38th International Convention on*, IEEE, pp.
964 1324–1330.
- 965 [34] G. Vert, A. L. Claesson-Vert, J. Roberts, E. Bott, A technology for detection
966 of advanced persistent threat in networks and systems using a finite angular
967 state velocity machine and vector mathematics, in: *Computer and Network*
968 *Security Essentials*, Springer, 2018, pp. 41–64.
- 969 [35] C. FIU, 25th geni engineering conference (gec25)-part iii, 2017.
- 970 [36] M. Lee, D. Lewis, Clustering disparate attacks: mapping the activities of the
971 advanced persistent threat, Last accessed June 26 (2013).
- 972 [37] P. Giura, W. Wang, Using large scale distributed computing to unveil ad-
973 vanced persistent threats, *Science J* 1 (2012) 93–105.


- 974 [38] Z. Saud, M. H. Islam, Towards proactive detection of advanced persistent
975 threat (apt) attacks using honeypots, in: Proceedings of the 8th International
976 Conference on Security of Information and Networks, ACM, pp. 154–157.
- 977 [39] N. Virvilis, D. Gritzalis, The big four-what we did wrong in advanced per-
978 sistent threat detection?, in: Availability, Reliability and Security (ARES),
979 2013 Eighth International Conference on, IEEE, pp. 248–254.
- 980 [40] Z. Zulkefli, M. M. Singh, N. H. A. H. Malim, Advanced persistent threat
981 mitigation using multi level security–access control framework, in: Interna-
982 tional Conference on Computational Science and Its Applications, Springer,
983 pp. 90–105.
- 984 [41] E. M. Hutchins, M. J. Cloppert, R. M. Amin, Intelligence-driven computer
985 network defense informed by analysis of adversary campaigns and intrusion
986 kill chains, *Leading Issues in Information Warfare & Security Research 1*
987 (2011) 80.
- 988 [42] P. Hu, H. Li, H. Fu, D. Cansever, P. Mohapatra, Dynamic defense strategy
989 against advanced persistent threat with insiders, in: *Computer Communica-*
990 *tions (INFOCOM), 2015 IEEE Conference on, IEEE, pp. 747–755.*
- 991 [43] I. Radware, Defenseflow - sdn based network, application ddos and apt pro-
992 tection, 2014.
- 993 [44] M. Ammar, M. Rizk, A. Abdel-Hamid, A. K. Aboul-Seoud, A framework
994 for security enhancement in sdn-based datacenters, in: *New Technologies,*
995 *Mobility and Security (NTMS), 2016 8th IFIP International Conference on,*
996 *IEEE, pp. 1–4.*
- 997 [45] M. Saher, J. Pathak, Malware and exploit campaign detection system and
998 method, 2014. US Patent App. 14/482,696.
- 999 [46] N. Foster, R. Harrison, M. J. Freedman, C. Monsanto, J. Rexford, A. Story,
1000 D. Walker, Frenetic: A network programming language, in: *ACM Sigplan*
1001 *Notices, volume 46, ACM, pp. 279–291.*
- 1002 [47] W. M. Eddy, *Tcp syn flooding attacks and common mitigations (2007).*
- 1003 [48] F. Gont, *Icmp attacks against tcp (2010).*

- 1004 [49] K. Tools, hping3. ICMP or SYN flooding tool,
1005 <https://tools.kali.org/information-gathering/hping3>, 2014.
- 1006 [50] S. Shekyan, SlowHTTPTest. Application Layer DoS attack ,
1007 <https://github.com/shekyan/slowhttpstest/wiki>, 2011.
- 1008 [51] Y. Zhang, S. Debroy, P. Calyam, Network-wide anomaly event detection
1009 and diagnosis with perfsonar, *IEEE Transactions on Network and Service*
1010 *Management* 13 (2016) 666–680.
- 1011 [52] Y. Zhang, P. Calyam, S. Debroy, M. Sridharan, Pca-based network-wide cor-
1012 related anomaly event detection and diagnosis, in: 2015 11th International
1013 Conference on the Design of Reliable Communication Networks (DRCN),
1014 pp. 149–156.
- 1015 [53] R. S. Boyer, J. S. Moore, Mjrty—a fast majority vote algorithm, in: *Auto-*
1016 *mated Reasoning*, Springer, 1991, pp. 105–117.
- 1017 [54] T. G. Dietterich, et al., Ensemble methods in machine learning, *Multiple*
1018 *classifier systems* 1857 (2000) 1–15.
- 1019 [55] Scapy: Packet manipulation tool, <http://www.secdev.org/projects/scapy/>,
1020 2007.
- 1021 [56] M. K. Daly, Advanced persistent threat, *Usenix*, Nov 4 (2009) 2013–2016.
- 1022 [57] W. A. M. T. Y. Enterprise, G. Hoglund, Advanced persistent threat (????).
- 1023 [58] E. Cole, Advanced persistent threat: understanding the danger and how to
1024 protect your organization, Newnes, 2012.
- 1025 [59] I. Ghafir, V. Prenosil, Advanced persistent threat attack detection: An
1026 overview, *International Journal of Advances in Computer Networks and*
1027 *Its Security (IJCNS)* (2014).
- 1028 [60] A. J. T.-F. Yen, Sherlock holmes and the case of the advanced persistent
1029 threat (2012).
- 1030 [61] M. Ask, P. Bondarenko, J. E. Rekdal, A. Nordbø, P. Bloemerus, D. Pi-
1031 atkivskiy, Advanced persistent threat (apt) beyond the hype, *Project Re-*
1032 *port in IMT4582 Network Security at Gjøvik University College*, Springer
1033 (2013).

- 1034 [62] C. J. Anderson, N. Foster, A. Guha, J.-B. Jeannin, D. Kozen, C. Schlesinger,
1035 D. Walker, Netkat: Semantic foundations for networks, in: Proceedings of
1036 the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Program-
1037 ming Languages, POPL '14.
- 1038 [63] A. Gervais, G. Karame, K. Wust, V. Glykantzis, H. Ritzdorf, S. Capkun,
1039 On the security and performance of proof of work blockchains, in: ACM
1040 SIGSAC Conference on Computer and Communications Security.
- 1041 [64] T. Neely, M. Vassel, N. Chettri, R. Neupane, P. Calyam, R. Du-
1042 rairajan, Dolus for ddos attacks defense - open repository
1043 <https://bitbucket.org/travisivart/dolus-defensebypretense>, 2018.
- 1044 [65] T. Neely, M. Vassel, N. Chettri, R. Neupane, P. Calyam,
1045 R. Durairajan, Dolus for apt attacks defense - open repository
1046 <https://github.com/travisivart/adapts>, 2018.



Roshan Lal Neupane received his MS degree in Computer Science from University of Missouri - Columbia and his BE degree in Computer Science and Engineering from Visvesvaraya Technological University, Karnataka, India. His research interests include Cloud Computing, Computer Networking, Internet of Things, and Cyber Security. He is a student member of IEEE.




Travis Neely is currently pursuing his MS degree in Computer Science from University of Missouri - Columbia. He received his BS degree in Computer Science from Maryville University of Saint Louis. His research interests include Software-defined Networking, Data Science, and Enterprise Security.



Prasad Calyam received his MS and PhD degrees from the Department of Electrical and Computer Engineering at The Ohio State University in 2002 and 2007, respectively. He is currently an Assistant Professor in the Department of Computer Science at University of Missouri-Columbia. Previously, he was a Research Director at the Ohio Supercomputer Center. His research interests include: Distributed and Cloud computing, Computer Networking, and Cyber Security. He is a Senior Member of IEEE.



Nishant Chettri received his MS in Computer Science degree from the University of Missouri - Columbia, and his BS in Computing degree from London Metropolitan University, Kathmandu, Nepal. His research interests include Mobile and Web Development, and Database Security.



Mark Vassell is currently pursuing his MS degree in Computer Science from University of Missouri - Columbia. He received his BS degree in Computer Science from University of Missouri - Columbia. His research interests include IoT technologies for public safety, Data Science, and Cyber Security.



Ramakrishnan Durairajan received his MS and PhD degrees from the Department of Computer Science at the University of Wisconsin - Madison in 2014 and 2017, respectively. He is currently an Assistant Professor in the Department of Computer Science at University of Oregon. His research interests include: Internet Measurements, Computer Networking, and Cyber Security.