

Expect More from the Networking: DDoS Mitigation by FITT in Named Data Networking

Zhiyi Zhang
UCLA

Vishrant Vasavada
UCLA

Siva Kesava Reddy K
UCLA

Eric Osterweil
GMU

Lixia Zhang
UCLA

Abstract

Distributed Denial of Service (DDoS) attacks have plagued the Internet for decades, but defenses have not fundamentally outpaced attackers. Instead, the size and rate of growth in attacks have actually outpaced carriers' and DDoS mitigation services' growth. In this paper, we comprehensively examine ways in which Named Data Networking (NDN), a proposed data-centric Internet architecture, fundamentally addresses some of the principle weaknesses in today's DDoS defenses in IP networking. We argue that NDN's architectural changes (even when incrementally deployed) can make DDoS attacks fundamentally more difficult to launch and less effective.

We present a new DDoS mitigation solution – Fine-grained Interest Traffic Throttling *FITT*, to leverage NDN's features to combat DDoS in the Internet of Things (IoT) age. *FITT* enables *the network* to detect DDoS directly from feedback from victims, throttle DDoS traffic along its exact path in the network, and perform reinforcement control over the misbehaving entities *at their sources*. In cases like the Mirai attacks, where smart IoT devices (smart cameras, refrigerators, etc.) were able to cripple high-capacity service providers using diverse DDoS Tactics Techniques and Procedures (TTPs), *FITT* would be able to precisely squelch the attack traffic at its distributed sources, without disrupting other legitimate application traffic running on the same devices. *FITT* offers an incrementally deployable solution for service providers to effectuate the application-level remediation at the sources, which remains unattainable in today's DDoS market. Our extensive simulations results show that *FITT* can effectively throttle attack traffic in a short time and achieve over 99% legitimate traffic.

1 Introduction

Distributed Denial of Service (DDoS) attacks have plagued the Internet for decades, and often capitalize on inherent properties of today's TCP/IP networking model. Starting

with early DDoS examples (e.g., attacks from the Trin00 botnet in 1999 [13]) through to recent attacks from the Mirai botnet [6], the remediation techniques used suggest that our defensive tactics may not be fundamentally keeping pace with attackers. Rather, with attacking botnets swelling in size to hundreds of thousands, and even millions, attacks have grown large enough that their attack volume rivals provisioned capacity of DDoS mitigation providers. The Mirai botnet serves as a quintessential example, in that it was used to launch some of the largest DDoS attacks in history, and it did so using compromised devices that primarily included Internet of Things (IoT) devices and household appliances that were both easily discoverable and poorly protected [6]. In many noteworthy modern instances, DDoS has evolved to being more distributed than ever and to using increasingly application-level semantics (e.g. reflective amplification attacks using DNS, NTP, memcached, etc.). Service operators, providers, and mitigation services [2, 31, 14] have had little recourse but to centralize defenses and backhaul and disrupt application semantics of malicious traffic, or to absorb undisrupted attack traffic ("packet love") in large DDoS mitigation service networks. In efforts to meet the distributed threat posed by DDoS with distributed remediation in incrementally deployable ways, approaches like BGP's FlowSpec [27], Remote Triggered Back-Holing (RTBH) [23], etc. have attempted to coordinate defenses at the network-level. However, DDoS Tactics, Techniques, and Procedures (TTPs) are sufficiently nuanced as to need Deep Packet Inspection (DPI), and network-level remediation lacks the necessary expressiveness to encode the TTPs in network-level remediations. This has, therefore, led to collateral damage or lack of adoption of these protocols and techniques.

In this paper, we show that a new incrementally deployable technique can be within our reach. We examine the Named Data Networking (NDN) [39] (Figure 1) to investigate whether DDoS defense can benefit from *its* architectural changes. NDN changes the basic network communication model and directly brings application-layer data names to




Figure 1: TCP/IP Architecture and NDN

the network layer: instead of pushing packets to an IP, in an NDN network users request named *Data packets* by sending *Interest packets* that carry the desired data name (Figure 2). The network forwarders will record the state of Interest packets, making breadcrumb traces for returning requested data packets. This stateful forwarding provides abundant traffic insights. These two key designs make NDN itself harder to be DDoS attacked and provide a solid foundation for DDoS defense mechanisms, which can be effective with incremental rollout of NDN. We argue that the architectural changes of the network, even with incremental deployment, can make DDoS attacks fundamentally more difficult to launch and less effective.




Figure 2: NDN Interest packet and Data packet

NDN’s receiver-driven model eliminates attacks by data packets, however attackers may attack a target in an NDN network by flooding various types of Interest packets. To address this attack vector, we present an NDN DDoS mitigation solution – Fine-grained Interest Traffic Throttling (*FITT*) that leverages NDN’s semantically meaningful names and stateful forwarding to combat Interest DDoS attacks. As we show later in this paper, FITT achieves the following desirable goals:

- Direct DDoS detection by feedback from *the victim*.
- Fine-grained traffic throttling at distributed sources on the specific attacking traffic flows at application level.
- Reinforcement control to further distinguish attack traffic and minimize the collateral damage to legitimate traffic.

In section 6, we elaborate how today’s service providers can deploy NDN with FITT incrementally over the existing network infrastructure to mitigate DDoS. We posit that the IoT industry (a major source of Mirai attack traffic) may even be incentivized to lead this deployment, and thereby poised to

offer de facto Internet-wide benefits. Indeed, finding operationally feasible security models and incentives to secure IoT devices have been an open research challenge [20, 42].

The contributions of this paper are twofold.

- To the best of our knowledge, this paper is the first comprehensive description of how NDN’s architectural design decisions lead to an inherently resilient foundation for DDoS defense.
- We propose a specific solution, FITT, to mitigate DDoS attacks more accurately than existing solutions for TCP/IP networks. We also compared FITT with other proposed countermeasures for NDN and other Information Centric Network (ICN) architectures to explain its superiority.

Our analysis illustrates that the network can do more to protect application services than we currently expect. By reconsidering the architectural design, we can augment network security in a fundamental way: *let applications instruct the network to squelch DDoS at its sources*.

2 DDoS Mitigation over TCP/IP Architecture

2.1 DDoS and Vulnerabilities in IP

Many of the largest volumetric DDoS attacks that we have seen today rely on architectural properties of TCP/IP. For example, the Mirai botnet was used to generate unprecedented DDoS traffic volume, causing millions of dollars (in losses and mitigation service costs). In part, this seems to derive from the large numbers of sources that are relatively easy to compromise: IP address enumeration and scanning trivialize the discovery of poorly protected IoT devices, which can then be used to build the biggest botnet ever. Another example is that reflective amplification attacks capitalize on source address spoofing and redirect traffic to any IP addresses [6]. DDoS attacks have large variances in their TTPs and combating them is often done by application-level remediation (in scrubbing systems, appliances, app configurations, etc.)

Previous works [22, 11, 37, 19, 33] have observed that DDoS attacks often exploit the following properties in IP:

- *Push-model Communication*: Any Internet node can send packets to any other IP address.
- *Destination-based Delivery*: Packet delivery is solely based on the destination address and there is no source address validation by default, thus source IP addresses can easily be misattributed.
- *Limited Expressiveness in IP Forwarding*: IP’s forwarding system cannot embody the semantic characteristics of application-level traffic, which makes it difficult for DDoS defense mechanisms to inspect DDoS attacks.

Based on these observations of IP networking, [19, 11] propose modifications to the existing Internet architecture to have greater DDoS resilience at an architectural level. A

number of desired features of a DDoS-resistant Internet architecture are presented, the major ones include (i) limiting the accessibility to a server based on the server’s capabilities, (ii) source address authentication to prevent source address spoofing, (iii) separating client and server address space to prevent unwanted traffic from client to client and server to server, and (iv) building symmetric traffic flows to prevent reflection attacks at the network layer. In this paper, our examination of NDN in Section 3 shows that NDN’s architecture natively embodies these notions.

2.2 Network-centric DDoS Mitigation

Deployable filter-based network-level remediation approaches in TCP/IP like FlowSpec [27], RTBH [23], the IETF’s Distributed Open Threat Signaling (dots) Working Group [29] etc., lack the necessary expressiveness of DDoS’s TTPs. Such mechanisms would additionally require Deep Packet Inspection (DPI) to gain insight into the ongoing traffic which IP’s stateless forwarding cannot provide. For instance, black-hole filtering blacklists entire network prefixes, which can cause collateral damage to: well-behaved sources, non-attack traffic that is sourced from compromised devices, etc. As another example, FlowSpec requires the proper n-tuple consisting of several matching criteria so that DDoS traffic can be classified; however, since most supported matching criteria is at the network layer (e.g., IPv4, IPv6, ICMP), FlowSpec can mistakenly drop legitimate traffic to other services deployed on the same victim server or block good traffic sent from the compromised bots.

Ioannidis and Bellovin proposed router-based Pushback [21] which utilizes a heuristics function to detect packets that *probably* belong to an attacker by checking the “congestion signature” (e.g., source/destination addresses) of the traffic. Due to routers’ coarse-grained inspection of traffic, the filtering can lead to collateral damage.

Another interesting work is Active Internet Traffic Filtering (AITF) [8], which requires the routers on the path of attacking traffic to mark traffic flows with route records (RRs), which are the IP addresses of routers who have forwarded the flow. In this way, the victim can report the unexpected flows to the network, and routers can filter the flows identified by the RRs. However, this may potentially cause damage to (i) all the traffic (both attack traffic and legitimate) sent along the path, because RRs cannot distinguish application-level flows, and (ii) legitimate hosts who are under the same first hop network as the attackers, because routers do not have fine-grain state to identify exact senders.


Some other network-based works like StopIt [26] and SIFF [36] require additional features that are missing in the existing TCP/IP architecture. For example, SIFF introduces privileged communication, which requires additional information carried by the IP header and each router on the connection path marking IP packets. Capability-based ap-

proaches like TVA (Traffic Validation Architecture) [37, 25] introduce authentication of the packet source into the network system. Taking TVA as an example, by embedding cryptographic authentication info into the IP packet, the routing system and servers are able to distinguish legitimate users from “bad” ones. The primary difficulties in operationally deploying these proposed solutions are adding extra functionality into the deployed TCP/IP architecture and an incentive misalignment similar to Pushback.

3 NDN’s Properties for DDoS Mitigation: A Comprehensive Examination

3.1 Named Data Networking

Named Data Centric Named Data Networking (NDN) makes named data the thin waist of the network architecture. More specifically, applications name their data at the application layer and NDN directly uses the namespace of applications for network layer data delivery. These data names are semantically meaningful and structured, e.g., a video produced by Alice’s device may have the name “/univ1/cs/alice/video.mp4”. In NDN, routing and forwarding the packets are based on *name prefixes*. Figure 3 shows a simple illustration of how an Interest packet is forwarded to fetch the data.



The Interest has a name “/univ1/cs/alice/video/demo.mp4”. Each forwarder along the path forwards the Interest packet based on the forwarding information table (FIB) using the longest prefix match.

Figure 3: Data Fetching with an Interest

Stateful Forwarding NDN utilizes a stateful forwarding plane: forwarders will record each Interest packet toward data sources, and the fetched Data packet will strictly follow, in reverse, the path taken by the corresponding Interest to get back to the requesting entity. Since an NDN network concerns about data instead of locations, multiple Interest packets requesting the same Data packet are merged in the network (called Interest Aggregation).

NDN’s forwarding module realizes stateful forwarding by introducing a Pending Interest Table (PIT) into each router. The PIT stores currently unsatisfied (pending) Interests together with their incoming/outgoing interfaces. When a Data packet arrives, the router sends the Data packet to all incoming interfaces recorded in the corresponding PIT entry and removes this PIT entry; the replied Data can be cached in

the router’s Content Store (CS) to satisfy future Interests requesting the same piece of data.

In addition to Interest and Data packets, either NDN routers or data producers may generate NACK packets, which serve as a hop-by-hop feedback mechanism to report a problem in further forwarding of Interests. When a router receives such a NACK packet, it takes appropriate action(s) based on reason code carried in the NACK packet.




Figure 4: NDN Forwarding

Built-in Security NDN builds communication security [41] into the architecture by requiring data producers to cryptographically sign all data packets at the time of production and, if needed for content confidentiality, encrypt them as well. Securing data packets directly enables routers to cache them as they pass along, and enables consumers to validate Data packet regardless of where and how they are fetched.

3.2 NDN’s DDoS Mitigation Properties

In this section, we examine NDN’s architectural advantages over DDoS defenses in the TCP/IP architecture in terms of DDoS resiliency. We identify the following advantages: (i) NDN’s Interest-Data packet exchange eliminates network-layer reflection attacks and DDoS attacks by flooding Data packets. (ii) Its data pull model and securing data directly make it more difficult for attackers to recruit “zombie armies”. (iii) NDN’s Interest aggregation and cache automatically relieve the overload caused by Interests for static and existent Data. (iv) NDN’s stateful forwarding provides rich insights into ongoing traffic for DDoS defense mechanisms.

3.2.1 Off By Design

For devices that serve content in TCP/IP servers, DDoS threats begin as soon as a service goes online and becomes reachable. By contrast, the communication in NDN follows a pull model and an application or a node is considered as “off by design” [11] for the following reasons: 1. One cannot send an Interest to a consumer application, or a producer application whose name is unreachable from the sender. By simply not announcing its prefix to solicit Interest packets, an application can still pull data from others, but can *not* be reached by an Interest packet, thus reducing the attack surface for malware infections and for DDoS attackers. 2. A

Data packet cannot go anywhere if it is not requested, because there is no corresponding Interest path. 3. Flow-parity: one Interest can at most bring one Data packet back. With the pull model, an attacker cannot launch DDoS by flooding Data packets, thus network layer DDoS attacks can only be carried out by Interest flooding. However, unlike in TCP/IP, as Data Packets follow the reverse path of their corresponding Interests, an attacker cannot redirect them to another consumer. By design NDN *fundamentally* eliminates reflection DDoS attacks at the network layer.

3.2.2 Data Name Instead of IP Address

In TCP/IP, even clients who do not run services can be attacked, compromised, and enslaved. One very common intrusion TTP of attackers is to scan the IP address space in order to discover devices, and then compromise them. In NDN, however, if an end device does not *serve* data, it does not even need a name and can eliminate the attack surface of being exposed at all. More so, routers in NDN forward an Interest by its name. NDN names are defined by the application semantics in an arbitrary format that are not enumerable. For example, a smart home device with an application-defined name “/my/name/home/refrigerator-02” is less exposed compared with an IP address with a default port number, like Mirai exploits [7], because in NDN the exposed network prefix may only be “/my/name” and “guessing” the exact name requires reconnaissance. If an Interest name does not match a specific prefix in the forwarding table, the packet will get dropped by the router. Using application-defined names fundamentally makes a source more difficult to be found and then compromised.

Another big benefit of using a name is to allow the network to inspect traffic at a much finer granularity. For example, a compromised smart home refrigerator, in a Mirai botnet, may be carrying out network transactions with its device-manufacturer while *also* being forced to participate in a DDoS attack. There, the legitimate Interest traffic might have prefix “/iot-provider/service” and attacking traffic might be towards prefix “/univ1/service/email”. Since names of data are directly exposed to the network, the infrastructure is able to identify specific application-level traffic flows and squelch just them (and not the legitimate traffic).

3.2.3 Reduce DDoS traffic with Cached Data and Interest Aggregation

NDN’s content-centric communication model provides enhanced data availability by enabling caching inside the network (i.e. the CS in NDN Forwarders). Because of the in-network caching, Data packets carrying static content (e.g. HTML files, CSS files, images) can be cached by routers to satisfy future Interest packets, thus reducing the number of Interests reaching the producer (victim).

In addition, Interests targeting the same piece of the named data will be aggregated by the router and later Interest packets will not be sent out. This feature makes it harder for DDoS attackers to flood the same Interest packet or a small set of Interest packets towards the producer in a short time.

As shown in previous work [32] and our simulation results in Section 7.1, in-network caching and Interest aggregation help to mitigate the Interest flooding where attackers send Interests for static or existing Data packets. However, if attackers flood a target prefix with a large set of Interest packets or even fake Interests with randomly generated components, the benefits will diminish. This is because churn and evictions in the cache will lead to diminishing cache hit-rates and the chance of two attacking Interests sharing the same name drops. However, it is also noteworthy that NDN is incrementally deployable and does not immediately require rich deployment and caching in the routed core of networks in order to function properly.

3.2.4 Rich Traffic Insight by Stateful Forwarding

As NDN’s deployment pervades more of the routing infrastructure, its stateful forwarding [38] provides rich insight into ongoing traffic. Different from a router in TCP/IP, which has little knowledge about which downstream interface attackers are behind, stateful forwarding in NDN helps forwarders to know exactly which interface the traffic is coming in from, by design. This helps NDN traceback to misbehaving clients and reinforces mitigation. By observing each Data packet and its corresponding pending Interest entry in the PIT, an NDN forwarder is able to measure the round-trip time, throughput, and name reachability of each outgoing interface. As mentioned in [1, 15], a forwarder can also learn the Interest satisfaction ratio, namely the proportion of Interests that successfully fetched a Data packet, and thus detect possible fake Interest DDoS attacks. Moreover, PIT entry timeouts also offer relatively cheap DDoS attack detection, as mentioned in [39].

3.3 A taxonomy of DDoS attack using NDN Interests

In NDN, attackers can only attempt to DDoS a target by flooding it with Interests. Specifically, inspired by the work [18], we categorize Interest packets used in DDoS attacks into I-1, I-2, and I-3 according to (i) whether the Interest is valid or not and (ii) whether the target Data requires the producer to generate data in real time. Among the three types, I-1 and I-3 Interests are *valid Interests* because they can fetch a Data back in theory while I-2 Interests are *fake Interests*.

I-1: Valid Interest for static or existing data The target Data packets do not require real-time processing; for instance, the Data packet carrying a static CSS file or a video

chunk. Type I-1 attacks can be naturally mitigated by NDN’s Interest aggregation and in-network cache as we discussed in Section 3.2.3 and Section 3.2.3. However, as analyzed and shown in Section 7.1, when flooding Interests to a large enough number of names, the effect of Interest aggregation and in-network cache becomes insufficient.

I-2: Invalid Interests When attackers send large volumes of Interests that cannot be satisfied by the producer, we consider such traffic to be invalid. Some possible reasons could be (i) the Interest format is incorrect (e.g., unexpected name components, invalid Interest signature) and (ii) the target Data packets are non-existent and can never be generated when sent in large volumes. This type of Interest is mostly useful to malicious adversaries. One possible way for attackers to generate such Interests is to append non-existent name components (e.g., garbled text generated randomly) to valid server prefixes. Since I-2 Interest names can be arbitrary and there is no target Data packets, Interest aggregation and caching cannot mitigate such attacks.

I-3: Valid Interests for dynamically-generated data When an Interest for dynamically-generated Data arrives, the server needs to process the requests (e.g., database queries, calculations, etc.) before it can generate Data and reply back. For example, a legitimate Interest name may contain variable components and thus there are no preexisting Data packets to match it. In this case, the server will need to process the request first and then generate a Data packet following with a reply. Since the Interest’s name is customized and the Data is generated in real time, hardly any Interests arriving at a forwarder would hit a pending Interest with the same name or a previously cached Data packet, so NDN itself does not effectively mitigate such attacks.

4 Fine-grained Interest Traffic Throttling

To address the attack surface of Interest flooding and to enable data sources to identify and push fine-grained network remediation to attack sources, we present FITT. In this section, we use the topology shown in Figure 5 to help explain the attack scenarios and the FITT design. The *server S* runs an NDN producer application and serves the data prefix *P* “/univ1/service/email”, which is the target of a DDoS attack. The *clients C1 to C6* are consumer nodes. We let *C1, C2, C3* to be compromised IoT devices in a botnet, e.g. Mirai, and *C4, C5, and C6* to be legitimate clients. The *routers R1 to R5* are routers that support NDN’s network stack. Importantly, these routers may be deployed as an *overlay* topology over the real world IP topology. Based on the Data packet flows, we call the routers towards the server “upstream routers” and routers towards the clients “downstream routers”, e.g., *R4 and R5* are downstream routers of *R3*.

Goals Our proposed system aims to defend not only from I-2 attacks, but also from I-1 attacks, I-3 attacks, and mixed attacks (mix of I-1, I-2, I-3) in an effective way without affect-




Figure 5: An example FITT topology as an overlay of current TCP/IP network

ing legitimate clients and the traffic under another prefix. For example, in our example topology, when the email service “/univ1/service/email” is under attack, our proposed approach should throttle the DDoS traffic away from the email service only, at the same time, legitimate traffic from clients under other prefixes (e.g., “/univ1/service/video”) should not be affected (even if that traffic is from attacking clients).

Assumptions The FITT design is based on the following assumptions that we argue are either reasonable or easy to be realized.

- *Victim S is best able to know its capacity to process incoming Interests under a specific prefix.*

The server’s capacity can be easily obtained based on the provisioned memory, CPU and other resources versus the time/space complexity of processing the Interest requests under the prefix.

- *The server S (the victim) is best able to know whether it is under a DDoS attack, the prefix that is under attack, and whether an Interest is fake or valid.*

The victim server inherently has the most accurate judgments of a DDoS attack: By simply inspecting whether Interests under a prefix overwhelm the processing power, S knows whether prefix is under attack. Moreover, the server may immediately know whether an Interest is fake or not when processing it because a fake Interest cannot fetch a Data.

- *A node knows whether it is a FITT edge router.*

In FITT, an edge router is either (i) a gateway router to which clients are connected or (ii) a far-most router from the DDoS victim that supports FITT and NDN in incremental deployment cases. This can easily be configured by the Internet Service Provider (ISP) at the network level (e.g. CPE, SOHO router, etc.) or by the device vendor (e.g., IoT vendor) at the overlaid NDN level. The information can also be obtained through automatic means; for

example, to learn whether it is connected to a client or not, the router can check the hop count of incoming packets from that interface.

4.1 A Design Overview

FITT is designed to be running on each NDN forwarder as a NDN forwarding strategy¹. Utilizing the topology in Figure 5, FITT works in a way that after receiving DDoS feedback (carried in a NACK packet from the victim S), the router R1 analyzes the feedback and triggers the FITT reaction. By checking the incoming interfaces stored in the PIT, R1 will trace back to the DDoS traffic sender routers R2 and R3 and notify them by NACK packets. R2 and R3 will perform the similar procedures as R1 does. In this way, FITT pushes remediation from S all the way to *edger routers* like R2 and R4 where exact traffic senders are connected. The edge routers will first notify these clients by NACKs and perform Interest throttling on suspect downstream interfaces within the specific prefix reported by S. During the traffic throttling, an edge router will check whether a client has changed its behavior following the NACK or not (i.e. whether it lowers down its sending rate to the required value under the specified prefix). The router can then relax or reinforce the limit, accordingly.




Figure 6: FITT System Overview

Multiple FITT reactions can be triggered at the same time for different traffic prefixes (located on the same server or different servers) and different types of attacks. When multiple reactions take place, for a specific traffic flow under the prefix and from a suspect interface, a FITT edge router will take the minimum allowed traffic value to perform the throttling.

¹In an NDN forwarding module, forwarding strategies decide the forwarding operations. Adding a new forwarding strategy requires no modification to the forwarding module design

4.2 Explicit Feedback from the Victim

A FITT reaction is triggered by a victim’s NACK packet. To be more specific, a NACK packet created by the victim server and sent downstream will carry the following information:

- *RSN*: The **reason** code used to notify routers whether the victim is under fake Interest (I-2) attack *FAKE* or valid Interest (I-1 or I-3) attack/overload *VALID*.
- *PREF*: The **prefix** under which the overwhelming traffic comes to the victim.
- *C*: The receiving rate of valid Interests that the server can handle currently under the prefix *PREF*. We call the number **capacity**. The capacity is carried in a NACK only in case of valid Interest overload.
- *FakeList*: **Fake Interest name list** for fake Interest attack only. The list contains fake Interest names under the prefix *PREF* that a server received within the last unit time interval which is defined by the server. Optimizations can be applied to reduce the space complexity by (i) sampling the list, (ii) utilizing Bloom Filter [12], or (iii) regular expression

4.3 Pushing Back to Exact Sources

After receiving a FITT NACK, the router closest to the victim will first check the *RSN* field and trigger different types of reactions.

If *RSN = FAKE*, the router will first check NACK’s *FakeList* and find out the corresponding PIT entries whose Interest name is in the *FakeList*. Through these entries, the router learns the exact incoming interfaces of the fake Interests. For each of these interfaces *i*, the router will generate a new fake Interest name list *FakeList_i* which only contains *FakeList*’s fake Interest names that were sent from interface *i*. After that, the router sends each interface *i* a new NACK carrying the *RSN*, *PREF*, and *FakeList_i*.

If *RSN = VALID*, there is no *FakeList* carried in the NACK. Different from the reaction to a fake Interest attack, since both attackers and legitimate clients send valid Interests, the router cannot distinguish the good traffic from the offending. Consequently, the router will check all the current pending Interests under the prefix *PREF* and get a set of suspect incoming interfaces. The router will then calculate a weight w_i for each suspect interface *i* to distribute the *C* to these interfaces. Since a router do not have the knowledge of the clients behind each interface (i.e., how many end hosts are behind each interface, whether these hosts are attackers or not), FITT adopts the simplest way of equally share the weight to all suspect interfaces:

$$w_i = \frac{1}{Num_Suspect} \quad (1)$$

where *Num_Suspect* denotes the number of the suspect interfaces. As described in Section 5.4, the reinforcement throttling will help to further identify attacking traffic and amend the potential unfairness caused by the equal share. After that, for each suspect interface *i*, the router then computes the weighted capacity $C_i = w_i \times C$ and sends a new NACK carrying new capacity value C_i as well as *RSN* and prefix *PREF*.

All the procedures mentioned above are within the prefix *PREF* reported by the victim, traffic under other prefixes will not be counted in.

In this way, all the routers along the Interest sending path recursively receive and generate new NACKs that will be propagated to further downstream routers. Finally, the DDoS report that originated from the server will arrive at all FITT edge routers.

4.4 Reinforcement Traffic Throttling

When a FITT edge router receives a NACK from the upstream, it will perform traffic throttling to downstream interfaces (connecting to the clients or further routers) where DDoS traffic is from. To be specific, the router will first calculate the permitted sending rate $Limit_i$ for each interface *i* under the prefix *P* as follow.

$$Limit_i = \begin{cases} C_i, & \text{if } i \text{ is suspicious when } RSN = VALID \\ 0, & \text{if } i \text{ is suspicious when } RSN = FAKE \\ \infty, & \text{if } i \text{ is not suspicious} \end{cases}$$

In case of *RSN = VALID*, similar to other FITT routers, the router will calculate the assigned capacity C_i . When *RSN = FAKE*, FITT doesn’t tolerate any fake Interest (I-2) and imposes zero permitted Interest sending rate under the prefix *P*.

The router then sends a NACK packet to each suspect interface as a notification. After the notifications, the edge router then starts throttling traffic by randomly dropping Interest packets to ensure:

$$\forall i \in range(1, n) \quad Rate_{(i, P)} \leq Limit_i$$

where $Rate_{(i, P)}$ is the Interest sending rate from interface *i* and under prefix *PREF* and $Limit_i$ is the allowed traffic sending rate for the interface *i*.

Once receiving a NACK from the gateway router, legitimate clients will comply and lower down their sending rate of the Interests under the prefix *PREF*, while the bots may not obey the rules, prompting the router to perform reinforcement throttling. In case of valid Interest (I-1 or I-3) attack, the router will monitor the sending rate of each suspicious client and perform the following control:

- If a sender lowers its Interest sending rate $Rate_{(i, P)}$ down to $Limit_i$, the router will remove the throttling over this client.

- If a sender does not follow the control, the router will reset the limit to $\frac{1}{2} \times Limit_i$.

All the throttling is within the DDoS traffic flow defined by the prefix *PREF*. This is how FITT squelches DDoS traffic at the source while still letting the sources communicate to other services (names).

4.5 Timers in FITT

FITT requires routers to keep the FITT records when pushing NACKs down to edges and when edge servers perform the throttling. FITT utilizes two types of timers **RevertTimer** and **RateLimitTimer**. These timers limit the overhead of FITT and will not affect the final result of the FITT reaction. Different routers can set these two timers differently based on their own needs.

A *RevertTimer* decides how long a router should keep the DDoS Records. The timer is set when a new NACK arrives. Whenever a new NACK arrives, the router checks whether there is an existing *RevertTimer* for the FITT reaction with the same reason *RSN* and same prefix *PREF*. If yes, the router will update the timer instead of creating a new one.

A *RateLimitTimer* is maintained by edge routers only. It decides how long it takes for a router to decide whether an individual is well-behaving or not. After the *RateLimitTimer*, the gateway router will remove the limit of “good” clients and strengthen the limit of bad ones. This timer is periodically reset until all clients either behave well or move into the black list.

5 FITT Design Rationale

FITT directly makes use of NDN’s properties to combat DDoS in the following ways:

- Utilizing structured names at the network layer allows FITT to perform service-level, per-prefix, reaction and monitoring.
- Stateful forwarding provides exact information about the traffic so that FITT can accurately identify suspicious clients.
- Remediating traffic at sources alleviates network congestion and capacity concerns of data producers and intermediate network infrastructure (like Internet eXchange Points, IXPs, etc.).

5.1 Victim’s Feedback enables Accuracy

The explicit feedback enables FITT to perform fine-grained DDoS mitigation. To be specific, the prefix *PREF* helps FITT narrow the target traffic scope and the capacity *C* can inform the downstream routers about the percentage volume of traffic that should be controlled in case of valid Interest (I-1 or I-3) overwhelm. Moreover, hearing the feedback

also makes the reaction more accurate: NACK’s *RSN* helps routers to learn the attack type and take different actions accordingly. In case of fake Interest attacks, by tracking the incoming interfaces of the Interests listed in *FakeList*, routers are able to directly identify the attackers and completely block them.

Since the design is built to mitigate DDoS for a variety of prefixes and types of Interest Attacks, FITT is able to handle complex scenarios. For example, more than one victim server can send multiple NACKs containing different prefixes and different reasons.

5.2 Minimizing Collateral Damage by Fine Granularity

The fundamental benefit of the fine granularity of FITT is twofold: (i) By explicitly setting *PREF* in a NACK, FITT squelches the traffic to *PREF* only. All the other services provided by the victim server will not be affected. (ii) FITT throttles traffic sent by a suspect client to the prefix *PREF* only, letting the clients communicate to other services.

Considering the example in Figure 5, we assume the “/univ1/service/email” is under attack, FITT will limit the traffic to this service to the expected volume *C*, which are configured by *S*, or block the traffic consisting of fake Interest (I-2). In the throttling, non-attacking clients like *C4* and *C5* can use the service “/univ1/service/video” as normal, reducing the collateral damage on *S*’s services that are not attacked.

Using the same example, we assume *C3* is a compromised smart home device, e.g., a home camera. Though *C3* is compromised to send attack traffic to *S*, FITT is able to stop its DDoS traffic and at the same time, does not bother its normal functions. For example, it can still upload the surveillance video records to the smart home controller.

5.3 Distributed Throttling at Sources

In our design, only FITT edge routers play the role of rate limiting. This is because, on one hand, we cannot trust a client’s device to take actions - it could be compromised as well. On the other hand, an upstream router should not perform rate limiting for the following reasons: (i) When the traffic volume under a target prefix increases, upstream routers don’t have enough knowledge to tell whether it is because of the misbehaving downstream routers or new clients have joined. (ii) When legitimate clients are behind a downstream router, compared with edge routers, upstream router actions will also hurt legitimate clients. Note that this is one reason why attempts to remediate DDoS in IP networks (using FlowSpec, RTBH, etc.) suffer and are often maligned, because of the collateral damage that can result from their use.

In cases when an edge router is connected to downstream routers or clients that are not equipped with NDN and FITT, our proposed approach still reduces the attack surface by reinforcing the limitation on the suspect traffic. Because FITT may cause collateral damage to the legitimate clients behind the edge routers (since clients can't be reached in this case) FITT uses best-effort to mitigate the DDoS traffic to the victim.

5.4 Monitoring Further Identifies Attackers

After an edge router sends out NACKs to suspicious clients, we believe that the legitimate clients will obey the DDoS control and lower their sending rate of Interests accordingly while attackers may not abide by this. It is possible that bots may use intelligence to analyze and attempt to circumvent the NACK, but FITT already succeeds if the bots cannot increase Interest sending rate, thus greatly reducing damage. Essentially, FITT forces bad entities to comply.

In valid Interest attack scenarios, a router calculates the weight by simply equally distributed the capacity among all suspect interfaces by equation 1, which may potentially cause unfair capacity assignment to clients. For example, C1 and C2 in Figure 5 will be assigned more capacity than C3 and C4 in a valid Interest attack reaction because C1 and C2 are closer to the victim in the topology, thus having larger weight. Though there is potential for unfairness, the monitoring will help FITT to further restrict the attacking traffic and relax the limit on legitimate traffic by adjusting the throttling to be fair. As shown in simulation results in Section 7.3, the reinforcement will quickly block all attacking traffic and let legitimate traffic recover from the throttling.

6 Incremental Deployment of FITT




Figure 7: An Example Topology

NDN and FITT present an option for immediate incremental deployment, which elevates many of their advantages to near-term objectives. In general, incremental deployment models become more realistic when they align their costs with incentives. That is, those who deploy new mechanisms are more likely to do so when they anticipate direct benefits from doing so. By contrast, deployments like

ingress and egress filtering (BCP-38 [17] and BCP-84 [10]) illustrate slow adoption, arguably, because those deploying them do not gain any direct benefits. Conversely, service providers already expend resources and money to combat DDoS by either provisioning large amounts of excess bandwidth or by contracting with commercial DDoS mitigation providers [2, 31, 14]. Service providers whose applications may already be in a position to benefit from migrating to NDN's architecture would gain *additional* benefits by deploying NDN with FITT and thereby being able to shed large amounts of DDoS traffic. In an incremental deployment, the overlay NDN routers and two ends speaking NDN are enough for FITT to effectuate the DDoS mitigation, as shown in Figure 7. What's more, service operators who migrate their services to NDN bolster each others' NDN deployments, as those clients independently augment each others' deployments (through facilities like shared caching and shared routing infrastructure). In particular, we observe that serendipitous IoT deployments of NDN, which may already be underway, could benefit other services whose providers have (or will) independently embraced NDN for this reason. That is, an NDN-enabled service may shed DDoS traffic from would-be attack nodes that might otherwise be bots in Mirai. IoT deployment of NDN would put the FITT mitigation machinery very near to some of the Internets most voluminous DDoS sources for all NDN applications (not just IoT). We believe that it is demonstrably feasible for independent service operators to overcome network protocol ossification and migrate (at least portions) of their production traffic to NDN.

6.1 Overcoming Ossification

A common lament in the Internet has been that network protocols evolve very slowly, or tend to be ossified [5]. Recent deployment successes in other network protocols [24] has illustrated that this impasse can be overcome by providers who control (i.e. implement and deploy) both ends of a service (the client and server sides). When implementing mobile apps, the provider has the ability to choose both ends of the network protocol. While applications that depend solely on web browsers must often remain backward-compatible with TCP/IP, mobile clients can often implement service-specific code. The deployment of Google's QUIC [24] provides a timely example of this flexibility.

In that case, deployment grew quickly with Google's ownership of the transactions.² Legacy TCP support was maintained, but QUIC was treated as preferred where QUIC-compatible clients were used. We observe that this tactic is equally available to NDN, through mobile applications. As a migration path, and to maintain backward compatibility, service providers could bifurcate their deployments and

²Google was able to implement QUIC on its mobile platforms and its Chrome browser, but maintained TCP/IP support for other browsers.

offer TCP/IP services on separate infrastructure. Then, under cases like large DDoS attacks, TCP/IP could be serviced by different infrastructure, and all NDN/FITT infrastructure could remain unencumbered by attack traffic, while TCP/IP remediations are enacted on the legacy infrastructure.

6.2 Using IoT to Bootstrap

With increasing attention being paid to the NDN architectures utility to problem-spaces like IoT and smart-homes [9, 3, 4, 40], grass-roots deployments of NDN may already be beginning. IoT, in particular, represents a fortuitous opportunity because there would be strategic advantages to having NDN’s and FITT’s protections in front of devices that have historically participated in some of the Internet’s largest DDoS attacks [6]. This possibility could be attainable in the near term because in the IoT space, it is not uncommon for firmware upgrades to be automated and consumers to update their “infrastructure” more often, and the device lifetimes are often shorter than some network infrastructure (like core routers, DNS infrastructure, CDN caches, etc.). Such a setting could suggest that a measurably aggressive deployment of new devices is feasible, and NDN deployments could grow more quickly than in other settings. What’s more, the ability of NDN and FITT to address DDoS in the network could be seen as an incentive model for IoT developers to embrace security in their development by simply choosing the NDN network architecture.

As such deployments grow, this could present a serendipitous opportunity for *other* NDN and FITT deployments. For example, consider a product line of refrigerators that were to be deployed using NDN for their communications, and only support TCP/IP for legacy communications (to ensure reachability in the home). These refrigerators could still be scanned, compromised, and become bots in a Mirai botnet (through their TCP/IP stack). These devices, may, therefore, still be enslaved into participating in DDoS attacks. However, if an unrelated email service were receiving DDoS traffic from those devices, and were to move to NDN with FITT, the IoT devices would be able to transact with that same service over NDN instead of TCP/IP. At that point, the device would have the FITT mitigation machinery right next to the attack traffic, and the NDN service could quell the DDoS traffic from it.

7 Evaluation of NDN’s DDoS Resilience and FITT

We implement FITT in C++ over ndnSIM [28], which is a NDN simulation platform based on NS-3. The network topology that we use for our experimentation and evaluation has four Autonomous Systems (ASes) with meshed connections. “/univ1/cs/server/email” is the prefix that is under attack and the service provider is node “/univ1/cs/server”.


We assume the server is globally reachable, which means all users, even outside the server’s local network, have means to learn the name. For sake of simplicity, we will just call the target prefix P in rest of the section.

In the simulation result plots, we use the red dashed line to represent attackers’ sending rate, the blue solid line to represent P ’s receiving traffic rate, the green dotdash line to represent legitimate clients’ sending rate.

We first demonstrate NDN’s DDoS resilience to I-1 Interest attack and then evaluate FITT. The simulation results show that after the DDoS starts, FITT can effectively control the traffic to the victim as expected within seconds (less than 2 seconds under our simulation settings), and ensure that over 99% of the attack target(s) incoming traffic is from legitimate clients after a short period of time.

7.1 NDN’s DDoS resilience to I-1 Attack

Figure 8 demonstrates NDN’s DDoS resilience to I-1 Interest attack with the help of Interest Aggregation and in-network caching.



Simulation Settings: 60 attackers are located across all the ASes with no FITT deployed. Each attacker start sending I-1 attacking Interests at 100 Interests/s from second 3. We simulate NDN’s resilience with target’s available Data names to be 1000 and 500.


Figure 8: NDN’s DDoS Resilience to I-1 Interest Attack

We first disabled cache in all routers so that the result will only be affected by NDN’s Interest aggregation. As shown in Figure 8a, Interest aggregation can withhold traffic from attackers (red dotted line) to the server (blue solid line). We then introduced cache to see how it can suppress traffic even more. As shown in Figure 8b, it is apparent that the number of Interests reaching P decreases because of the caching capacity (i.e. forwarder’s Content Store), which is because intermediate nodes along the path will serve future same Interests with cached Data (the freshness of cached Data is 4 seconds in our simulation).

The two figures indicates that the effect of Interest aggregation and cache is lower when an attacker can use a bigger set of Interest names to attack the victim. This is because larger the name set, smaller the chance of two Interests car-

rying the same name and smaller the chance to hit a previous cached Data packets.

7.2 FITT: Fake Interest Attack




Simulation Settings: There are 60 attackers with sending rate 100 Interests/s and 12 legitimate clients with sending rate 40 Interests/s across all the ASes. I-1 attacking Interests starts from second 3.

Figure 9: FITT: Fake Interest Attack

We first study FITT’s performance against fake Interest (I-2) DDoS attack. As shown in Figure 9a, initially, P only receives Interests from legitimate clients (green dotted line). After second 3, attackers start the DDoS by sending I-2 Interests to P (the red dashed line goes to 6000 Interests/s). As depicted by the plot, FITT eliminate the DDoS traffic in seconds and P ’s receiving traffic line soon merges the legitimate clients’ outgoing traffic line, meaning good traffic from doesn’t get affected. The effectiveness is because in fake Interest attack, FITT can accurately identify attackers and throttle their traffic flows under the specific prefix. Figure 9b shows that after FITT reaction, all the traffic received by P is from legitimate clients.

7.3 FITT: Valid Interest Attack

The simulation results of the valid Interest Attack are shown in Figure 10. Different from the fake Interest attack simulations, after the reaction, since both legitimate clients and attackers send out valid Interest packets, as we discussed in Section 4, the router cannot tell good traffic from bad traffic. Therefore both legitimate clients and attackers will be limited. After receiving the FITT NACK, legitimate clients will abide by the control placed and lower down their sending rate until the router determines them to be legitimate and free the limits, explaining why the green dotted line goes down in the first several seconds of the attack and then back to the normal later. As for attackers, as shown, the traffic received by the victim drops periodically (every 3 seconds), which confirms the FITT’s reinforcement throttling: FITT will halve the limit on attackers until all the attackers’ traffic to the reported prefix are totally blocked. At the end of the




Simulation Settings: Using the same simulation settings as Figure 9, attackers send valid (I-1 and I-3) Interests to P from second 3. We let the victim’s capacity of handling Interests under P to be 1500 Interests/s and RateLimitTimer to be 3 seconds.

Figure 10: FITT: Valid Interest Attack

Pushback, almost all the Interests received by the victims are from legitimate clients (Figure 10b).

7.4 FITT: Mixed Interest Attack



Simulation Settings: Using the same simulation settings as Figure 10, attackers send both fake and valid Interests after the DDoS starts.

Figure 11: FITT: Mixed Interest Attack

Figure 11 shows how FITT handle mixed Interest attack where attackers will send both fake and valid Interests towards the server. One obvious difference from the fake and valid Interest attack scenarios is that, after the attack starts, FITT will limit the traffic to be much lower than the black line (the tolerance plus the capacity plus the good Interests). This is because the edge routers will take the smaller value from the limits for fake Interest attack reaction and valid Interest attack reaction. Similar as the valid Interest attack scenario, after a short time period, FITT will place the limit to misbehaving clients only and the legitimate clients will recover. As shown, in the end, FITT will only pass legitimate traffic to the server (Figure 11b).

7.5 FITT: Multiple DDoS to Different Prefixes

Previous plots indicate that FITT works well with all types of Interest attacks, keeping the traffic well below the thresholds when there is one victim that is under attack. FITT is also designed to be able to handle multiple DDoS attacks (i.e., attacks to different prefixes, started at different times, using different types of Interests) at the same time. To evaluate FITT’s performance in multiple-DDoS cases, we use valid Interest attack as an example. Using the same topology settings, we let half of the attackers attack prefix P starting from second 2 and another half attack the another server P' starting from second 4. Note that P and P' can be located on the same node or different nodes. We set the capacity to be 750 for both servers so that the attacking traffic (3000 Interests/s for each victim) will trigger the FITT reactions.




Figure 12: FITT: Fake Interest Attack with 2 Victims

As shown, when multiple FITT reactions take place, FITT can effectively control the DDoS traffic from both attacks at the same time. For each victim server, the incoming Interests are throttled in the similar way as that when there is only one victim server under attack. Two servers’ incoming traffic lines quickly go below the threshold after the attack started and soon merge the legitimate client traffic line and finally all the traffic received by the two servers are from legitimate clients.

7.6 FITT: Throttling Granularity

In DDoS mitigation, collateral damage may ruin the legitimate traffic sent from the compromised devices. As a real-world example, IoT devices compromised as part of the Mirai botnet will generate attacking traffic to the DDoS target, but these IoT devices still function as expected, e.g., a smart home device still communicate with the IoT service provider normally. FITT throttles Interest traffic at a granularity of the name prefix.

We evaluate FITT in terms of the granularity of the traffic throttling. We simulate that 60 clients are compromised to attack the P with 100 Interests/s sending rate; at the same time,

the clients keep the normal communication with another service provider P' at the reasonable rate of 20 Interests/s. In this case, the attacking traffic will overwhelm P but the legitimate traffic will not go beyond the capacity of P' . As shown in the Figure 12b, compared with the two-victim scenario where traffic to both P and P' will be throttled, FITT only squelches clients’ traffic under P while the traffic towards P' will not be affected.

8 Discussion

8.1 Authenticity of Victim’s FITT NACK

Authenticity of the feedback is of vital importance since FITT relies on victim’s explicit feedback to start a reaction. In cases when attackers send fake NACKs to abuse FITT to deny good services, the router can easily detects such attacks by comparing the reported prefix with the routing prefix of the interface where NACK was received. For example, in topology 5, when router $R1$ receives a NACK from the link that connects S , reporting that the prefix “/isp0/service” is under attack. Assuming in $R1$ ’s routing table only prefix “/univ1” is registered on the interface towards S , $R1$ can drop the NACK because “/isp0/service” cannot be matched. Another approach to ensure NACK’s authenticity is to let the victim server appends a digital signature to the NACK, which requires the pre-configured trust model and key distribution between the server and the network.

8.2 Misbehaving Routers in FITT

In FITT’s design, we don’t have the assumption that all the NDN routers are trustworthy and perform FITT properly. However, it is possible that illegitimate routers (e.g., free WiFi access point deployed by attackers) will not forward NACKs or not throttle the attacking traffic at all. To handle the evil routers, service providers should carefully set which nodes to be edge routers. For example, in the Figure 5, if $R4$ is not a trustworthy router (e.g., is not fully controlled by the service provider), the service provider can set the $R3$ as the FITT edge router. In this case, even though $R4$ does nothing to combat the DDoS attack, the $R3$ will throttle the traffic sent from $R4$, including both attacking traffic from $C3$ and legitimate traffic from $C4$. Therefore, misbehaving routers will not ruin FITT’s performance but degrade the granularity of the traffic throttling.

8.3 Comparison with Previous Solutions

As we discussed in Section 2, existing DDoS mitigation solutions for TCP/IP networks either (i) base DDoS traffic signature on matching criteria from network headers (e.g., IETF dots, Pushback, BGP FlowSpec, RTBH), or (ii) add additional features to the network architecture for finer flow clas-

sification (e.g., AITF, SIFF, StopIt, TVA). Due to strict layering, identifying specific application-level traffic flows is unattainable.

By contrast, NDN enables FITT to identify and throttle specific application-level DDoS traffic flows at fine granularity by leveraging NDN’s architectural features— (i) application-semantic meaningful names and (ii) stateful forwarding, victim is able to report the DDoS at name prefix granularity, and FITT is able to identify exact application-level attack traffic flows and attackers at the network level (as shown in the simulation scenario in 7.6).

There have also been various proposed approaches to mitigate Interest DDoS over NDN/ICN. Specifically, [1, 15] leverage the “success ratio” (how many Interests get satisfied by Data) to detect a presence of fake Interest DDoS. [16, 34] propose to detect Interest flooding by monitoring the PIT size or PIT utilization rate. They mainly focus on one specific type of attack – I-2 Interest attack (i.e. Interests carrying false data names). To be effective, they also require routers must be able to set proper threshold values for the detection function, and these threshold values can be non-trivial to configure when underlying traffic composition is complex. In comparison, FITT directly takes input from attack targets and can handle valid Interest (i.e. I-1 and I-3 Interests towards existent or dynamically generated data) DDoS and mixed Interest attack scenarios where attackers can send both fake and valid Interests towards the target. The explicit feedback from the victim enables accurate traffic throttling and reinforcement rate limiting to misbehaving consumers, removing the need to configure proper threshold values.

9 Conclusion

DDoS attacks have been an asymmetric threat since they first became significant, roughly 20 years ago. In that time, we have seen an increasing trend of application-specific TTPs, most clearly visible in reflective amplification attacks that use application semantics. As the IP network architecture is being abused to facilitate distributed attacks, our mitigation techniques and defenses have struggled with fundamental misalignments between the essential functions and forwarding semantics needed for effective mitigation and IP’s stateless forwarding, to address and remediate the traffic from large DDoS attacks. Although many solutions have been proposed to add those missing components, misalignments of incentives make their rollout difficult. Consequently, the current mitigation techniques necessarily backhaul DDoS attack traffic across the Internet to centralized mitigation servers that do DPI. This results in congesting links, costing operational overhead, and framing an unmaintainable capacity mismatch (in which transit capacity of centralized mitigation must match the aggregate DDoS traffic from distributed attack sources). The ability to remediate DDoS attacks close

to their sources has been among the goals that have long been sought after.

In this work, we have illustrated how an architectural change can lead to effective solutions to DDoS mitigation, a challenge that has faced Internet for at least 20 years. Our solution, called FITT, utilizes NDN’s basic properties and fundamentally (and architecturally) addresses the TTPs of the Internet’s largest DDoS attacks, like those from the Mirai botnet. By utilizing NDN’s stateful forwarding and structured names, FITT is capable of actively responding to all three types of Interest flooding. Evidence from our effort to mitigate the Interest flooding suggests that NDN provides a solid foundation for DDoS defense. As next step, we plan to add the FITT forwarding strategy into the NDN forwarding module implementation [35, 40] and perform a live deployment of FITT over both the NDN testbed [30] and in NDN’s experimental IoT edge networks.

What’s more, our analysis also shows that NDN’s architectural design is incrementally deployable today, allows independent edge-in deployments to natively work together, leading toward inherent DDoS resilience. The incremental deployment path for NDN with FITT could reasonably be expected to put DDoS defense mechanisms directly in front of major sources of DDoS, and would let applications directly instruct the network to squelch DDoS traffic at its sources.

References

- [1] AFANASYEV, A., MAHADEVAN, P., MOISEENKO, I., UZUN, E., AND ZHANG, L. Interest flooding attack and countermeasures in named data networking. In *IFIP Networking Conference, 2013* (2013), IEEE, pp. 1–9.
- [2] AKAMAI. Akamai ddos protection, 2019. Online; Available at <https://www.akamai.com/us/en/resources/ddos-protection.jsp>.
- [3] AMADEO, M., CAMPOLO, C., IERA, A., AND MOLINARO, A. Named data networking for iot: An architectural perspective. In *2014 European Conference on Networks and Communications (EuCNC)* (June 2014), pp. 1–5.
- [4] AMADEO, M., CAMPOLO, C., IERA, A., AND MOLINARO, A. Information centric networking in iot scenarios: The case of a smart home. In *Communications (ICC), 2015 IEEE International Conference on* (2015), IEEE, pp. 648–653.
- [5] ANDERSON, T., PETERSON, L., SHENKER, S., AND TURNER, J. Overcoming the internet impasse through virtualization. *Computer* 38, 4 (April 2005), 34–41.

- [6] ANTONAKAKIS, M., APRIL, T., BAILEY, M., BERNHARD, M., BURSZTEIN, E., COCHRAN, J., DURUMERIC, Z., HALDERMAN, J. A., INVERNIZZI, L., KALLITSIS, M., ET AL. Understanding the mirai botnet. In *USENIX Security Symposium* (2017).
- [7] ANTONAKAKIS, M., APRIL, T., BAILEY, M., BERNHARD, M., BURSZTEIN, E., COCHRAN, J., DURUMERIC, Z., HALDERMAN, J. A., INVERNIZZI, L., KALLITSIS, M., ET AL. Understanding the mirai botnet. In *USENIX Security Symposium* (2017).
- [8] ARGYRAKI, K. J., AND CHERITON, D. R. Active internet traffic filtering: Real-time response to denial-of-service attacks. In *USENIX annual technical conference, general track* (2005), pp. 135–148.
- [9] BACCELLI, E., MEHLIS, C., HAHM, O., SCHMIDT, T. C., AND WÄHLISCH, M. Information centric networking in the iot: Experiments with ndn in the wild. In *Proceedings of the 1st ACM Conference on Information-Centric Networking* (New York, NY, USA, 2014), ACM-ICN '14, ACM, pp. 77–86.
- [10] BAKER, F., AND SAVOLA, P. Ingress filtering for multihomed networks. BCP 84, RFC Editor, March 2004. <https://tools.ietf.org/html/bcp84>.
- [11] BALLANI, H., CHAWATHE, Y., RATNASAMY, S., ROSCOE, T., AND SHENKER, S. Off by default!
- [12] BLOOM, B. H. Space/time trade-offs in hash coding with allowable errors. *Communications of the ACM* 13, 7 (1970), 422–426.
- [13] CERT COORDINATION CENTER. Cert incident note in-99-04, 1999. Available at https://web.archive.org/web/20081115163511/http://www.cert.org/incident_notes/IN-99-04.html.
- [14] CLOUDFLARE. Cloudflare advanced ddos attack protection, 2019. Online; Available at <https://www.cloudflare.com/ddos/>.
- [15] COMPAGNO, A., CONTI, M., GASTI, P., AND TSUDIK, G. Poseidon: Mitigating interest flooding ddos attacks in named data networking. In *Local Computer Networks (LCN), 2013 IEEE 38th Conference on* (2013), IEEE, pp. 630–638.
- [16] DAI, H., WANG, Y., FAN, J., AND LIU, B. Mitigate ddos attacks in ndn by interest traceback. In *Computer Communications Workshops (INFOCOM WKSHPS), 2013 IEEE Conference on* (2013), IEEE, pp. 381–386.
- [17] FERGUSON, P., AND SENIE, D. Network ingress filtering: Defeating denial of service attacks which employ ip source address spoofing. BCP 38, RFC Editor, May 2000. <http://www.rfc-editor.org/rfc/rfc2827.txt>.
- [18] GASTI, P., TSUDIK, G., UZUN, E., AND ZHANG, L. Dos and ddos in named data networking. In *2013 22nd International Conference on Computer Communication and Networks (ICCCN)* (July 2013).
- [19] HANDLEY, M., AND GREENHALGH, A. Steps towards a dos-resistant internet architecture. In *Proceedings of the ACM SIGCOMM workshop on Future directions in network architecture* (2004), ACM, pp. 49–56.
- [20] HEER, T., GARCIA-MORCHON, O., HUMMEN, R., KEOH, S. L., KUMAR, S. S., AND WEHRLE, K. Security challenges in the ip-based internet of things. *Wireless Personal Communications* 61, 3 (Dec 2011), 527–542.
- [21] IOANNIDIS, J., AND BELLOVIN, S. M. Implementing pushback: Router-based defense against ddos attacks. In *NDSS* (2002), vol. 2.
- [22] JIN, C., WANG, H., AND SHIN, K. G. Hop-count filtering: an effective defense against spoofed ddos traffic. In *Proceedings of the 10th ACM conference on Computer and communications security* (2003), ACM, pp. 30–41.
- [23] KUMARI, W., AND MCPHERSON, D. Remote triggered black hole filtering with unicast reverse path forwarding (urpf). RFC 5635, August 2009.
- [24] LANGLEY, A., RIDDOCH, A., WILK, A., VICENTE, A., KRASIC, C., ZHANG, D., YANG, F., KOURANOV, F., SWETT, I., IYENGAR, J., ET AL. The quic transport protocol: Design and internet-scale deployment. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication* (2017), ACM, pp. 183–196.
- [25] LIU, X., LI, A., YANG, X., AND WETHERALL, D. Passport: Secure and adoptable source authentication. In *NSDI* (2008), vol. 8, pp. 365–378.
- [26] LIU, X., YANG, X., AND LU, Y. To filter or to authorize: Network-layer dos defense against multimillion-node botnets. In *ACM SIGCOMM Computer Communication Review* (2008), vol. 38, ACM, pp. 195–206.
- [27] MARQUES, P., SHETH, N., RASZUK, R., GREENE, B., MAUCH, J., AND MCPHERSON, D. Dissemination of flow specification rules. RFC 5575, August 2009.

- [28] MASTORAKIS, S., AFANASYEV, A., AND ZHANG, L. On the evolution of ndnSIM: an open-source simulator for NDN experimentation. *ACM Computer Communication Review* (July 2017).
- [29] MORTENSEN, A., ANDREASEN, F., ET AL. Distributed-denial-of-service open threat signaling (dots) architecture. Internet-Draft draft-ietf-dots-architecture-10, IETF Secretariat, December 2018.
- [30] NDN RESEARCHERS. Ndn testbed, 2019. Available at <https://named-data.net/ndn-testbed/>.
- [31] NEUSTAR. Neustar defense and performance, 2019. Online; Available at <https://www.security.neustar/digital-defense/ddos-protection>.
- [32] PSARAS, I., CHAI, W. K., AND PAVLOU, G. Probabilistic in-network caching for information-centric networks. In *Proceedings of the second edition of the ICN workshop on Information-centric networking* (2012), ACM, pp. 55–60.
- [33] ROSSOW, C. Amplification hell: Revisiting network protocols for ddos abuse. In *NDSS* (2014).
- [34] SALAH, H., AND STRUFE, T. Evaluating and mitigating a collusive version of the interest flooding attack in ndn. In *Computers and Communication (ISCC), 2016 IEEE Symposium on* (2016), IEEE, pp. 938–945.
- [35] SHI, J., AFANASYEV, A., ET AL. Named data networking forwarding daemon, 2019. Online; Available at <https://github.com/named-data/NFD>.
- [36] YAAR, A., PERRIG, A., AND SONG, D. Siff: A stateless internet flow filter to mitigate ddos flooding attacks. In *Security and Privacy, 2004. Proceedings. 2004 IEEE Symposium on* (2004), IEEE, pp. 130–143.
- [37] YANG, X., WETHERALL, D., AND ANDERSON, T. A dos-limiting network architecture. In *ACM SIGCOMM Computer Communication Review* (2005), vol. 35, ACM, pp. 241–252.
- [38] YI, C., AFANASYEV, A., MOISEENKO, I., WANG, L., ZHANG, B., AND ZHANG, L. A case for stateful forwarding plane. *Computer Communications* 36, 7 (2013), 779–791.
- [39] ZHANG, L., AFANASYEV, A., BURKE, J., JACOBSON, V., CROWLEY, P., PAPADOPOULOS, C., WANG, L., ZHANG, B., ET AL. Named data networking. vol. 44, ACM, pp. 66–73.
- [40] ZHANG, Z., LU, E., ET AL. Ndn-lite library, 2019. Online; Available at <https://github.com/named-data-iot/ndn-lite>.
- [41] ZHANG, Z., YU, Y., ET AL. An overview of security support in named data networking. *IEEE Communications Magazine* 56, 11 (November 2018), 62–68.
- [42] ZHANG, Z.-K., CHO, M. C. Y., WANG, C.-W., HSU, C.-W., CHEN, C.-K., AND SHIEH, S. Iot security: ongoing challenges and research opportunities. In *Service-Oriented Computing and Applications (SOCA), 2014 IEEE 7th International Conference on* (2014), IEEE, pp. 230–234.