

## Manuscript Details

<b>Manuscript number</b>	FGCS_2018_716
<b>Title</b>	Intelligent Defense using Pretense against Targeted Attacks in Cloud Platforms
<b>Article type</b>	Full Length Article

### 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 (SDxl). 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 SDxl-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 SDxl 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 SDxl 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

**Corresponding Author** Prasad Calyam

**Corresponding Author's Institution** University of Missouri-Columbia

**Order of Authors** Roshan Lal Neupane, Travis Neely, Prasad Calyam, Nishant Chettri, Mark Vassell, Ramakrishnan Durairajan

**Suggested reviewers** Kaiqi Xiong, Sejun Song, Flavio Esposito, Khaled Salah

## Submission Files Included in this PDF

### File Name [File Type]

highlights.pdf [Highlights]

dolus-intelligent-defense-pretense-033118.pdf [Manuscript File]

Author\_Bios.pdf [Author Biography]

## Submission Files Not Included in this PDF

### File Name [File Type]

photos.zip [Author Photo]

To view all the submission files, including those not included in the PDF, click on the manuscript title on your EVISE Homepage, then click 'Download zip file'.

# Intelligent Defense using Pretense against Targeted Attacks in Cloud Platforms

Cover Letter for Submission to Elsevier FGCS Journal, March 2018;  
Special Issue on Cyber Threat Intelligence and Analytics

Roshan Lal Neupane, Travis Neely, Prasad Calyam, Nishant Chettri  
Mark Vassell, and Ramakrishnan Durairajan

Department of Electrical Engineering and Computer Science, University of Missouri, Columbia, MO 65211 USA  
University of Oregon, Eugene, OR 97403 USA

Email: {rlnzq8, nc5ff, mdvy96}@mail.missouri.edu, {calyamp, neelyt}@missouri.edu, ram@cs.uoregon.edu

This document is a cover letter provided in response to the following requirement: Authors are required to include with their submission a letter in which they identify all prior publications on which their submission may be based, provide pointers to publicly available versions of those publications, and articulate any changes made to improve and/or expand on those conference publications.

## NOVEL CONTRIBUTIONS

An earlier version of this article was published in the proceedings of the ACM 19th International Conference of Distributed Computing and Networking (ICDCN) [1]. This manuscript extends the conference version with an extra  $\approx 60\%$  of new material: We extend novelty of our Dolus system for DDoS defense using pretense, from our ICDCN 2018 publication, by including the contribution of an automation defense mechanism against Advanced Persistent Threats (i.e., ADAPTs) to safeguard cloud-hosted services against APT attacks.

- The Title, Abstract, Introduction (Section 1) sections were rewritten to reflect the novelty of the enhanced pretense theory, new ADAPTs work extensions and experiment findings.
- The Related Work (Section 2) section also has updated descriptions and the manuscript now has 28 new references that include related state-of-the-art approaches.
- The Dolus Defense Methodology (Section 3) has extended information on pretense theory, and its use in the design of our Dolus system for both DDoS and APT attacks defense.
- A new section “APT Attack Defense with Dolus” (Section 5) is added in the manuscript that includes the APT attack model and a novel defense by pretense scheme.
- In addition, new experiment testbed setup and experiment results are added in the Performance Evaluation (Section 6) section.
- Lastly, the Conclusion (Section 7) section is updated with new findings and future work suggestions.
- The source code of our Dolus implementations and experiment scripts for DDoS and APT attacks defense are publicly available under GNU license at [2] and [3].

## REFERENCES

- [1] Roshan Lal Neupane, Travis Neely, Nishant Chettri, Mark Vassell, Yuanxun Zhang, Prasad Calyam, and Ramakrishnan Durairajan. Dolus: Cyber defense using pretense against ddos attacks in cloud platforms. In *ICDCN*, 2018.
- [2] Travis Neely, Mark Vassel, Nishant Chettri, Roshan Neupane, Prasad Calyam, and Ramakrishnan Durairajan. Dolus for ddos attacks defense - open repository <https://bitbucket.org/travisivart/dolus-defensebypretense>, 2018.
- [3] Travis Neely, Mark Vassel, Nishant Chettri, Roshan Neupane, Prasad Calyam, and Ramakrishnan Durairajan. Dolus for apt attacks defense - open repository <https://github.com/travisivart/adapts>, 2018.

# Intelligent Defense using Pretense against Targeted Attacks in Cloud Platforms

Roshan Lal Neupane<sup>1a</sup>, Travis Neely<sup>1b</sup>, Prasad Calyam<sup>b</sup>, Nishant Chettri<sup>a</sup>, Mark Vassell<sup>a</sup>, Ramakrishnan Durairajan<sup>c</sup>

<sup>a</sup>{rlnzq8, nc5ff, mdvy96}@mail.missouri.edu

<sup>b</sup>{neelyt, calyamp}@missouri.edu

<sup>c</sup>ram@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

---

<sup>1</sup> 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 (ADAPT). Our goal in ADAPT’s 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 We evaluate the efficacy of our Dolus using two GENI Cloud [11] testbeds,  
113 one for DDoS detection and the other for APT attacks detection. The DDoS de-  
114 tection testbed contains three SDN switches, two slave switches and a single root  
115 switch. The slave switches are each attached to users and attackers, a quarantine  
116 VM, and a connection to the root switch. Likewise, the root switch is connected  
117 to elastic VMs, each of which could serve as a candidate for the target application  
118 (i.e., a video gaming portal) hosting that could be compromised by the attackers.  
119 All switches are connected to a unified SDN controller located in the cloud ser-  
120 vice provider domain, which directs the policy updates. Our experiment results  
121 demonstrate the real-time capabilities of our Dolus system to: (a) detect DDoS  
122 attacks and redirect attack traffic to quarantine resources to engage the attacker  
123 under pretense, and (b) coordinate SDxI policies to possibly block DDoS attacks  
124 closer to the attack source(s) without affecting the (benign) cloud users/customers.  
125 The APT detection testbed with ADAPTs is similar to the DDoS detection testbed  
126 but features dynamic traffic manipulation, monitoring, and analysis to calculate  
127 Suspiciousness Scores for each device on the network.

128 Another testbed development contribution that is used in our evaluation exper-  
129 iments is an Administrative User Interface (Admin UI) which we developed for a  
130 network administrator to defend against targeted attacks. The Admin UI acts as a  
131 central analytics hub for defense against both types of targeted attacks considered  
132 in this paper. The Admin UI informs the administrator of total bytes being trans-  
133 mitted through each switch connected to the controller. The administrator can  
134 also update policies dynamically and on-the-fly, thus allowing for customization  
135 of the network data flows allowing for human control of any data flowing through  
136 the network. The Admin UI also displays suspiciousness scores for each device  
137 on the network, as well as overall network suspiciousness. The administrator can  
138 then use this information to determine if a device has cause for closer investigation  
139 to determine if an APT exists on the device or if the device has been compromised  
140 by other means.

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

## 147 **2. Related Work**

### 148 *2.1. Attack Defense using Trickery*

149 There have been efforts that seek to implement defense mechanisms using  
150 some form of ‘trickery’ to engage an attacker. For example, authors in [12] in-  
151 troduce the notion of tricking the attackers through IP randomization methods  
152 in decoy-based MTD efforts. In contrast, the notion of pretense in our Dolus  
153 approach is akin to Honeypots and Honeynets which are effective in gaining in-  
154 formation about possible attacks based on minimal active interactions with attack-  
155 ers [13]. Primarily they are used in a setting to either gain more information about  
156 potential attacks or the behavior of attackers.

157 Our work is complementary to Honeypots/Honeynets: we employ pretense  
158 to deceive attackers by rerouting and responding to attack traffic using quaran-  
159 tine VMs. Dolus system’s pretense theory is mainly built upon the work in [8]  
160 and [9] belonging to the field of child pretend play psychology. Our novel *defense*  
161 *by pretense* mechanism for effective mitigation of targeted attacks is inspired by  
162 the authors’ experiments where they show children (analogous to our attackers)  
163 various pictures of the animals along with a mismatch of the sounds made by the  
164 associated animals. Observations are made on how a pretense is effective based on  
165 how long it takes for a child to understand/protest that the information portrayed  
166 is actually false. In our case, the longer an attacker is tricked by our pretense, the  
167 more time a cloud service provider has to perform MTD mechanisms, strategize  
168 on patching identified vulnerabilities, as well as implement a SDxI-based infras-  
169 tructure policy coordination for mitigation of the impact of a targeted attack.

### 170 *2.2. Defense against DDoS Attacks*

171 Defense against flooding attacks such as DDoS typically involves attack traf-  
172 fic feature learning that provides intelligence on where the attack is coming from,  
173 and the specific attack type(s) [14] [15] [16]. Analysis of features such as source  
174 IP, destination IP, source port, destination port, size of packets, packet identifiers  
175 commonly help in subsequent filtering of flooding attacks. Authors in [17] show  
176 that the Internet traffic patterns are distinguishable, which can help filter and iso-  
177 late attack traffic flows. Once attack flows are filtered, blacklists are created [18],  
178 which can then be used to “scrub” the flows through scrubbing SaaS services as a  
179 low-cost solution [19].

180 A number of other network-based defense strategies have been proposed in  
181 efforts that involve analysis of traffic and dynamic updation of rules to effectively



182 reroute malicious traffic. Such efforts include [20], where a network reacts to tar-  
183 geted attacks using accountability and content-aware supervision concepts. Simi-  
184 larly, using volume counting, authors in [21] provide a DDoS defense mechanism  
185 that involves monitoring SDN traffic flows in OpenFlow-enabled switches. In  
186 the context of programmability of SDN switches to mitigate targeted attacks, au-  
187 thors in [22] present a programming framework. In another similar effort, authors  
188 in [23] propose a memory-efficient system that uses Bloom filter and monitoring  
189 tools to dynamically update SDN rules to mitigate DDoS attacks. Also leveraging  
190 the dynamic rule update feature of SDN, authors in [24] analyze the probability  
191 that a flow is traced back across multiple ASes' hops by sampling the probability  
192 and the analyzing signatures of attack traffic flows.

193 Alternately, cloud service providers allow mitigation of DDoS attacks by uti-  
194 lizing the elastic capacity provisioning capabilities in the cloud platforms that  
195 allow "moving target defense" (MTD) techniques to be implemented. MTD ba-  
196 sically involves replication and live migration [25] of compromised application  
197 services (with pre-attack state information) in new VM(s) to redirect legitimate  
198 users, and keep attackers in a quarantine VM(s) [6]. As an added defense strat-  
199 egy, authors in works such as [26] present a survey of SDN-based mechanisms to  
200 detect attacks closer to the attackers/attack sources.

### 201 2.3. Defense against APT Attacks

202 Techniques to detect APTs have been of interest to the community [27, 28,  
203 29, 29, 10, 30, 31, 32, 33, 34, 35, 36]. This includes finite angular state velocity  
204 machines and vector mathematics to model benign versus attack traffic, allowing  
205 a network operator to easily view the differences [33], assessing the outbound net-  
206 work traffic [10, 30], using honeypots [37] and using distributed computing [36].  
207 Another APT detection technique is based on a ranking system where all inter-  
208 nal hosts are ranked based on number of bytes sent outside the network, number  
209 of data transfers initiated to an entity outside the network, and number of dis-  
210 tinct destinations contacted outside the network per host [10]. Yet, another APT  
211 detection technique is to monitor attack traffic using a detector in an enterprise  
212 network [34].

213 Potential countermeasures against APTs are discussed in [4, 38, 39, 40, 41,  
214 42, 43, 44]. Defense strategies include: (a) running routine software updates to  
215 avoid backdoors, bugs and vulnerabilities; (b) strengthening network access con-  
216 trol and monitoring services; (c) enabling strict Internet access policies; and (d)  
217 dropping encrypted traffic from unknown hosts. Similarly, authors in [4] discuss  
218 a number of counter measures against APTs including training users about social

219 engineering attacks, blacklisting hosts, dropping packets, etc. Furthermore, SDN-  
220 based defense [42] involves: (i) defining and maintaining a network baseline to  
221 identify any deviation from the baseline through analytical tools, and (ii) updation  
222 of flow policies for (re)directing and blocking traffic in any of the network seg-  
223 ments. A framework that realizes such an SDN-based defense is discussed in [43].  
224 In a similar vein, authors in [44] provide a sandbox environment using which a  
225 security professional can emulate the propagation of APTs across an enterprise  
226 network environment.

### 227 **3. Dolus Defense Methodology**

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

#### 230 *3.1. Pretense Theory*

231 The pretense in the Dolus system is designed to create stimulus from the target  
232 side that matches the initial expectation of an attacker that a high-value target has  
233 not yet been compromised through an automated bot activity. Pretense theory  
234 concepts from [8] motivate us to address the issue of how a cognitive agent can  
235 present a pretense world, which is different from the real world using the following  
236 four steps:

- 237 (a) The basic assumption(s) or premise(s) that is used by a pretender on *what*  
238 is being pretended.
- 239 (b) Inferential elaboration which details of what goes into or what actually hap-  
240 pens in the process of pretense.
- 241 (c) Appropriate behavior production which answers the question of whether the  
242 pretender was successful on the audience being tricked.
- 243 (d) Balancing and steering the effects of pretense.

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

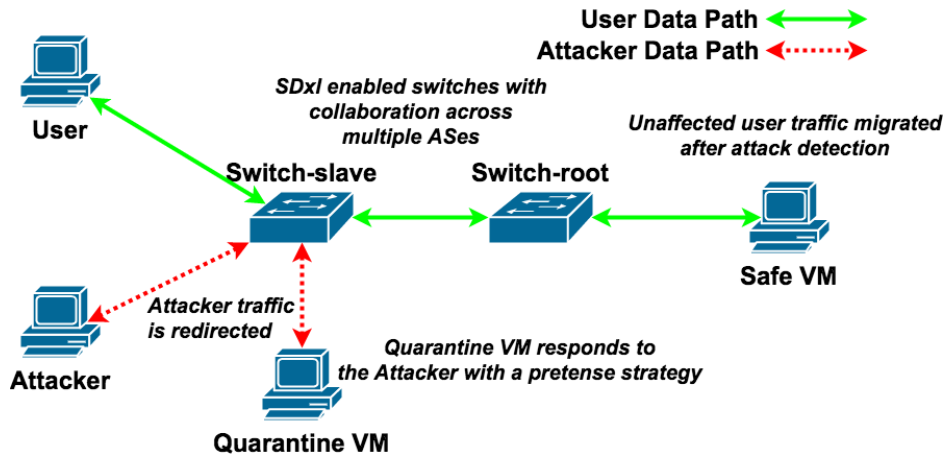


Figure 1: Illustration of the proposed Dolus system scheme wherein the attacker is *tricked* by redirection of the attack traffic to a quarantine VM for pretense initiation, while the providers work collaboratively to block the attack traffic closer to the source side.

252 3.2. Pretense in Dolus System Design

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

260 From the time gained through such a pretense initiation, Dolus enables cloud  
 261 service providers to decide on a variety of policies by dynamically generating net-  
 262 work policies using Frenetic [45] to mitigate the attack impact, without disrupting  
 263 the cloud services experience for legitimate users. In the worst case, destination-  
 264 side blocking can be enforced. Alternately, if the cloud service provider uses the  
 265 attack intelligence information and successful pretense time to coordinate the ‘hu-  
 266 mans in the loop’ of neighboring SDN-enabled domains, together they can direct  
 267 a unified SDN controller that directs SDN-enabled switches to actually enforce  
 268 attack traffic blocking measures closer to the attack source side.

269 The goal of our Dolus approach is to model the notion of pretense as a zero-  
 270 sum game. Specifically, a zero-sum game is one in which the sum of the individual  
 271 payoffs for each outcome is zero. That is, (1) loss to an attacker is gain for a de-  
 272 fender and vice versa, and (2) total sum of gain and loss is (roughly) zero. There

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

<b>Objective</b>	<b>Attacker</b>	<b>Defender</b>
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

273 are two strategies to play a zero-sum game, one from an attacker’s perspective  
 274 and the other from the defender’s perspective. Specifically, we plan to employ  
 275 the following two strategies: (a) *minimax* strategy: minimizing defenders own  
 276 maximum loss (from defenders perspective), and (b) *maximin* strategy: max-  
 277 imize attacker’s own minimum gain (from attacker’s perspective). We explore  
 278 these two strategies in our Dolus system on a number of objectives, which are  
 279 summarized in Table 1.

280 Our guiding strategy for targeted attack defense using pretense is to use a form  
 281 of *pretense machine learning* which we propose to be understood as - *If you don’t*  
 282 *know the enemy and don’t know yourself, then you will succumb in every battle.*  
 283 *If the attacker does not know you but you know the attacker, for some victories*  
 284 *gained you will suffer some defeat. If the enemy knows you and you know yourself,*  
 285 *you need not fear the result of a hundred battles.*<sup>2</sup>

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

---

<sup>2</sup>A quote modified from “The Art of War” by Sun Tzu.

## 293 4. DDoS Attack Defense with Dolus

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

### 297 4.1. Attack Model

298 DDoS attacks aim to overwhelm network-accessible devices such as networks,  
299 firewalls and end-systems in enterprises by sending packets at excessively high  
300 rates from multiple attack points. With cloud-hosted applications with large mon-  
301 etary value becoming highly common, DDoS attacks can cause LoA for users and  
302 customers, and can be used for extortion from vulnerable online businesses. Com-  
303 mon DDoS attack event types are amongst the 5 common attack vectors: DNS  
304 (Domain Name System), UDP (User Datagram Protocol) fragmentation, NTP  
305 (Network Time Protocol), SYN (short for synchronize), SSDP (simple service  
306 discovery protocol). For the purposes of our work, we assume the DDoS attacker  
307 uses SYN [46] and ICMP/Ping [47] flooding. Such attacks typically inundate a  
308 networks’ resources with Echo Request packets. We also assume that the attack-  
309 ers’ traffic is sent constantly and may or may not solicit a response in return. Such  
310 attacks can bring the network to a standstill due to the high volume of both incom-  
311 ing and outgoing traffic. To effectively capture the semantics of this attack model  
312 and to exhaust the target application services, we generate and emit synthetic ping  
313 and HTTP traffic using `hping3` [48] and `SlowHTTPTest` [49] tools, respec-  
314 tively. Furthermore, to capture the dynamics of an attacker, we randomly change  
315 the number of attack packets emitted by these tools.

### 316 4.2. Defense by Pretense Scheme

317 .  
318 Figure 2 depicts the cross-domain setup in a Dolus system deployment to im-  
319 plement a defense by pretense scheme. To complement Figure 2, interactions  
320 between different phases of a Dolus system configured for spoofing pretense are  
321 shown in Figure 3 and Algorithm 1, respectively.

322 **Attack Detection.** First, traffic within a cloud provider’s network (which is gen-  
323 erated by the SDN switches) or across multiple transit provider ASes (which are  
324 composed of SDX plus SDN switch substrates) is monitored using a Frenetic run-  
325 time [45]-enabled monitoring subcomponent (line 24 of Algorithm 1). Next, in  
326 order to learn and classify the attacks (line 25 of Algorithm 1), we employ a two-  
327 stage ensemble learning scheme on the incoming traffic, both from the attackers

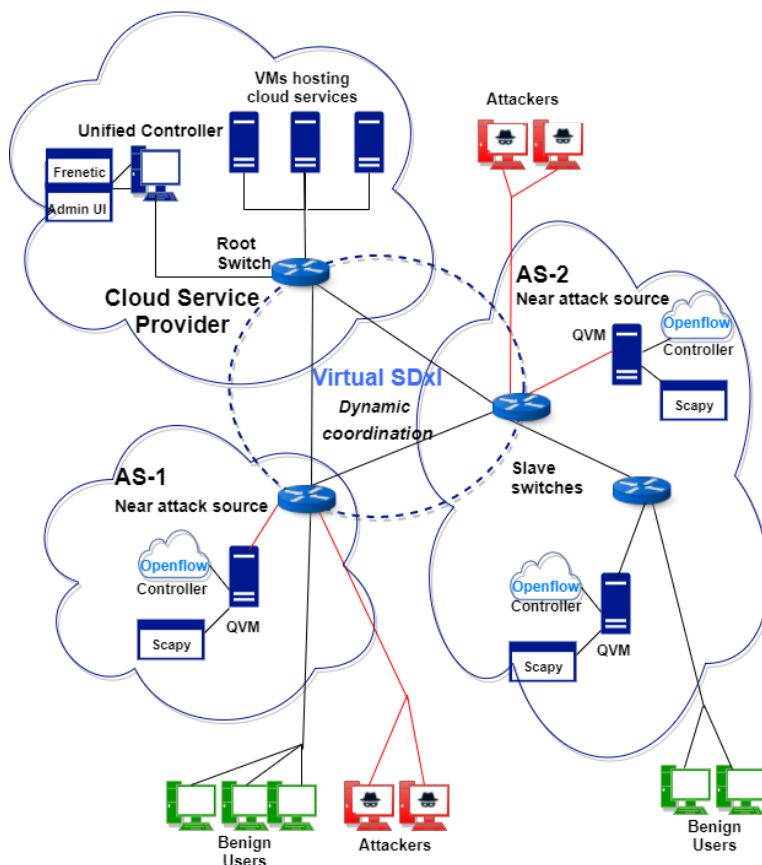


Figure 2: Cross-domain physical setup in a Dolus system deployment to share threat intelligence for a unified controller to coordinate policy management with a federation of ASes to block attack traffic closer to the source side.

328 and from the benign users. In order to differentiate attackers from benign users,  
 329 the first stage handles outlier detection to identify salient events of interest (e.g.,  
 330 connection exhaustion), whereas the second stage handles outlier classification to  
 331 distinguish different event types (e.g., DDoS attack).

332 *Outlier Detection.* We use basic/static methods such as multivariate Gaussian to  
 333 detect outliers and build upon our prior work on detecting network-wide correlated  
 334 anomaly events [50, 51] that are typical of the traffic from multiple attack sources.  
 335 Specifically, the outlier detection is a composition of many efficient, multivariate  
 336 outlier detectors or hypotheses functions:  $\mathcal{H} = \{h_1, h_2, \dots, h_n\}$  and the result,  
 337  $\mathcal{F}$ , is an ensemble of the different hypotheses. Furthermore, we note that the

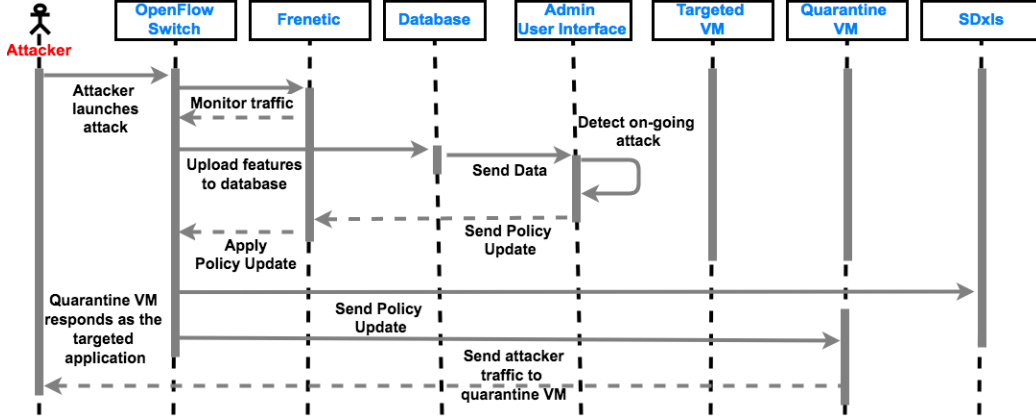


Figure 3: Sequence diagram of the Dolus system interactions for attack detection, quarantine setup, pretense initiation/maintenance and DDoS attack impact mitigation.

338 traditional methods for ensemble learning use averaging or majority voting [52].  
 339 In our case, to achieve higher accuracy with a minimum size of the training dataset  
 340  $D$ , we use the Bayesian voting scheme [53] as the ensemble method to predict the  
 341 result for new data  $x$ , which can be represented as Equation 1.

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

342 Final ensemble result  $\mathcal{F}$  consists of all of the hypotheses in  $\mathcal{H}$ , and each hypo-  
 343 thesis  $h$  weighted by its posterior probability  $P(h|D)$ . The posterior probability  
 344 is proportional to the likelihood of the training data  $D$  times the prior probability  
 345 of  $h$  (2).

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

346 *Outlier Classification.* The outliers detected are classified into either interesting  
 347 events (e.g., attacks) or erroneous conditions (e.g., router failure). We use a simple  
 348 classifier to this end: if the final ensemble results of consecutive events (detected  
 349 in the first stage) fall in the same range, we classify them as an attack; otherwise,  
 350 we ignore those events. We remark that the above two-stage ensemble learning  
 351 scheme requires a sizable amount of data to classify the attacks effectively. To  
 352 overcome this challenge, we initially let the attacker(s) to attack the cloud ser-  
 353 vices. However, we also monitor the incoming traffic carefully and make sure that  
 354 the attack does not disrupt the network resources. Once an attack is classified,

355 which are shown separately in Figure 2, we reroute the attack traffic using Fre-  
356 netic runtime to quarantine VM (QVM) along with sample server responses (see  
357 lines 1 through 22 of Algorithm 1).

358 **Quarantine Setup for Pretense.** Dolus calls the quarantine setup procedure  
359 (lines 1 to 9) where a new QVM is instantiated using a cloud platform’s elas-  
360 tic provisioning capability and the update policy routine is invoked (line 3). In  
361 the update policy routine (lines 10 to 14), we log the attack traffic to prevent fu-  
362 ture attack events as well as invoke the Frenetic runtime to generate new policies  
363 (line 12). Frenetic executes Python scripts to identify suspicious packets, learn  
364 from patterns and directs switches to redirect packets to QVMs. We then adver-  
365 tise this information (attack intelligence) to the neighboring switches (line 13),  
366 where, apart from the policy updates, the IP addresses of the attackers are black-  
367 listed. Following this, based on the stored attack traffic logs, the QVM uses Scapy  
368 libraries [54] to generate responses with spoofed IP addresses and pretends as the  
369 targeted VM under attack from the perspective of the attacker(s) (lines 20 to 22).

370 Subsequently, depending on the nature and volume of the incoming data, we  
371 decide either to move forward with the pretense or drop the traffic—which is the  
372 third step of production of appropriate behavior in pretense theory (lines 28 to  
373 30). In order to gain more information about the attackers/attacks, we typically  
374 continue the process of pretense. While we continue the pretense, we routinely  
375 update threat intelligence such as the attacker’s IP, targeted VM’s IP where ser-  
376 vice(s) under attack is hosted, type of attacks, etc. Furthermore, we assume that  
377 an attacker has enough knowledge on how a successful attack should affect our  
378 system, which is another reason why we keep the attacker involved in the system  
379 as long as is usually expected. If we drop the attack traffic too early or maintain it  
380 for too long, attacker might potentially infer our pretense.

381 Finally, we redirect the flow of the attack traffic by pushing a new policy from  
382 the unified controller running in the cloud to the switch(es). This will redirect  
383 the attacker’s traffic that is intended for the targeted VM towards the QVM. The  
384 QVM then responds to the attacker’s traffic as though it is the targeted VM/server  
385 under attack with spoofed IP address and hostname of the target, which creates  
386 the pretense effect, from an attacker’s perspective, that the targeted DDoS attack  
387 is successful. Depending on the nature of the attack, we want the attacker to  
388 believe that services are no longer up/available on the targeted VM. We therefore  
389 allow the QVM to continue to respond to the attacker for a limited amount of time  
390  $t$ . We tune  $t$  based on the type of attack traffic and how the targeted VM would  
391 respond if it was under attack. For example, if the targeted VM went down after  
392 10 seconds of attack, the QVM would do the same by not responding at the same



393 time with a variable random delay factor of  $[-1,1]$  seconds added. This allows the  
394 attacker to see that the services are available until, suddenly, they no longer are.  
395 **Policy Decision Making.** In this sense, our defense maintains the pretense: gives  
396 the attacker the confirmation of a successful attack, when in reality the service has  
397 not been affected at all considering the scenario that the user is running a video  
398 gaming portal application. This also gives us sufficient time to collect information  
399 about the attackers and their attack patterns. We use the collected information to  
400 create a blacklist of attacker information. To help network administrators effec-  
401 tively manage the network in the face of attacks, our system also consists of a  
402 Administrator User Interface (Admin UI) module and a unified controller module  
403 that can be customized in a Dolus system instance deployment depicted in Fig-  
404 ure 2. The Admin UI shown in Figure 4 can be used for e.g., to enforce users  
405 to adhere to the policies generated by Frenetic runtime when they connect to the  
406 cloud. Policies generated by Frenetic internally are updated through the User  
407 Interface using JSON arrays. These policies (e.g., open/block flows) could be in-  
408 stalled in the switches using the unified controller module, which is also linked  
409 with a back-end database that logs traffic characteristics and user profiles.

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

423 **Threat Intelligence Sharing.** Algorithm 1 runs in the monitor component and  
424 coordinates/shares intelligence with the switches deployed in the network and  
425 across different providers. This in turn enables a collaborative environment among  
426 providers such that the targeted attacks can be detected closer to the source *with-*  
427 *out* affecting the cloud infrastructure. A natural question is why would a provider  
428 share the attack intelligence, especially in a business that is driven by competi-  
429 tion? We posit that the coordination among different ASes/providers is mutually  
430 beneficial for all the entities involved. Of course, a particular AS/provider can de-

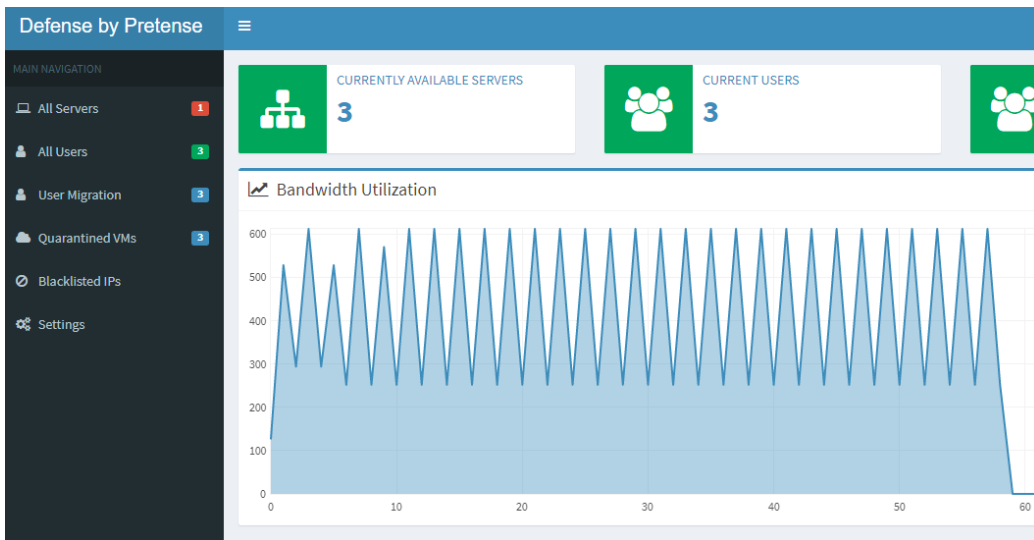


Figure 4: Administrator User Interface of an Dolus system instance.

431 cide not to share the attack intelligence to others. However, if an AS experiences  
 432 an attack and if it shares the intelligence with other ASes, a global and unified  
 433 hardening of infrastructure against such targeted attacks can be achieved. In addi-  
 434 tion, any downtime is money lost in a business; sharing the attack intelligence in  
 435 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
```

## 436 5. APT Attack Defense with Dolus

### 437 5.1. Attack Model

438 APTs are long-term attacks and affect a target in four stages: *preparation, ac-*  
439 *cess, resident, and harvest* [7, 55, 56, 57, 58, 59, 38]. In the preparation stage,  
440 attackers apply a reconnaissance tactic through social engineering (e.g., via so-  
441 cial networks) to bootstrap the attack [4]. Once the attack is bootstrapped, at-  
442 tackers identify a vulnerability, and/or a vulnerable target and send malwares ei-  
443 ther through email (e.g., spear phishing) or through third-party software/service  
444 (e.g., watering-hole attack) in the access stage. Subsequently, the malwares estab-  
445 lish external communication paths with attackers' Command and Control (C&C)  
446 server(s), and spread across other targets in the resident stage; which is a slow and  
447 a stealthy phenomenon. Finally, in the harvest stage, attackers extract any vital  
448 information in an on-going fashion for extended periods of time.

### 449 5.2. Defense by Pretense Scheme

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

463 **Attack Detection.** Inspired by the work of authors in [10] to identify hosts ex-  
464 hibiting suspiciousness in a network, we propose a Suspiciousness Score ( $SS$ ) in  
465 a similar vein. We calculate  $SS$  based on captured network traces (.pcap) using  
466 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

467 Table 2 shows the list of values/features captured in network traces for APT  
468 attack defense analytics. For each packet trace, a trace\_id  $t$  is assigned. For each  
469  $t$ , we perform the following: the features are normalized and their combined Root  
470 Mean Square Error (RMSE) values are calculated. Using the RMSE values, we  
471 calculate the Suspiciousness Scores of each device as follows. The *Min* and *Max*  
472 values (below) are assumptions made per device type on what one may expect the  
473 minimum and maximum values to be on the type of device, network and traffic  
474 expectations. These values are determined by the system/network administrators  
475 and could vary vastly depending on each ASeS or domain's threat monitoring  
476 objectives.

477 Destination suspiciousness for trace  $t$ :

$$dst_i = \frac{numDst_i - numDistMin_i}{numDstMax_i - numDstMin_i} \quad (3)$$

478 Flow suspiciousness for trace  $t$ :

$$flows_i = \frac{numFlows_i - numFlowsMin_i}{numFlowsMax_i - numFlowsMin_i} \quad (4)$$

479 Bytes suspiciousness for trace  $t$ :

$$bytes_i = \frac{numBytes_i - numBytesMin_i}{numBytesMax_i - numBytesMin_i} \quad (5)$$

480 Device suspiciousness for trace  $t$  is based on equations 3, 4 and 5 as shown  
 481 below.

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

482 Note that for each device on the network  $i$  we calculate a Suspiciousness Score  
 483 and the overall network suspiciousness for trace  $t$  is calculated based on  $ss$  for  
 484 each individual device (equation 6) that is connected. That is, the sum of all  $ss$   
 485 for each devices on the network  $n$  is the *overall network suspiciousness*  $SS$  for that  
 486 particular  $t$ .

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

487 Relative change in device  $i$ 's suspiciousness score on new traffic  $t$  is simply  
 488 given by -

$$\Delta ss_{it} = \frac{ss_{it} - \sqrt{\frac{ss_{i1}^2 + ss_{i2}^2 + \dots + ss_{it-1}^2}{t-1}}}{\sqrt{\frac{ss_{i1}^2 + ss_{i2}^2 + \dots + ss_{it-1}^2}{t-1}}} \quad (8)$$

489 **Quarantine Setup for Pretense.** VMs which are internal to an organization and  
 490 which implement minimal versions of honeypot-like hosts are internal quarantine  
 491 VMs (iQVMs), whose Suspiciousness Scores are monitored continuously. These  
 492 are also the hosts that play the game of pretense i.e., they create a false notion of  
 493 *high-value targets within an organization with sensitive data* to the external world.  
 494 An attacker is lured to attack iQVMs first; they maintain pretense by sending data  
 495 similar to what a host with sensitive data would send. Apart from monitoring the  
 496 data sent out of iQVMs, they also add weights to the calculated Suspiciousness  
 497 Scores, overcoming the threshold evasion problem.

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

- 505 1. SS view: `Flot.js`-based bar and line graphs as depicted in Figure 5 ex-  
 506 cerpted from the Admin UI. SS per device or for the overall network can be

507 viewed in temporal fashion as shown in Figure 6 extracted from the Admin  
508 UI. Moreover, when a suspiciousness score of a blacklisted device is shown  
509 to be above a certain threshold, an administrator can block all traffic from  
510 that device to the network or take an appropriate action.

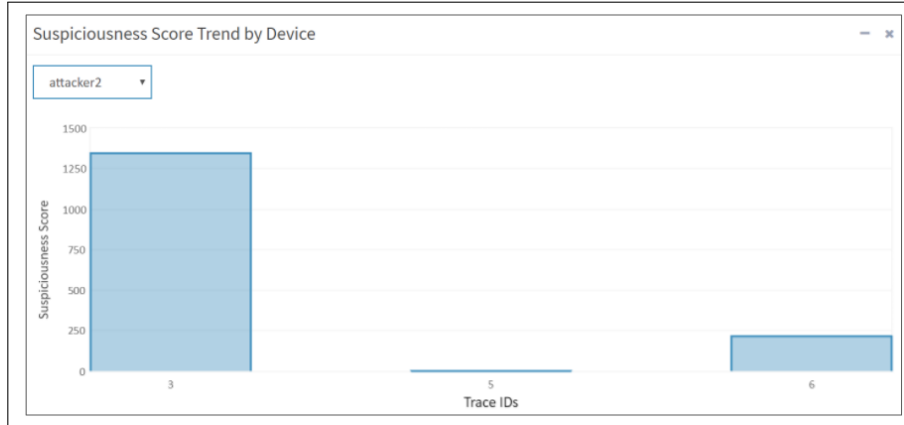


Figure 5: Suspiciousness score per device over time.




Figure 6: Overall network suspiciousness score over time.

511 2. Upload policy view: This view on the user interface enables administra-  
512 tors to push NetKAT-based policies [60] to a centralized database, which  
513 stores device configurations, thresholds, policies maintained by the organi-  
514 zation. Interfaces are provided to select a specific device and a correspond-  
515 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.

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

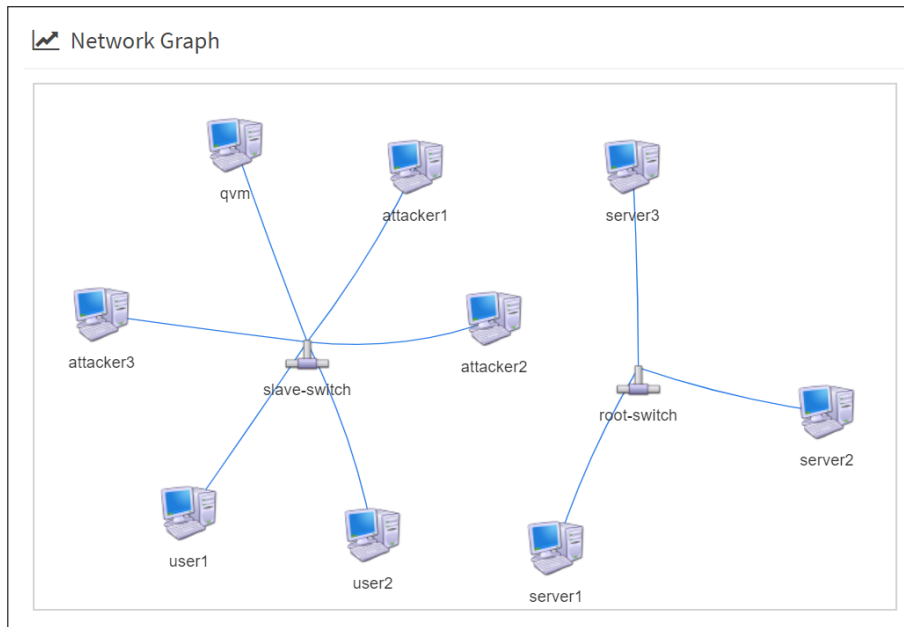


Figure 8: Network graph of all the connected devices.

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



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

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

528 Both these policies and the revocation/enabling functionalities are instantiated  
529 based on the policy updater mechanism, whose main objective is to simplify the  
530 learning curve for users/administrators to get proficient at writing policies (e.g.,  
531 using network programming languages such as Frenetic [45])—a daunting and te-  
532 dious task. With this in mind, the updater component can auto-generate policies  
533 based on simplified inputs that are provided via the user interface. For example, to  
534 minimize the process, the policy updater takes a generic command such as “user1  
535 to server1” and all possible configuration policies would be generated by the up-  
536 dater. The updater works with the centralized database and is pre-programmed  
537 with the network architecture.

538 **Threat Intelligence Sharing.** Our iQVM monitors also coordinate and share the  
539 APT threat intelligence such as *SS* thresholds, policy updates, etc. with other  
540 hosts in the network. Apart from providing a collaborative environment amongst  
541 pertinent hosts to effectively counter APTs, the mechanism also provides a way  
542 to drill down on specific segments of the network with suspicious hosts. Further-  
543 more, we believe that the coordination mechanism will pave the way to achieve  
544 a global and unified hardening of the enterprise network against APTs. In addi-  
545 tion, any sensitive data sent out is money lost in a business; sharing the threat  
546 intelligence in turn provides a cheaper alternative to lost data and host/business  
547 downtimes.

## 548 6. Performance Evaluation

549 In this section, we describe the evaluation of our Dolus methodology in GENI  
550 Cloud testbeds for DDoS and APT attacks. For showing effectiveness of Dolus  
551 for each targeted attack type, we start by describing our testbed, followed by the  
552 experiments and results discussion. The source code and instructions to replicate  
553 below experiments are openly available at [61] [62].

554 6.1. Dolus Experiments for DDoS Attack Defense

555 6.1.1. Testbed Setup

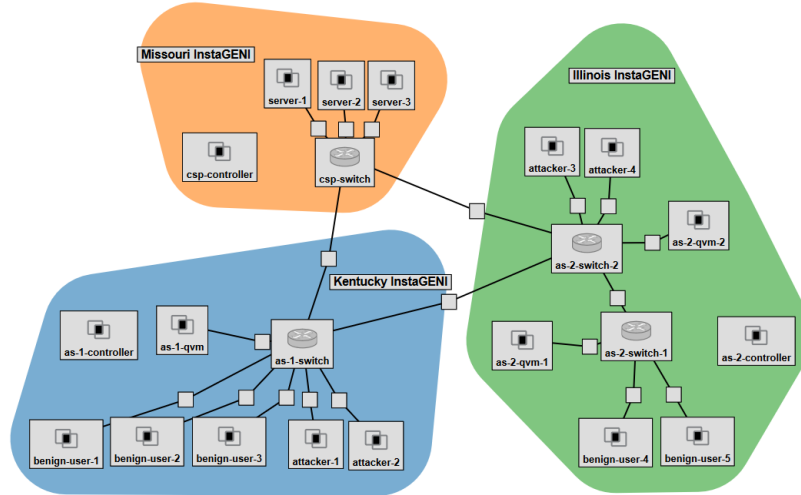


Figure 9: GENI Cloud testbed used to evaluate Dolus for DDoS attack defense.

556 We evaluate the efficacy of our Dolus system using a realistic, GENI Cloud [11]  
557 testbed as shown in Figure 9. The testbed contains three SDN switches, two slave  
558 switches and a single root switch. Such a system could also be extended to host  
559 many more switches and devices. . The slave switches are each attached to users  
560 and attackers, a quarantine VM, and a connection to the root switch. Likewise, the  
561 root switch is connected to elastic VMs, each of which could serve as a candidate  
562 for the target application (i.e., a video gaming portal) hosting that could be com-  
563 promised by the attackers. All switches are connected to a unified SDN controller  
564 located in the cloud service provider domain, which directs the policy updates.  
565 In the following, we show the attack detection and classification accuracy using  
566 our two-stage ensemble learning scheme and then present results from two sets of  
567 experiments that were run for a maximum of 28 *seconds* to show how our Dolus  
568 implementation can be used in real-time to restore cloud services under DDoS  
569 attack situations.

570 6.1.2. Attack Detection and Classification Results

571 Using the Dolus system, we monitor different types of data that are permitted  
572 to enter the GENI Cloud testbed depicted in Figure 9. We send both normal and

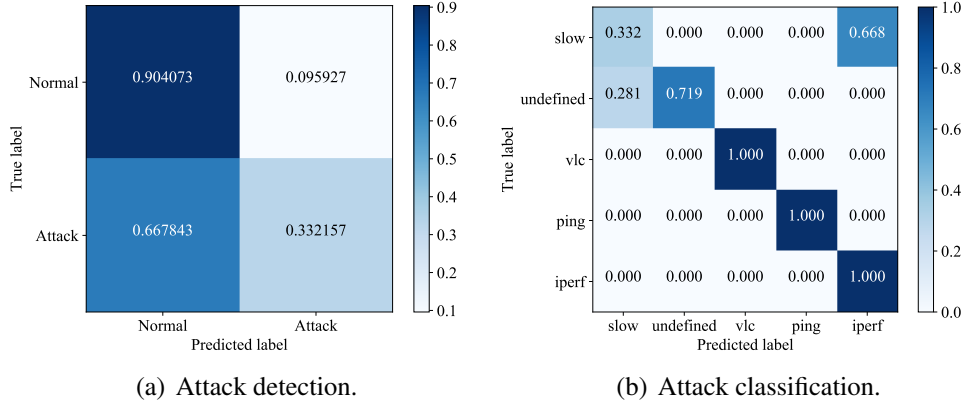


Figure 10: Confusion matrices for attack detection and classification for multiple traffic flows sent to multiple hosts.


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

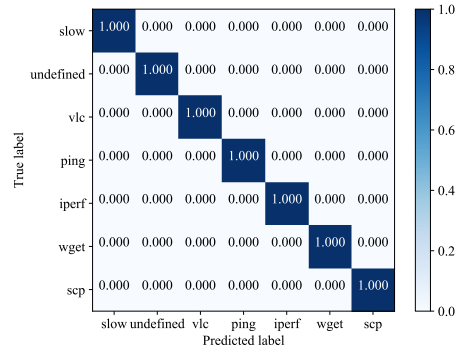
573 attack traffic (i.e., our datasets) to the targeted server to test the efficacy of our two-  
574 stage ensemble learning scheme. Our evaluation results span over two instances  
575 of learning of datasets as explained in the following.

576 The first instance shows multiple traffic types from a single attacker VM to  
577 a single target node. For this instance, we divide  $\sim 180,000$  lines of data into  
578 two sets, one for training and the other to test the accuracy of our scheme. Fur-  
579 thermore, the types of traffic used to create these instances are composed of  
580 SlowHTTPTest, iperf, VLC and ICMP ping. Figure 15 shows the two confu-  
581 sion matrices for attack detection and classification in a normalized fashion. We  
582 note that both the detection and the classification of attack took less than a sec-  
583 ond. In addition to the rapid detection and classification, our approach is highly  
584 accurate as shown in Table 3, where stage 1 is the detection stage and stage 2  
585 is the classification stage.

586 In the second instance, we consider multiple traffic types to multiple hosts.  
587 This instance is composed of 2.5 million rows per test, totaling 5 million rows of



(a) Attack detection.



(b) Attack classification.

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

588 data. The types of traffic that we use to create this dataset include SlowHTTPTest,  
 589 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  
 590 evaluate the effectiveness of our two-stage ensemble learning scheme. Figure 10  
 591 shows the two confusion matrices in normalized form for attack detection and  
 592 classification. Detection and classification of attack took  $\sim 7$  and  $\sim 13$  seconds,  
 593 respectively. Despite the slowdown in attack detection/classification in comparison  
 594 with the first instance, the accuracy of our approach is still high as shown in  
 595 Table 3.  
 596

597 While the two-stage ensemble learning scheme is effective in detecting test  
 598 data, a new attack that has not been used in training could initially go undetected  
 599 and impact services. However, with pertinent labeling of attack traffic flows during  
 600 training, the accuracy of the ensemble learning scheme can be improved significantly.  
 601 We depict the outlier detection and classification for a trained case in Figure 11,  
 602 where we make use of 60% of the data as training data and 40%  
 603 as test data for the same dataset used in the 2<sup>nd</sup> instance. For the purpose of our  
 604 evaluation, the sorted dataset has randomized time stamps.

605 Though the dataset that we use is discrete with differences in traffic such as  
 606 protocol, bytes transmitted, number of packets, source and destination addresses,  
 607 our two-stage ensemble learning scheme is effective in detecting the attacks with  
 608 good accuracy and efficiency. The ensemble learning scheme can further be modified  
 609 based on other characteristics of network traffic, and such modifications are  
 610 beyond the scope of the work in this paper.




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.

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

Figure 12 compares the time taken by our Dolus system to stop a DDoS attack versus MTD-based and no defense strategies. After a warm-up period of 6 seconds, we start the SlowHTTPTest and hping3 at the 7<sup>th</sup> second from the attackers. In a SDxI-based cloud network with no defense strategy, the services are immediately affected by the attack traffic. Similarly, the MTD-based defense strategy takes  $\sim 6$  seconds to mitigate the attack traffic impact. However, our Do-




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

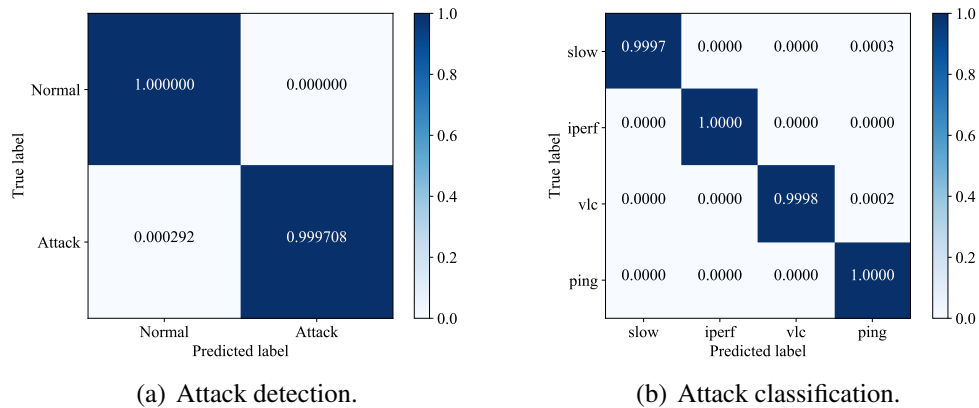


Figure 15: Confusion matrices for attack detection and classification for multiple traffic flows sent to a single server.

618 lus system supported service on the other hand, does not suffer from any loss of  
 619 availability in comparison with the other two strategies. This is due to the sharing  
 620 of attack intelligence between the slave switches and redirection of attack traffic  
 621 to quarantine VMs closer to the attackers, making the cloud network completely  
 622 oblivious to the attackers.

#### 623 6.1.4. Amount of Traffic Processed at the Root Switch

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

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

#### 638 6.2. Dolus Experiments for APT Attack Defense

##### 639 6.2.1. Testbed Setup

640 For the purposes of APT attack detection and defense, a modified GENI Cloud  
641 testbed was setup as shown in Figure 16. Since an APT is not a distributed attack,  
642 there was no need to consider multiple attack vectors from many directions. How-  
643 ever, due to the nature of an APT attack being secretive and stealthy, we assume  
644 that an APT can be hiding anywhere in an enterprise network in our setup con-  
645 figurations. Our testbed is comprised of multiple open vSwitches (a slaves and  
646 a single root), numerous nodes (which are hosts), and a controller. The slave  
647 switches connect all the user nodes, and the root switch connects all the servers  
648 hosting the application system and related services to the slave switches. The con-  
649 troller in the setup is a standalone node, running the monitor and policy updaters,  
650 calculating  $SS$  thresholds for nodes and the overall network, managing all the  
651 traffic and defense by pretense mechanisms of the Dolus system.

##### 652 6.2.2. Suspiciousness Score Calculation Results

653 In the first experiment, we randomly selected three hosts, and compromised  
654 them by running `slowhttp` attacks from attacker 1 and attacker 3, and a secure  
655 copy (`scp`) from attacker 2. This configuration allows us to compare the suspi-  
656 ciousness between a DDoS attack, and a file exfiltration attack. Before running  
657 the experiment, we specify minimum and maximum values for flows, connections,

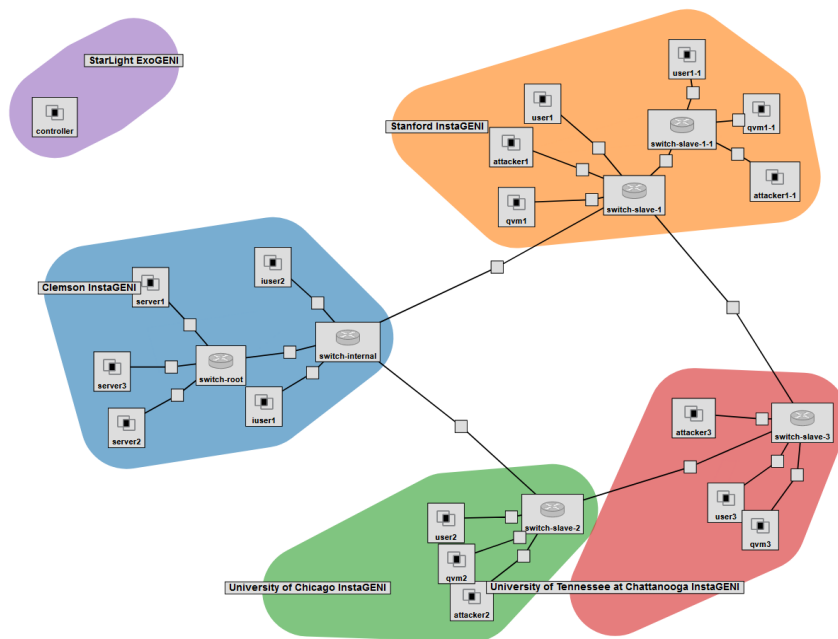


Figure 16: GENI Cloud testbed used to evaluate Dolus for APT attack defense.

658 and bytes: the user and attacker nodes are each set to a minimum of 1 and a max-  
 659 imum of 10 connections, a minimum of 100 and a maximum of 1,000 flows, and  
 660 a minimum of 10 and a maximum of 100,000 bytes. The servers had a minimum  
 661 of 10 and a maximum of 1000 connections, a minimum of 1,000 and maximum  
 662 of 10,000 flows, and minimum of 100,000 and a maximum of 100000000 bytes.

663 From the controller, we obtain the  $SS$  for these three attackers before (see  
 664 Table 4) and after (see Table 5) whitelisting. Note that all devices have  $SS$  calcu-  
 665 lated for them, as we don't initially whitelist any devices or traffic on our testbed  
 666 network. Attacker 2 exhibited the highest  $SS$  out of three, due to data exfiltration  
 667 [10]. The traffic that is being exfiltrated generates a much higher score than  
 668 the regular traffic in the network.

669 The purpose of whitelisting is to allow administrators to ignore traffic, which  
 670 is not going outside of the network. For example, if we consider that both server  
 671 1 and user 1 are within our own network, then any data transmitted between those  
 672 two machines would not be data being exfiltrated from the enterprise network.  
 673 Therefore, we can consider such traffic as benign. However, whenever we con-  
 674 sider attacker 1 and server 1, since attacker 1 is compromised, we consider all  
 675 traffic from attacker 1 to be possible data exfiltrated from the enterprise network.



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

Table 5: Suspiciousness Scores after Whitelisting

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

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

685 **6.2.3. Time Overhead for Suspiciousness Score Calculation**

686 Table 6 shows the time taken by ADAPTs to calculate the *ss* for devices, each  
687 running on a single core. It also shows the number of traces, and their correspond-  
688 ing processing times. As high as 1.8 Million traces for 8 devices can be processed  
689 under 100s, which demonstrates the efficacy of ADAPTs. However, there is a lin-  
690 ear increase in time as the number of traces grow. If such a linear increase does  
691 not meet the threat monitoring objectives of a domain, a parallel implementation  
692 of ADAPTs can be extended and used on nodes with multi-core settings to reduce  
693 the computation times in the Suspiciousness Score calculations.

Table 6: Processing time taken by ADAPT on single core nodes.

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

694 **7. Conclusion**

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

708 Our above approach for DDoS and APT attacks defense takes advantage of  
709 elastic capacity provisioning in cloud platforms to implement moving target de-  
710 fense techniques that does not affect the cloud-hosted application users, and con-  
711 tains the attack traffic in a quarantine VM(s). With the time gained through effec-  
712 tive pretense initiation in the case of DDoS attacks, cloud service providers could  
713 coordinate across a unified SDxI infrastructure involving multiple ASes to decide  
714 on policies that help in blocking the attack flows closer to the source side. Perfor-  
715 mance evaluation results of our Dolus system in a GENI cloud testbed for DDoS  
716 attacks show that our approach can be effective in filtering, detection and imple-  
717 mentation of SDxI-based infrastructure policy coordination for mitigation of the  
718 impact of the DDoS attacks. In addition, we also showed how the Dolus system  
719 can be an effective defense using pretense against APTs. Using the concept of  
720 Suspiciousness Scores, we proposed novel threat intelligence collection and ana-  
721 lytics of subtle and secretive attacks at a device level and also at a network-wide  
722 level. Further, we found that our Admin UI capability can greatly help network  
723 operators and cloud service providers to overcome their difficulty in determin-  
724 ing which devices on an enterprise network or a cloud service deployment may  
725 be compromised and how effective a pertinent defense strategy is functioning to  
726 mitigate the impact of targeted attacks.

727 Future work can be pursued to investigate more sophisticated pretense schemes  
728 that use threat intelligence collection on effectiveness of a working pretense, and  
729 initiate more involved adaptations. In addition, data analytics extensions can be  
730 pursued for more sophisticated targeted attacks with significantly larger number  
731 of features that need to be involved in effective detection and defense schemes.

## 732 **Acknowledgement**

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

## 737 **References**

- 738 [1] Verizon. State of the Market: Enterprise Cloud,  
739 [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)  
740 [enterprise-cloud-2016\\_en\\_xg.pdf](http://www.verizonenterprise.com/resources/reports/rp_state-of-the-market-enterprise-cloud-2016_en_xg.pdf), 2017.

- 741 [2] SDxCentral, SDxCentral. Software Defined Everything Part 3: SDx  
742 infrastructure, [https://www.sdxcentral.com/cloud/definitions/software-  
defined-everything-part-3-sdx-infrastructure](https://www.sdxcentral.com/cloud/definitions/software-<br/>743 defined-everything-part-3-sdx-infrastructure), 2017.
- 744 [3] A. Gupta, N. Feamster, L. Vanbever, Authorizing network control at soft-  
745 ware defined internet exchange points, in: Proceedings of the Symposium  
746 on SDN Research, SOSR '16, ACM, New York, NY, USA, 2016, pp. 16:1–  
747 16:6.
- 748 [4] P. Chen, L. Desmet, C. Huygens, A study on advanced persistent threats, in:  
749 IFIP International Conference on Communications and Multimedia Security,  
750 Springer, pp. 63–72.
- 751 [5] P. Kampanakis, H. G. Perros, T. Beyene, Sdn-based solutions for moving  
752 target defense network protection., in: WoWMoM, IEEE Computer Society,  
753 2014, pp. 1–6.
- 754 [6] S. Debroy, P. Calyam, M. Nguyen, A. Stage, V. Georgiev, Frequency-  
755 minimal moving target defense using software-defined networking, 2016  
756 International Conference on Computing, Networking and Communications  
757 (ICNC) 00 (2016) 1–6.
- 758 [7] N. Virvilis, D. Gritzalis, T. Apostolopoulos, Trusted computing vs. advanced  
759 persistent threats: Can a defender win this game?, in: Ubiquitous intel-  
760 ligence and computing, 2013 IEEE 10th international conference on and  
761 10th international conference on autonomic and trusted computing (uic/atc),  
762 IEEE, pp. 396–403.
- 763 [8] S. Nichols, S. Stich, A cognitive theory of pretense, *Cognition* 74 (2000)  
764 115–147.
- 765 [9] J. W. V. de Vondervoort, O. Friedman, Young children protest and correct  
766 pretense that contradicts their general knowledge, *Cognitive Development*  
767 43 (2017) 182–189.
- 768 [10] M. Marchetti, F. Pierazzi, M. Colajanni, A. Guido, Analysis of high vol-  
769 umes of network traffic for advanced persistent threat detection, *Computer*  
770 *Networks* 109 (2016) 127–141.

- 771 [11] M. Berman, J. S. Chase, L. Landweber, A. Nakao, M. Ott, D. Raychaudhuri, R. Ricci, I. Seskar, Geni: A federated testbed for innovative network  
772 experiments, *Comput. Netw.* 61 (2014) 5–23.  
773
- 774 [12] A. Clark, K. Sun, R. Poovendran, Effectiveness of IP address randomization  
775 in decoy-based moving target defense, in: *Proceedings of the 52nd IEEE  
776 Conference on Decision and Control, CDC 2013, December 10-13, 2013,  
777 Firenze, Italy*, pp. 678–685.
- 778 [13] T. Sochor, M. Zuzcak, Study of internet threats and attack methods using  
779 honeypots and honeynets, in: *International Conference on Computer Net-  
780 works*, Springer, pp. 118–127.
- 781 [14] S. Seufert, D. O’Brien, Machine learning for automatic defence against  
782 distributed denial of service attacks, in: *2007 IEEE International Conference  
783 on Communications*, pp. 1217–1222.
- 784 [15] J. Choi, C. Choi, B. Ko, P. Kim, A method of ddos attack detection using  
785 http packet pattern and rule engine in cloud computing environment, *Soft  
786 Computing* 18 (2014) 1697–1703.
- 787 [16] T. Thapngam, S. Yu, W. Zhou, S. K. Makki, Distributed denial of service  
788 (ddos) detection by traffic pattern analysis, *Peer-to-Peer Networking and  
789 Applications* 7 (2014) 346–358.
- 790 [17] Y. Chen, S. Jain, V. K. Adhikari, Z. L. Zhang, K. Xu, A first look at inter-data  
791 center traffic characteristics via Yahoo! datasets, pp. 1620–1628.
- 792 [18] L. Yang, T. Zhang, J. Song, J. S. Wang, P. Chen, Defense of ddos attack  
793 for cloud computing, in: *2012 IEEE International Conference on Computer  
794 Science and Automation Engineering (CSAE)*, volume 2, pp. 626–629.
- 795 [19] Internet2’s ddos mitigation strategy, <https://www.internet2.edu/blogs/detail/12234>,  
796 2016.
- 797 [20] Y. Choi, Implementation of content-oriented networking architecture (cona):  
798 a focus on ddos countermeasure.
- 799 [21] C. YuHunag, T. MinChi, C. YaoTing, C. YuChieh, C. YanRen, A novel  
800 design for future on-demand service and security, in: *2010 IEEE 12th Inter-  
801 national Conference on Communication Technology*, pp. 385–388.

- 802 [22] S. Shin, P. A. Porras, V. Yegneswaran, M. W. Fong, G. Gu, M. Tyson, Fresco:  
803 Modular composable security services for software-defined networks., in:  
804 NDSS, The Internet Society, 2013.
- 805 [23] Y. Yu, C. Qian, X. Li, Distributed and collaborative traffic monitoring in  
806 software defined networks, in: Proceedings of the Third Workshop on Hot  
807 Topics in Software Defined Networking, HotSDN '14, ACM, New York,  
808 NY, USA, 2014, pp. 85–90.
- 809 [24] H. Tian, J. Bi, An incrementally deployable flow-based scheme for ip trace-  
810 back., IEEE Communications Letters 16 (2012) 1140–1143.
- 811 [25] C. Clark, K. Fraser, S. Hand, J. G. Hansen, E. Jul, C. Limpach, I. Pratt,  
812 A. Warfield, Live migration of virtual machines, in: Proceedings of the 2nd  
813 Conference on Symposium on Networked Systems Design & Implementa-  
814 tion - Volume 2, NSDI'05, USENIX Association, Berkeley, CA, USA, 2005,  
815 pp. 273–286.
- 816 [26] Q. Yan, F. R. Yu, Q. Gong, J. Li, Software-defined networking (sdn) and dis-  
817 tributed denial of service (ddos) attacks in cloud computing environments:  
818 A survey, some research issues, and challenges, IEEE Communications Sur-  
819 veys Tutorials 18 (2016) 602–622.
- 820 [27] J. Kim, T. Lee, H.-g. Kim, H. Park, Detection of advanced persistent threat  
821 by analyzing the big data log, Cloud Security Alliance 13 (2013) 101–107.
- 822 [28] P. K. Sharma, S. Y. Moon, D. Moon, J. H. Park, Dfa-ad: a distributed frame-  
823 work architecture for the detection of advanced persistent threats, Cluster  
824 Computing 20 (2017) 597–609.
- 825 [29] P. Bhatt, E. T. Yano, P. Gustavsson, Towards a framework to detect multi-  
826 stage advanced persistent threats attacks, in: Service Oriented System En-  
827 gineering (SOSE), 2014 IEEE 8th International Symposium on, IEEE, pp.  
828 390–395.
- 829 [30] B. Binde, R. McRee, T. J. OConnor, Assessing outbound traffic to uncover  
830 advanced persistent threat, SANS Institute. Whitepaper (2011).
- 831 [31] J. de Vries, H. Hoogstraaten, J. van den Berg, S. Daskapan, Systems for  
832 detecting advanced persistent threats: A development roadmap using intel-  
833 ligent data analysis, in: Cyber Security (CyberSecurity), 2012 International  
834 Conference on, IEEE, pp. 54–61.

- 835 [32] J. Vukalović, D. Delija, Advanced persistent threats-detection and defense,  
836 in: Information and Communication Technology, Electronics and Micro-  
837 electronics (MIPRO), 2015 38th International Convention on, IEEE, pp.  
838 1324–1330.
- 839 [33] G. Vert, A. L. Claesson-Vert, J. Roberts, E. Bott, A technology for detection  
840 of advanced persistent threat in networks and systems using a finite angular  
841 state velocity machine and vector mathematics, in: Computer and Network  
842 Security Essentials, Springer, 2018, pp. 41–64.
- 843 [34] C. FIU, 25th geni engineering conference (gec25)-part iii, 2017.
- 844 [35] M. Lee, D. Lewis, Clustering disparate attacks: mapping the activities of the  
845 advanced persistent threat, Last accessed June 26 (2013).
- 846 [36] P. Giura, W. Wang, Using large scale distributed computing to unveil ad-  
847 vanced persistent threats, Science J 1 (2012) 93–105.
- 848 [37] Z. Saud, M. H. Islam, Towards proactive detection of advanced persistent  
849 threat (apt) attacks using honeypots, in: Proceedings of the 8th International  
850 Conference on Security of Information and Networks, ACM, pp. 154–157.
- 851 [38] N. Virvilis, D. Gritzalis, The big four-what we did wrong in advanced per-  
852 sistent threat detection?, in: Availability, Reliability and Security (ARES),  
853 2013 Eighth International Conference on, IEEE, pp. 248–254.
- 854 [39] Z. Zulkefli, M. M. Singh, N. H. A. H. Malim, Advanced persistent threat  
855 mitigation using multi level security–access control framework, in: Interna-  
856 tional Conference on Computational Science and Its Applications, Springer,  
857 pp. 90–105.
- 858 [40] E. M. Hutchins, M. J. Cloppert, R. M. Amin, Intelligence-driven computer  
859 network defense informed by analysis of adversary campaigns and intrusion  
860 kill chains, Leading Issues in Information Warfare & Security Research 1  
861 (2011) 80.
- 862 [41] P. Hu, H. Li, H. Fu, D. Cansever, P. Mohapatra, Dynamic defense strategy  
863 against advanced persistent threat with insiders, in: Computer Communica-  
864 tions (INFOCOM), 2015 IEEE Conference on, IEEE, pp. 747–755.


- 865 [42] I. Radware, Defenseflow - sdn based network, application ddos and apt protection, 2014.  
866
- 867 [43] M. Ammar, M. Rizk, A. Abdel-Hamid, A. K. Aboul-Seoud, A framework  
868 for security enhancement in sdn-based datacenters, in: New Technologies,  
869 Mobility and Security (NTMS), 2016 8th IFIP International Conference on,  
870 IEEE, pp. 1–4.
- 871 [44] M. Saher, J. Pathak, Malware and exploit campaign detection system and  
872 method, 2014. US Patent App. 14/482,696.
- 873 [45] N. Foster, R. Harrison, M. J. Freedman, C. Monsanto, J. Rexford, A. Story,  
874 D. Walker, Frenetic: A network programming language, in: ACM Sigplan  
875 Notices, volume 46, ACM, pp. 279–291.
- 876 [46] W. M. Eddy, Tcp syn flooding attacks and common mitigations (2007).
- 877 [47] F. Gont, Icmp attacks against tcp (2010).
- 878 [48] K. Tools, hping3. ICMP or SYN flooding tool,  
879 <https://tools.kali.org/information-gathering/hping3>, 2014.
- 880 [49] S. Shekyan, SlowHTTPTest. Application Layer DoS attack ,  
881 <https://github.com/shekyan/slowhttpstest/wiki>, 2011.
- 882 [50] Y. Zhang, S. Debroy, P. Calyam, Network-wide anomaly event detection  
883 and diagnosis with perfsonar, IEEE Transactions on Network and Service  
884 Management 13 (2016) 666–680.
- 885 [51] Y. Zhang, P. Calyam, S. Debroy, M. Sridharan, Pca-based network-wide cor-  
886 related anomaly event detection and diagnosis, in: 2015 11th International  
887 Conference on the Design of Reliable Communication Networks (DRCN),  
888 pp. 149–156.
- 889 [52] R. S. Boyer, J. S. Moore, Mjrty—a fast majority vote algorithm, in: Auto-  
890 mated Reasoning, Springer, 1991, pp. 105–117.
- 891 [53] T. G. Dietterich, et al., Ensemble methods in machine learning, Multiple  
892 classifier systems 1857 (2000) 1–15.
- 893 [54] Scapy: Packet manipulation tool, <http://www.secdev.org/projects/scapy/>,  
894 2007.



- 895 [55] M. K. Daly, Advanced persistent threat, Usenix, Nov 4 (2009) 2013–2016.
- 896 [56] E. Cole, Advanced persistent threat: understanding the danger and how to  
897 protect your organization, Newnes, 2012.
- 898 [57] I. Ghafir, V. Prenosil, Advanced persistent threat attack detection: An  
899 overview, International Journal of Advances in Computer Networks and  
900 Its Security (IJCNS) (2014).
- 901 [58] A. J. T.-F. Yen, Sherlock holmes and the case of the advanced persistent  
902 threat (2012).
- 903 [59] M. Ask, P. Bondarenko, J. E. Rekdal, A. Nordbø, P. Bloemerus, D. Pi-  
904 atkivskiy, Advanced persistent threat (apt) beyond the hype, Project Re-  
905 port in IMT4582 Network Security at Gjøvik University College, Springer  
906 (2013).
- 907 [60] C. J. Anderson, N. Foster, A. Guha, J.-B. Jeannin, D. Kozen, C. Schlesinger,  
908 D. Walker, Netkat: Semantic foundations for networks, in: Proceedings of  
909 the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Program-  
910 ming Languages, POPL '14.
- 911 [61] T. Neely, M. Vassel, N. Chettri, R. Neupane, P. Calyam, R. Du-  
912 rairajan, Dolus for ddos attacks defense - open repository  
913 <https://bitbucket.org/travisivart/dolus-defensebypretense>, 2018.
- 914 [62] T. Neely, M. Vassel, N. Chettri, R. Neupane, P. Calyam,  
915 R. Durairajan, Dolus for apt attacks defense - open repository  
916 <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.

**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.