Security and Privacy Analysis of National Science Foundation Future Internet Architectures

Moreno Ambrosin, Alberto Compagno, Mauro Conti, *Senior Member, IEEE*, Cesar Ghali, and Gene Tsudik, *Fellow, IEEE*

***Abstract*—The Internet protocol (IP) is the lifeblood of the modern Internet. Its simplicity and universality have fueled the unprecedented and lasting global success of the current Internet. Nonetheless, some limitations of IP have been emerging in recent years. Furthermore, starting in mid-1990s, the advent of mobility, wirelessness, and the Web substantially shifted Internet usage and communication paradigms. This accentuated long-term concerns about the current Internet architecture and prompted interest in alternative designs. The U.S. National Science Foundation (NSF) has been one of the key supporters of efforts to design a set of candidate next-generation Internet architectures. As a prominent design requirement, NSF emphasized "security and privacy by design" in order to avoid the long and unhappy history of incremental patching and retrofitting that characterizes the current Internet architecture. To this end, as a result of a competitive process, four prominent research projects were funded by the NSF in 2010: nebula, named-data networking, MobilityFirst, and expressive Internet architecture. This paper provides a comprehensive and neutral analysis of salient security and privacy features (and issues) in these NSF-funded future Internet architectures. Prior surveys on future Internet architectures provide a limited, or even no, comparison on security and privacy features. In addition, this paper also compares the four candidate designs with the current IP-based architecture and discusses similarities, differences, and possible improvements.**

***Index Terms*—Network security, privacy, trust, future Internet architectures.**

I. Introduction

THE ORIGINAL Internet was intended to support thousands of users, mainly in North America, accessing shared resources via dumb terminals. Nowadays, the Internet connects over 3 billion of mobile and desktop devices with a variety of applications ranging from simple Web browsing to video conferencing and content distribution. These extreme changes in Internet usage accentuated limitations of the current IP-based architecture and prompted research into alternative internetworking architectures.

Manuscript received August 18, 2016; revised February 15, 2017, June 21, 2017, and October 15, 2017; accepted November 27, 2017. Date of publication January 25, 2018; date of current version May 22, 2018. *(Corresponding author: Alberto Compagno.)*

M. Ambrosin and M. Conti are with the Department of Mathematics, University of Padua, 35121 Padua, Italy (e-mail: ambrosin@math.unipd.it; conti@math.unipd.it).

A. Compagno is with the Department of Computer Science, University of Rome "La Sapienza," 00198 Rome, Italy (e-mail: compagno@di.uniroma1.it).

C. Ghali and G. Tsudik are with the Department of Computer Science, University of California at Irvine, Irvine, CA 92697-3435 USA (e-mail: cghali@uci.edu; gene.tsudik@uci.edu).

Digital Object Identifier 10.1109/COMST.2018.2798280

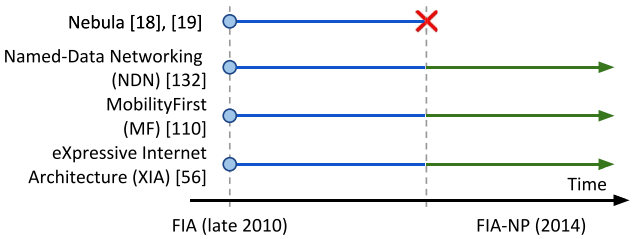

Fig. 1. Timeline of FIA & FIA-NP programs.

In 2010, the National Science Foundation (NSF) launched its "Future Internet Architecture" (FIA) program [95]. Originally, FIA was a 5-year program with the goal of designing a set of candidate next-generation Internet architectures. In 2015, NSF renewed its commitment with a follow-on "Future Internet Architecture – Next Phase" (FIA-NP) program. Unlike FIA which focused on architectural research, FIA-NP emphasizes evaluation, via prototypes, testbeds, trial deployments, and extensive experimentation.

FIA originally included four research projects: Nebula [18], [19], Named-Data Networking (NDN) [132], MobilityFirst (MF) [110], and eXpressive Internet Architecture (XIA) [56]. Each project focuses on a new Internet architecture with a distinct vision and design principles. Nebula envisions a highly-available and extensible core network interconnecting numerous data centers that enable new means of distributed communication and computing. NDN focuses on scalable and efficient data distribution – thus addressing inadequacies of the current Internet's host-centric design – by naming data instead of its location. MF concentrates on scalable and ubiquitous mobility and wireless connections. Meanwhile, XIA stresses flexibility and addresses the need to support different communication models by creating a single network that offers inherent support for communication between various principals (including hosts, content and services) while remaining extensible to future ones. Only three of the original four FIA architectures were selected for continued funding under FIA-NP: NDN, MF and XIA. Figure 1 illustrates the timeline of each project in FIA and FIA-NP programs.

Security and Privacy *by design* is one main goals of all FIA projects. It is mainly motivated by increasing reliance on Internet services, growing range and sophistication of

attacks,[1] and increasing demand for privacy. Given the rocky history of security and privacy in the current Internet, this goal is both very sensible and extremely important.

Contribution. In this paper, we survey and evaluate security and privacy features in the aforementioned four FIA projects. In doing so, we consider the network layer of the current Internet architecture as a point of reference. This is because the network layer reflects the most innovative choices and differences with respect to today's Internet. We also show how each FIA architecture succeeds or fails with respect to security and privacy features of current Internet's network layer, i.e., Internet Protocol (IP) [102] and IP Security Extensions (IPsec) [109]. We also discuss potential vulnerabilities that can be exploited to attack transmission channels, end-nodes, and the network infrastructure. Since different types of resolution services are needed in all FIA architectures, we compare security and privacy features of such services to those of Domain Name System (DNS) [88] and its prominent security extensions, such as the DNS Security Extensions (DNSSEC) [21], and DNSCurve protocol [44].

While we recognize that there are several future Internet architecture proposals outside the FIA program, such as DONA [66], PSIRP [47], NetInf [39] and SINET [131], we decided to focus our survey only on the four involved in the FIA project. We believe that this choice allows us to provide a fair and clear plane for comparison since all FIA projects share the same design principles, which includes security and privacy from the outset.

To the best of our knowledge, this paper represents the first comprehensive security and privacy treatment of four FIA architectures. Since it is impossible to predict which FIA architecture(s), if any, will ultimately succeed, we strive to remain neutral, i.e., to provide a complete and fair analysis.

Prior FIA surveys. An early article by Pan *et al.* [97] overviews Global Environment for Network Innovations (GENI). Unlike this paper, [97] provides a general overview and does not dwell on security and privacy aspects. The work in [64] provides a broad security analysis of the four NSF-founded FIA architectures, plus Recursive InterNetwork Architecture (RINA), Service Oriented Network Architecture (SONATE), and Netlet-based Node Architecture (NENA). The analysis in [64] considers four security features: confidentiality, integrity, availability and authentication. This paper provides a more in-depth security analysis and comparison of the four NSF-funded FIA architectures. In contrast with [64], it: (1) offers a thorough description of the said architectures and their resolution services; (2) treats a larger set of security features; and (3) discusses in detail security mechanisms at the network layer including the resolution services.

Other more focused surveys addresses security and privacy aspects of Information-Centric Networking (ICN) architectures [6], [7], [77], [120]; [77] analyzes security and privacy of NDN alone. Reference [6] investigates denial of service attacks in NDN, while [7] and [120] give a broader security analysis of several ICN architectures, considering several types of attacks on: naming, routing, and caching. Finally, [120] separately considers ICN security, privacy and access control.

Other, more general, surveys of ICN architectures do not focus on security and privacy aspects [14], [23], [121], [122]. Reference [14] analyze Data-Oriented Network Architecture (DONA), Named Data Networking (NDN), Publish-Subscribe Internet Routing Paradigm (PSIRP), and Network of Information (NetInf). The work concentrates on naming, routing and forwarding, caching, and mobility. It marginally considers security and privacy aspects. Reference [23] compares various naming and routing schemes in DONA, NetInf, PURSUIT and PSIRP architectures. References [121] and [122] compare mobility features of NDN, DONA, NetInf, and PURSUIT. None of them [23], [121], [122] discusses security and privacy in much detail.

Organization. We begin by overviewing IP, IPsec, DNS and its security extensions in Section II. Next, Sections III–VI summarize Nebula, NDN, MF and XIA, respectively. Section VII evaluates security and privacy features of these architectures, and compares them with those of IP and IPsec. Section VIII analyzes security and privacy of resolution services used by each new architecture. Section IX summarizes our comparative analysis, and highlights open issues, and possible future research directions. Finally, Section X concludes our paper.

Given familiarity with any FIA architectures, the corresponding sections, can be skipped without the loss of continuity.

II. The Internet of Today

Today's Internet architecture was designed over three decades ago to seamlessly inter-connect multiple heterogeneous networks. At the core of today's Internet is the TCP/IP protocol suite, which puts together protocols, applications and network mediums, and organizes them into four abstraction layers: Link, Internet, Transport and Application. This design leads to an hourglass shape with IP as the network layer as its "thin waist" [15].

We consider IP, which operates at the Internet layer, to be our point of reference when analyzing security and privacy of FIA architectures. IP is responsible for forwarding packets (a.k.a. datagrams) from the source IP interface to its destination counterpart. A host may have one or more IP interfaces, while a router has at least two. Each IP interface is identified by at least one distinct fixed-length IP address.

Another fundamental component of today's Internet architecture, and subject of our analysis, is DNS. DNS is a distributed service that translates application-specific domain names (specified in URLs) into their corresponding IP addresses, allowing hosts to communicate using meaningful names, rather than IP addresses.

In what follows, we briefly describe IP and DNS.

A. Internet Protocol

The cornerstone of IP is addressing of network devices. Every network-layer entity (router or host) is identified by at least one IP address which consists of a network prefix and a

[1] https://www.theguardian.com/technology/2016/oct/21/ddos-attack-dyn-internet-denial-service

host identifier. The boundary between them is flexible, which allows IP addressing to scale.

An IP datagram contains source and destination addresses along with other fields that convey control information. Actual data is carried in the payload field. When a packet is received, a router searches its Forwarding Information Base (FIB) to identify the next hop for that packet. A FIB contains a set of entries, each mapping one or more network prefixes to a router's interface and a next-hop IP address. This allows routers to perform longest-prefix matching on the destination address to identify the next hop. If a packet can not be forwarded, it is dropped and an error message is generated via the Internet Control Message Protocol (ICMP) [101].[2]

One important IPv4 feature is packet fragmentation. If the size of an IP packet is larger than the forwarding interface's Maximum Transmission Unit (MTU) [30], the packet must be divided into smaller chunks, called fragments. A destination host must reassemble fragments to recover the original packet. Other network entities, such as Network Address Translation tables (NATs) [55], [117] and firewalls *might* also assemble fragments.

As mentioned earlier, IP was originally designed for a small and relatively amicable research community. Neither its longevity nor its popularity was foreseen. Thus, it is unsurprising that IP lacks any security and privacy features. In late 1980-s and early 1990-s, as IP started to gain global popularity and the Internet transcended into the commercial sector, IPsec suite [109] was designed to provide basic security services, such as: origin authentication, data integrity, and confidentiality for IP datagrams. The first two are attained via Authentication Header (AH) protocol [62], while Encapsulation Security Payload (ESP) protocol [63] provides all three security features. IPsec supports two modes of operation:

- *Transport:* provides end-to-end communication, e.g., client-server communications. Only packet payloads are encrypted and authenticated in transport mode. Transport and application layers of packets are secured by a hash, thus, they can not be modified, e.g., using NAT. NAT-Traversal (NAT-T) [65] is developed to overcome this issue.
- *Tunnel:* typically used between gateways to provide a secure connection (pipe) between physically separate networks, e.g., different sites of the same organization. Tunnel mode also supports secure host-to-gateway communication. An IP packet is encrypted in its entirety and encapsulated as a datagram with a new outer IP header. One popular application of tunnel mode is Virtual Private Networks (VPN) [83].

IPv6 [42] is a newer version of IP developed to overcome some limitations of IPv4. One of its main new features is extended 128-bit address space (as opposed to 32 bits in IPv4). Another departure from IPv4 is lack of in-network fragmentation. Before sending an IP datagram must first discover the smallest MTU on the path to the destination and fragment the datagram accordingly. To help with this, the Path MTU Discovery protocol [85] was designed and implemented. IPv6 also takes into consideration security and privacy by implementing some features similar to IPsec – such as AH and ESP – as extension headers [1].

In the rest of this paper, we use the term "IP" to refer to both IPv4 and IPv6, unless otherwise specified.

[2]ICMP is also used for sending control messages, such as routing redirect for networks and hosts.

B. Domain Name System

The purpose of DNS is translation of domain names (e.g., those found in URL prefixes) into IP addresses. Domain names are organized in a hierarchical fashion: a top-level domain (e.g., "`.com`") is followed by many sub-level domains (e.g., second-level domain "`example.com`", and third-level domain "`sub.example.com`"). For each domain, DNS assigns an *authoritative name* server that stores information of, and responds to queries for, a specific contiguous portion of the domain name space, called *DNS zone*. This information is contained in *Resource Records* (RR-s) – basic DNS elements which are also carried in DNS replies. Moreover, authoritative name servers might delegate authority over sub-domains to other name servers, thus increasing DNS's scalability.

A user interacts with DNS by issuing a query to a local *resolver*: a process running on the end-user's device which forwards the query to the appropriate name server(s). The resolver sends a UDP (User Datagram Protocol [100]) packet containing the query to the DNS server, which is usually located in the resolver's local network. The server then checks if it can reply to the query from its cache. Otherwise, it fetches the response from other local or remote DNS servers.

DNS queries can be of two types: iterative or recursive. An iterative query allows a DNS server to return the best answer to the resolver, based on its local information, i.e., either a cached RR or an RR belonging to its zone. If the server does not have an exact match for the queried name, it returns a *referral*: a pointer to a DNS server authoritative for a lower level of the domain namespace. The resolver then queries the DNS server in the referral which can also reply with a referral. This process continues until the resolver receives requested information, or an error is generated. In the recursive query, DNS servers reply with either the requested RR or an error. If the DNS server does not have the requested information, it recursively queries other DNS servers.

Although DNS was originally designed as a static distributed database, it now allows dynamic records updates [29], [128] and zone transfers [69]. Also, a recent proposal envisions DNS as a distributed database to store IP related information [8]. For instance, [106] proposes storing IPsec keys related information in DNS records and mapping them to IP addresses.

The original DNS did not include any security or privacy features. DNS Security Extensions (DNSSEC) [21] was added to provide data integrity and origin authentication for DNS messages. In DNSSEC, RR-s are signed by their authoritative servers' keys. The basic mechanism and the query-response protocol of DNS remain unaltered.

There have been other attempts to secure DNS, e.g., DNSCurve [44]. It provides link-level security between clients and DNS servers, using elliptic-curve cryptography. DNSCurve guarantees hop-by-hop query/response confidentiality, authenticity and integrity.

III. NEBULA

Nebula [17]–[19] is a FIA project focused on providing a secure and cloud-oriented networking infrastructure. Its architecture is composed of three tiers:

- *Network core* (NCore) is a collection of routers and interconnections that provide reliable connectivity between routers and data centers. NCore is based on high-performance core routers and rich interconnected topologies [72].
- *Nebula Virtual and Extensible Networking Techniques* (NVENT) represents the control plane of Nebula. NVENT helps in establishing trustworthy routes based on policy routing [22] and service naming [94].
- *Nebula Data Plane* (NDP) is responsible for routing packets along the paths established by NVENT. To guarantee confidentiality, availability, and integrity, NDP ensures that packets for a specific communication can only be carried when all parties, i.e., end nodes and routers in between, have agreed to participate.

A. Nebula Network Layer

The original design of Nebula specifies different candidate network layer stacks for NDP [19], e.g., ICING [91], TorIP [73], and Transit-as-a-Service (TaaS) [99]. From this list, ICING was picked as the most suitable candidate and was included in the Zodiac Nebula prototype implementation [18].

ICING provides a new primitive, called *Path Verification Mechanism* (PVM), which guarantees the following two properties:

- *Path Consent* – every entity in a path between two hosts *consents* the use of the whole path before the communication starts.
- *Path Compliance* – the possibility for each node in a path between two hosts to verify that a received packet: (1) follows the approved path; and (2) has been "correctly" forwarded by all the previous nodes in the path, i.e., according to a specific pre-established policy.

ICING can be deployed either at the network layer or as an overlay on top of IP. In the former case, service providers can deploy ICING nodes as ingress gateways to their networks. However, in the latter case, ICING nodes may become *waypoints*, interconnected using IP, providing waypoint-level path guarantees.

To start communication, a sender must first establish a complete path. Such a path can be provided by DNS with policy enforcement [91]. Figures 2(a) and 2(b) show how forwarding works in ICING and a high-level representation of how the ICING header evolves.

Once a path is selected, the sender requests a *Proof of Consent* (PoC_j), for each node j on the path (action ① in Figure 2(a)). PoCs are cryptographic tokens created by each node transit provider, which attest to the provider's consent to carry packets along the specified path. Each PoC certifies that the corresponding network provider consents to (1) the full path, and (2) a specific policy-based set of local actions (e.g., forwarding) to be performed on packets traversing the path. PoCs are generated by a *consent server*, which is owned by the transit provider or acts on its behalf. Such servers share secret keys with each node (router) in their corresponding providers. Once all PoCs are received, the path is established and packet transmission can begin.

Each packet contains a header (shown in Figure 2(b)) including: (1) the path P consisting of all ICING nodes N_j forming it, and, (2) a list of verifiers V_j, one per node N_j in the path except the sender. This allows each verifier to prove that the packet passed through all previous nodes.

A sender builds a packet header as follows:

1) *Proof of Provenance* (PoP) token, one for each node on the path (action ② in Figure 2(b)), is generated using a PoP key k_j shared with the corresponding node j. In Figure 2(b), PoPs are denoted as $\text{PoP}_{i,j}$, where i is the index of the node generating the PoP and j is the index of the node for which PoP is generated. Specifically, $\text{PoP}_{0,j}$ is computed by node 0 using k_j, path P, and message M itself.
2) Authenticator A_j is computed for each node j using PoC_j, P and M.
3) Verifiers V_j, one per node, are computed by XORing the corresponding A_j and $\text{PoP}_{i,j}$.

PoP tokens are used by each node on the path to prove that downstream nodes have handled the received packets based on the established policies. When an intermediate node N_i receives a packet, it performs the following actions:

1) Computes the corresponding PoC_i.
2) Computes $\text{PoP}_{j,i}$ using k_j, for each downstream node N_j.
3) Verifies that the received PoC_i and $\text{PoP}_{j,i}$ match the two values computed in the previous two steps. If this verification fails, N_i drops the packet.
4) Derives a shared PoP key k_l, for each upstream node N_l, and computes $\text{PoP}_{i,l}$ as described above.
5) Modifies the verifiers to include the computed PoP, and forwards the packet upstream (actions ③ and ④ in Figures 2(a) and 2(b)).

The previous steps allow any node to guarantee that all packets are forwarded by all the consenting nodes while establishing the path.

B. Nebula Control Plane

The control plane in Nebula is provided by NVENT. NVENT uses declarative networking [75], [76], and allows administrators to provide high-level specifications of their routing policies. NVENT also involves special interfaces, called *service interfaces*, that enable service access and specify the required level of availability. For instance, an emergency service can request high availability, which can be provided by multi-path interdomain routing. A distributed resolution service is used for discovery of other NVENT services. This

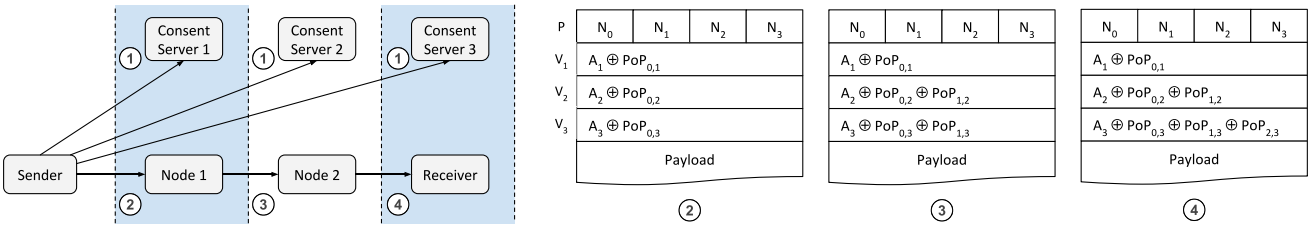

(a) Packet forwarding in ICING

(b) ICING packet high-level structure

Fig. 2. ICIN [91] architecture.

service is populated by service providers, e.g., NCore data centers [17].

Serval is an implementation of NVENT based on the concept of service-centric networking [53], [86], which decouples service instances (e.g., Web or email services) from their physical locations (i.e., IP address and port). Serval introduces a new layer, the Service Access Layer (SAL), between the network layer and the transport layer. With Serval, each service is identified by a *serviceID*, a unique identifier that applications use to communicate with the service. In addition, each local traffic flow, representing a connection between two hosts, is identified by a unique *flowID*. The request is handled by SAL, which uses local control plane policies to map the *serviceID* to a service instance. SAL eventually creates a new *flowID* that identifies the established connection. This *flowID* is delivered to the destination host during connection setup, and used by both parties for connection identification. Finally, SAL routes the packet based on specific control plane rules contained in its *SAL table*. For instance, a host application that wants to connect to a specific service might direct the first request to a default Serval router (using its IP address). The SAL of the router then processes the request and take further decisions based on its *SAL table* (e.g., forward to another router or send directly to a known service instance). Moreover, Serval does not directly provide clients a way to learn *serviceID*s: It simply suggests the use of directory services or search engines [94].

Figure 3 presents a view of how all Nebula components integrate to allow a user to negotiate a custom end-to-end path to a specific data center and send the desired packets. First, the user (either the mobile phone or the laptop in the figure) contacts NVENT to request a path to NCore. NVENT determines a suitable path that complies with each transit network's policies and contacts the corresponding consent servers to obtain the necessary PoCs. Once the path and all PoCs are delivered to the user, the latter generates appropriate packet headers and forwards them, using the NDP forwarders network, to the nearest NCore router. This router ensures that all header fields are valid (as described above) and verifies that the negotiated path has actually been traversed. Once verified, the core router forwards received packets to the correct data center using its NCore links.

IV. Named-Data Networking

While IP traffic consists of packets sent between communicating end-points, NDN traffic is comprised of explicit

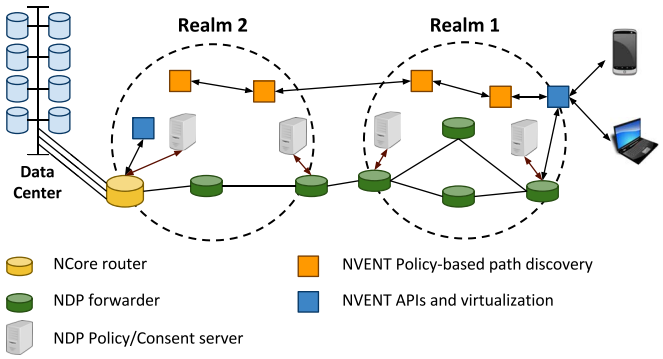

Fig. 3. High-level view of Nebula components integration [17]–[19].

requests for, and responses to, named content objects. NDN is based on the principle of Content-Centric Networking, where content, rather than hosts, occupies the central role in the architecture. NDN is primarily oriented towards efficient large-scale content distribution. Rather than directly addressing specific hosts, NDN users (called consumers) request pieces of content by name. The network is in charge of finding the closest copy of the content, and delivering it. This decoupling of content and location allows NDN to efficiently implement multicast, content replication and fault tolerance.

A. NDN Network Layer

The NDN network layer uses hierarchical structured names to directly address content. Names are composed of a number of components, e.g., `/ndn/bbc/frontpage/news` where "/" represents the boundary between two components. Since names are opaque to the network, they can contain binary or human-readable components.

To support content distribution, NDN defines two types of packets: interest and content (the latter is also called data packet). NDN communication adhered to the *pull* model, that is: every content is delivered to consumers only upon explicit request. Specifically, a consumer issues an *interest* packet carrying the name of the desired content. The network will then forward the interest towards the producer.

One important feature of NDN is in-network caching: any router can store a copy of the content it receives or forwards, and use it to satisfy subsequent interests. Therefore, an NDN interest might be satisfied by the actual content producer or any intermediate router. Along with in-network caching, NDN introduces another important feature called interest collapsing:

only the first of multiple closely spaced (and timed) interests requesting the same content is forwarded by each router.

Each NDN entity (not only routers) maintains the following three components [132]:

- *Content Store* (CS) – cache used for content caching and retrieval. A router's cache size is determined by local resource availability. Each router unilaterally determines what content to cache and for how long. From here on, we use the terms *CS* and *cache* interchangeably.
- *Forwarding Interest Base* (FIB) – table of name prefixes and corresponding outgoing interfaces. FIB is used to route interests based on longest-prefix matching of their names.
- *Pending Interest Table* (PIT) – table of outstanding (pending) interest names and a set of corresponding incoming interfaces, denoted as *arrival-interfaces*

When an NDN entity receives an interest, it searches its PIT to determine whether another interest for the same content is pending. There are three possible outcomes:

1) If a PIT entry for the same name exists, and the arrival interface of the present interest is already in *arrival-interfaces*, the interest is discarded.
2) If a PIT entry for the same name exists, yet the arrival interface is new, the router appends the new incoming interface to *arrival-interfaces*, and the interest is not forwarded further.
3) Otherwise, the router looks up its cache for a matching content. If it succeeds, the cached content is returned and no new PIT entry is needed. Conversely, if no matching content is found, the router creates a new PIT entry and forwards the interest using its FIB.

Upon receipt of the interest, the producer, or an intermediate router, responds with a matching content, thus *satisfying* the interest. The content is then forwarded towards the consumer, traversing the reverse path of the preceding interest. Each router on the path flushes the corresponding PIT entry and forwards the content out on all interfaces specified by that entry. If a content is received by a router with no prior matching interest, the content is considered unsolicited and is discarded. Since no additional information is needed to deliver content, interests do not carry any form of *source addresses*.

The last component at the end of content name can carry an implicit digest (hash) component of the content that is recomputed at every hop. This effectively provides each content with a unique name. Names carrying such digest forms what is called as Self-Certifying Names (SCNs). If an interest is issued using SCN, the retrieved content is guaranteed, due to longest-prefix matching, to be the same content requested by the consumer. However, in most cases, the hash component is not present in interest packets, since NDN does not provide any secure mechanism to learn a content hash *a priori*.

Apart from the name, content packets carry a `Signature` generated by the content producer and covering the entire content. For this reason, each producer is required to have at least one public key, represented as a *bona fide* named content object. Other notable fields in content packets are: the `Payload` containing the actual data of the content and the `ContentType` defining the type of the content, e.g., data or key. Other important fields in interest packets are: the `KeyLocator` which references to the public key required to verify the signature, and the `InterestLifetime` which specifies the lifetime of an interest before it expires (and its corresponding PIT entry is flushed).

Similar to IP, fragmentation of NDN packets can not be avoided. The fact that names can grow arbitrary long might cause interests length to span beyond some link MTU values. In this case, fragmentation must occur. However, since FIB forwarding is based on the availability of the entire name, reassembling of fragmented interests at every hop is a must. Furthermore, interest collapsing can cause content objects to be fragmented (or even re-fragmented) by routers [49]. The question remains to whether to perform a hop-by-hop reassembly [12], or cut-through processing of content fragments [49]. Regardless of its claimed benefits, it is trivial to see that hop-by-hop reassembly incurs unnecessary overhead and end-to-end latency.

Not all interests result in content being returned. If an interest encounters either: (1) a router that can not forward it further or (2) a producer that has no matching content, no error is generated. PIT entries in intervening routers simply expire when no matching content is received. In such case, the consumer can choose to re-issue the same interest after a timeout.

B. NDNS Distributed Database

Since content can be addressed using human-readable names, NDN in principle does not require a resolution service that translate user-friendly names into network addresses. However, as discussed in [11], a distributed database similar to DNS, if existed, provides several benefits to the NDN architecture:

- Cryptographic credential management: Since each data packet is required to be signed, a distributed database is optimal to store and serve security information (e.g., keys and certificates) for namespaces.
- Namespace regulation in the global routing: Similar to the ROVER project [48], a DNS-like service can store information that certifies the authorization of ASes to announce a particular prefix in the global routing.
- Scaling NDN routing: The fact that NDN names can be arbitrary long renders the namespace infinitely large. This exceeds the number of possible routable IP prefixes. Therefore, a DNS-like service can be used to implement a Map-n-encap solution to increase scalability in NDN routing [13].

Two distributed database systems that resemble the DNS design, KRS and NDNS, are proposed in [11] and [79], respectively. These two proposals adopt a similar design and provide the same features. In the rest of this paper we use NDNS to refer to such a distributed system.

Similar to domain names in DNS, NDNS organizes namespaces in a hierarchical set of zones and assigns replicated authoritative servers for each of them. NDNS queries are expressed via interest, in which, names carry all query's necessary information. NDNS responses are carried in content

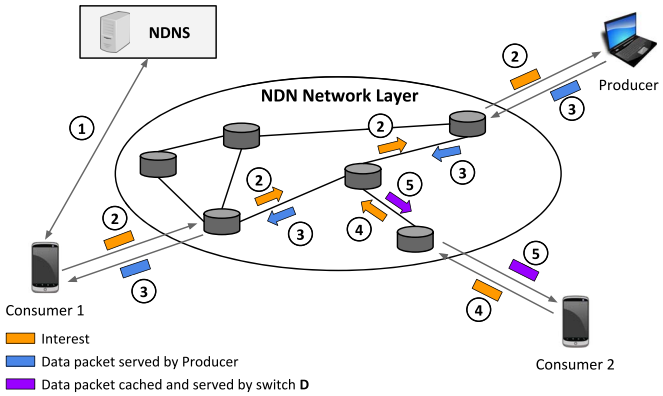

Fig. 4. NDN architecture and forwarding.

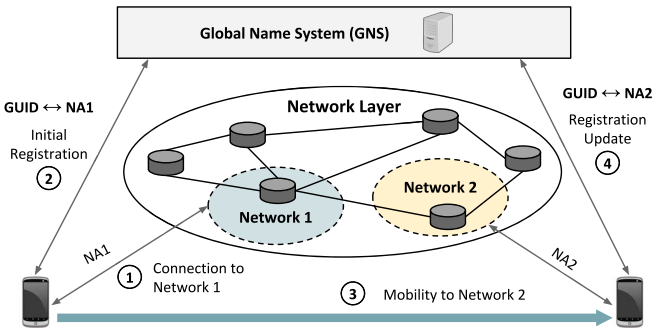

Fig. 5. MF architecture and GUID-NA mapping registration at GNS.

objects where their payloads contain the information requested by the corresponding query.

NDNS reflects many of the DNS protocol machinery: a resolver issues an iterative or recursive query to a local NDNS server. In case of iterative query, the server can reply with the answer (if known), a referral, or a negative response. In case of recursive query, if the NDNS server does not know the answer, it recursively queries another NDNS server until it receives an answer (i.e., the requested data or a negative response). Moreover, secure dynamic updates are provided as in DNS [126].

Figure 4 shows the basic dynamics of NDN naming resolution and forwarding. Consumer 1 retrieves a routable content name from NDNS (Step ① in Figure 4), then issues an interest, which is routed to its Producer (Step ② in Figure 4). The Producer, then responds with a data packet matching the interest, which is routed back to Consumer 1 on the reverse path (Step ③ in Figure 4). Subsequent interests for the same data packet (Step ④ in Figure 4) may be satisfied by intermediate network caches (Step ⑤ in Figure 4).

V. MobilityFirst

MobilityFirst (MF) architecture aims to overcome the inefficiencies and limitations of today's Internet due to mobility. It focuses on scenarios where wireless connections are *ubiquitous* and *pervasive*. To this end, MF has been designed around the concepts of *mobility* and *trustworthiness*. All endpoints must be able to seamlessly switch network connection, and the network must be resilient to compromised endpoints and routers.

MF treats *principals* – devices, content, interfaces, services, human end-users, or a collection of identifiers – as primary addressable network entities. To promote mobility, the (constant) identity of a principal and its (dynamic) network location are strictly separated. This requires a distributed Global Name Service (GNS) to bind principal identities to network addresses. Furthermore, identity and network address separation: (1) facilitates service implementation and deployment; and (2) supports designing routing protocols that overcome link fluctuation and disconnections [93].

We now briefly describe MF's network layer and its GNS.

A. Network Layer

Two types of identifiers are used to differentiate between principal identities and their physical locations.

- *Global Unique Identifier (GUID):* a flat self-certifying identifier that uniquely identifies a principal. GUIDs can be generated using multiple methods depending on the provided service type. For instance, they can be derived from the public key of a host or a service principal or the hash of a content principal. For the sake of usability, a human readable name can be assigned to a principal and later resolved (by GNS) to the corresponding GUID.
- *Network Address (NA):* a flat address that identifies a *network* to which a particular principal (GUID) is connected. MF networks are equivalent to ASes on today's Internet. NAs can be used to identify finer-grained networks such as subnets or organizations. In cases where principals are connected to multiple networks (e.g., using 3G and WiFi simultaneously), multiple NAs can correspond to the same principal.

As a consequence of this addressing scheme, MF defines a new packet type called Packet Data Unit (PDU). PDUs contain source and destination GUIDs, lists of source and destination NAs, payload, and other control fields.

Figure 5 shows a simplified architecture of MF, and registration and mobility handling. When connecting to a network with network address `NA1` (Step ① in Figure 5), a MF host first registers itself at GNS under a specific GUID (Step ② in Figure 5). When such host moves into a different network and obtains address `NA2` (Step ③ in Figure 5), it updates GNS database to reflect its new location (Step ④ in Figure 5).

In order to communicate with a specific GUID, endpoints need to query GNS to obtain the corresponding NA. The retrieved tuple (GUID, NA) is then carried in the PDU header as a routable destination identifier. PDUs are first delivered to their corresponding destination NAs (using inter-domain routing), and then to the destination GUIDs (using intra-domain routing). In case of delivery failure, the packet is stored inside the network (in routers) and GNS is periodically queried for a new or updated GUID-NA mapping.

Multihoming, anycast, and multicast are supported by multicast GUIDs (MIDs). MID has the same format as a regular GUID, except its resolution results in a *set of NAs* (instead of, at most, one). Technically, GNS associates one MID with

several GUIDs (the ones belonging to the multicast group). Resolving all of them results in one or more elements of the output NAs set.

MF can also support content distribution networks. In this case, GUIDs are composed of two parts:

- Content GUID (CID): uniquely identifies the content and is generated by computing the hash of the corresponding content.
- Publisher GUID (PID): points to the network entity providing the content. Such an entity can be the actual content provider, or a third-party content repository.

A router may be equipped with a cache. This opportunistic caching feature facilitates content distribution at the network layer by reducing end-to-end latency and bandwidth consumption. Moreover, MF exploits in-network caching to implement a per-segment (i.e., a continuous set of links with caching routers at each end) reliable chunk (few hundred of megabytes) transfer. Each chunk is fragmented and transmitted according to the segment MTU. Then, the caching router at the other end assembles the entire chunk and stores it. In case of transferring failure, caching routers can re-transmit a chunk via the same, or even a different, path.

B. Global Name Service

GNS is an essential part of the MF architecture. Its main task is to map endpoint identifiers (GUIDs or human readable names) to a set of attributes including the endpoint network address. GNS relies on the following two services:

- *Name Certification Service (NCS):* is equivalent to a Certificate Authority (CA). Its purpose is to (1) assign GUIDs to human-readable names and (2) attest this mapping by generating certificates. MF allows multiple NCSs without a global root of trust. Moreover, if GUID space is large enough, the need for coordination between different NCSs is eliminated.
- *Global Name Resolution Service (GNRS):* a distributed naming service similar to Domain Name System (DNS) that stores the mapping between GUIDs and NAs [74], [89], [125].[3] Two GNRS implementations are evaluated: (1) a distributed hash table maintained among all ASes of the Internet (DMap [127]), and (2) a number of replica-controllers that migrate data (GUID-NA mappings) between a variable number of active replicas (Auspice [111]).

Regardless of its implementation, GNRS clients interact with the service by issuing the following requests to the GNRS resolver:

- `insert`: register a new GUID-NA mapping when a principal joins the network.
- `update`: keep the GUID-NA mapping up-to-date when the corresponding principal migrates to a new network location.
- `query`: retrieve the list of NAs associated with a specific GUID.

[3]GNRS is the actual GNS service that is responsible for maintaining GUID-NA mappings.

In [74], a secure version of the above three GNRS request types is proposed. The secure `insert` and `update` requests adopt a two-step approach to check validity of a GUID-NA mapping. Four network entities are involved in this process: (1) the user issuing the new GUID-NA mapping, (2) the local router to which the user is connected, (3) the border gateway router that connects the user's AS to the rest of the Internet, and (4) the DHCP server which assigns the user's address.

The user generates and signs the request containing the GUID-NA mapping. Local and border routers are in charge of verifying validity of the announced mapping. This is achieved by verifying that the announced NA is the network connected to the user (and the local router), and querying the DHCP server to ensure that the returned NA corresponds to the announced GUID. If the NA matches the one contained in the `update` or the `insert` request, the mapping is accepted and added or updated in the GNRS table.

In the secure `query` request, the protocol involves three entities: the user, the border gateway, and GNRS. The user issues an authenticated request and the border router checks its validity. The router then forwards the request to the appropriate GNRS replica. On receipt, the GNRS satisfies the request with a signed GUID-NA mapping response.

There are several differences [111] between GNRS and DNS [88]. First, GNRS does not restrict the structure of the names, while DNS only supports hierarchical names. Second, scalability of GNRS does not rely on TTL-based caching, which has been proven to be ineffective in the presence of high mobility. Third, GNRS does not statically give the authority to a replicated server for a specific set of names. Active and on-demand replication reduce reliance on passive caching and ensure that mapping replicas are always accessible close to clients.

VI. Expressive Internet Architecture

eXpressive Internet Architecture (XIA) is another research effort aiming to design a new architecture. XIA is based on three types of principals. *Host*-centric networking can support end-to-end communication, such as video conferencing and file sharing. *Service*-centric networking allows users to access various network services such as printing and data storage services. Meanwhile, *content*-centric networking can support Web browsing and content distribution. However, XIA's design is extensible in that it can adaptively provide network evolution and support any new principal type that might emerge in the future.

A core architectural property of XIA is *intrinsic security* of all principals. Any entity should be able to authenticate the principal it is communicating with, without trusted third parties. This can be achieved by binding one or more security properties with principal names. For instance, using the hash of a service (or a host) public key as its name allows entities to verify that they are communicating with the desired principal. Similarly, binding content with its name can be achieved using the hash of the content as its name, allowing users to verify the integrity of a requested content.

XIA defines three main design requirements:

1) All network entities must be capable of clearly expressing their intent. This is achieved by designing the network to be *principal*-centric and allowing in-network optimization. Routers can perform principal-specific operations when receiving, processing, and forwarding packets.
2) The network must be able to adapt to new types of principals. This is essential to support network evolution.
3) Principal identifiers must be intrinsically secure. This depends on the principal type, e.g., authenticating hosts in *host*-centric networking is different than verifying content integrity in *content*-centric networking.

Principal identifiers are denoted as XID, where X defines the type of principle. For instance, HID identifies a host, SID a service, NID a network, and CID a content.

A. Expressive Internet Protocol

In order to comply with the aforementioned requirements, eXpressive Internet Protocol (XIP) is designed. XIP defines packet format, addressing schemes, and behavior of all nodes while processing incoming and outgoing packets from/to various principal types. One of the main features of the XIP addressing scheme is flexibility of defining multiple (fallback) paths to destinations. This prevents downtime and service interruption, especially while gradually deploying new principal types. An XIP address is a directed acyclic graph (DAG) with several properties:

- Each address is a single connected component.
- Each DAG starts with an untyped entry node and ends with one or multiple "sink" nodes. Thus, each node in the address graph has a unique XID except for the entry node.
- Edges define next hops in the path.
- Multiple outgoing edges of a single node are processed in the order they are listed.
- Out-degree of each node is upper bounded to restrict performance overhead.

Using DAGs as a basis for XIP addresses allows applications to build several "styles" of addresses, such as:

- *Shortcut routing* – this style, shown in Figure 6(a) is best suitable for requesting content principal. Each node has a direct edge to the destination principal CID_1, which enables in-network caching. If a node does not have the content cached, the fallback path is processed and the packet is forwarded to the next hop.
- *Binding* – some services require that communication is bound to a specific source or destination. For instance, a service hosted in multiple geographical locations. Users can establish a session with the closest host providing this service. Then, all further communications must be directed to this particular host. Figure 6(b) shows an example of this addressing style. The first packet is destined to SID_1, i.e., the closest host, while the second packet is destined to SID_1 provided by a specific host HID_1.
- *Infrastructure evolution* – as mentioned above, XIA supports gradual network evolution for emerging principal types using fallback paths. Figure 6(c) shows an example of this style. Assume that NID_1 is gradually deploying service SID_1. All NID_1 routers that are not yet updated to recognize and process SID_1 use the fallback path through HID_1 and HID_2.
- *Source routing* – Figure 6(d) gives an example of this addressing style, in which the source routes the packet to the destination through a third party domain and service, NID_a and SID_a, respectively.
- *Multiple paths* – this supports recovery from link failures. An example of this style is shown in Figure 6(e).

Figure 7 shows a high level overview of an XIA router. Its modular design allows efficient multi-principal processing and supports network evolution. Each router contains two main XID-specific processing modules:

- *Source XID-specific processing:* necessary for certain XID types. For instance, in case of a reply to a CID request, the "CID processing" unit can implement in-network content caching.
- *Next Destination XID-specific processing:* invoked by the *Next Destination XID-specific Classifier* which determines the appropriate forwarding action. Similar to source processing, this module consists of several units that carry on XID-specific operations right before forwarding the packet.

If all outgoing DAG edges of a node lead to unrecognizable XIDs, the packet is dropped and an unreachable destination error is generated. It is the responsibility of user applications to provide appropriate fallback paths to avoid forwarding failures at any router. Usually fallback paths are built using well-supported principals, e.g., HID and NID.

B. Principals

As mentioned above, principals in XIA support emerging communication paradigm on the current Internet. When introducing a new principal, the following issues arise:

- What does it mean to communicate with a principal of this type?
- How is the principal's unique XID generated and how does it map to intrinsic security properties?
- What are the source and next destination XID-specific processing actions that routers should perform and how can such actions be implemented?

We now describe several principal types and discuss their addressing schemes, in-router processing behaviors, and security properties.

1) Network and Host: Network and host principal identifiers are denoted as NID and HID, respectively. They are generated by using the public key hash of the network or the host. Unlike hosts on current Internet, each XIA host has a unique HID regardless of the interface it is communicating through. This feature helps support host mobility. In order to support fallback paths, all XIA routers should implement NID and HID processing modules.

As mentioned above, the fact that the network and host addresses are derived from their corresponding public keys allows users to verify the identity of entities

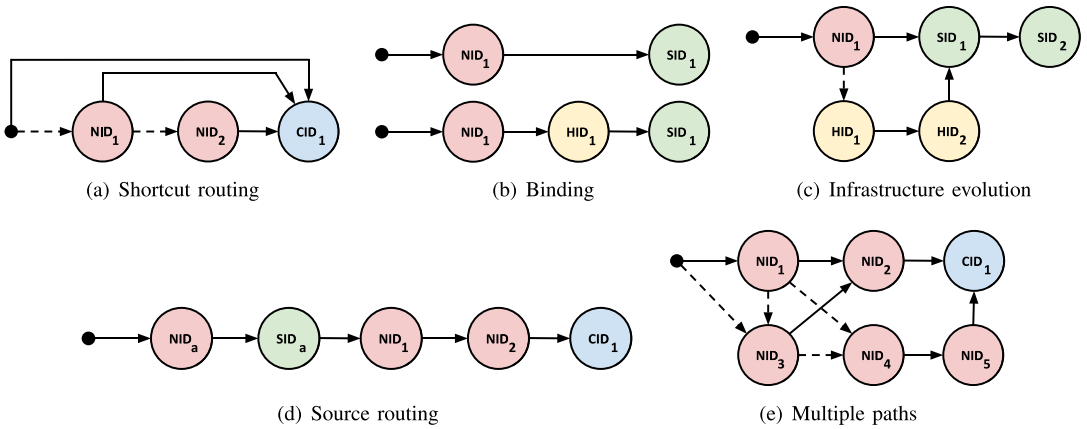

(a) Shortcut routing (b) Binding (c) Infrastructure evolution

(d) Source routing (e) Multiple paths

Fig. 6. XIP Addressing Styles [56].

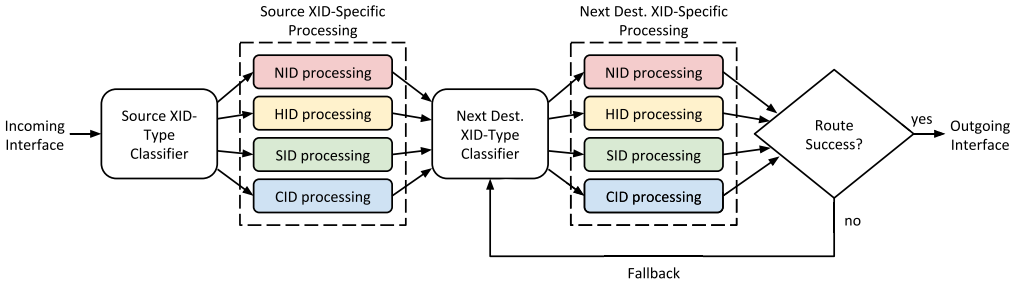

Fig. 7. XIA Router Diagram [56].

with whom they are communicating. Furthermore, this security requirement helps defend against address spoofing, Denial of Service (DoS), and cache poisoning attacks.

2) Service: Services in XIA represent applications in today's Internet. Users communicating with a service SID can use a destination address of the form NID:HID:SID. In today's terminology, this is analogous to sending a packet to a specific host in a specific network and indicating the associated protocol and port number.

Since different services might require different specialized processing, implementing in-router source and next destination processing modules is a challenge. Therefore, routers are only required to perform default processing, routing, and forwarding of SID packets. All other specialized processing should be handled by end-nodes.

SIDs are generated by computing the hash of the service public key. This inherits security properties similar to NIDs and HIDs.

3) Content Principals: This principal type signifies user's intent to retrieve content. Packets carrying content identifiers (CID) as destination addresses will be routed all the way to the node hosting the content. Routers can use a cached version of the content as a reply to such packets. As mentioned above, caching is implemented by routers source XID-specific processing module.

CIDs are generated based on the cryptographic hash of the content they address. This binds the content to its name, forming a self-certifying name.

VII. Network-Layer Security and Privacy

As mentioned earlier, security and privacy by design is one of the key NSF-stipulated guidelines for all FIA projects. In this section, we provide a comparison between the security and privacy features offered at the network layer of each architecture introduced above, and compare them with IP/IPsec. We consider the following security and privacy features, which we consider to be essential [112].

- *Trust:* confidence (sometimes based on inconclusive evidence) (1) that an entity will behave as expected; and (2) that a content comes from a trusted source.
- *Data origin authentication:* corroboration that the source of the received data is as claimed.
- *Peer entity authentication:* corroboration that a peer entity in an association is the one claimed.
- *Data integrity:* assurance that data has not been tampered with, in any (unauthorized or accidental) manner.
- *Authorization and access control:* respectively: (1) a secure means for an entity to access (e.g., read, write, delete) some resource, and (2) protection of a resource against unauthorized access.
- *Accountability:* assurance that all actions can be securely traced to their source(s).
- *Availability:* accessibility and usability of a resource upon demand by an authorized entity.
- *Data confidentiality:* unavailability of data to unauthorized entities.
- *Traffic flow confidentiality:* a set of countermeasures to traffic analysis.

- *Anonymous communication:* inability to determine identities of communicating entities.

We consider the last three as instrumental to privacy at the network layer. In the rest of this section, we discuss each feature separately. Tables I and II summarize our comparison on security and privacy features, respectively.

A. Trust

IPsec defines trust as a one-way relationship between two or more entities (hosts or networks). This relationship is captured by a Security Association (SA). An SA can be viewed as a "contract" between the peers. It describes security services and contains security information needed to protect communication.

Entities involved in secure communication in IPsec establish SAs via the ISAKMP[4] [84] protocol and exchange necessary cryptographic material using the Internet Key Exchange (IKE) protocol [32]. Host authentication in ISAKMP and IKE can be achieved via either digital signatures, or pre-shared keys. Digital signatures require the use of certificates to bind entity identifiers to public keys. This implies the existence of a CA to create, revoke, and distribute certificates.

Nebula's ICING-based network layer defines trust orthogonally to IPsec: between a host and all nodes forwarding its packets. As described in Section III, a host agrees on a "contract" with the network providers carrying its data (i.e., the path negotiated using NVENT) to specify the operation executed at each hop. Such contracts are cryptographically enforced. ICING assumes mutual trust between forwarding nodes and consent servers that are responsible for creating PoC tokens. Therefore, this notion of trust does not require any PKI [91]. However, ICING does not provide an end-to-end definition of trust, which can be added by adopting an IPsec-like approach.

Unlike IP, the notion of trust in NDN is directly associated with contents and not with hosts and networks. Trust in content can be expressed at different levels of granularity: from a single content object to an entire namespace. Recall that a content object is signed by its producer, which allows anyone to verify its origin and authenticity. Origin refers to the content producer and not to any entity that might store a copy of that content. To authenticate a content and its origin, its signature must be verified. This requires the verifier to retrieve, and establish trust in, the corresponding public key.[5] However, network-layer trust management is unspecified and relegated to individual applications. A detailed discussion of this topic can be found in [51].

MF places trust in the principal. Depending on the principal type, trust may be established with: (1) hosts, similar to IPsec, (2) content, similar to NDN, and (3) centralized or distributed services. Trust semantics in XIA also vary depending on the principal type. However, the intrinsic security feature of these principals (described in Section VI) increases trustworthiness of end-to-end communication and content retrieval. For instance, ensuring that a content hash matches its identifier allows receivers (and caching routers) to trust that content.

As shown in Section VI, an XIA address consists of a DAG containing a (partial) path to the destination. To provide trusted path selection for host-to-host communication, SCION is integrated with XIA [92]. SCION [133] is an architecture that provides control and isolation for secure and highly available end-to-end communication. The network is divided into multiple trust domains consisting of several Autonomous Systems (ASes) that trust each other. Each domain has a trusted root AS responsible for relaying packets to and from other domains. Roots initiate path establishment to all hosts in their domains based on local policies and available bandwidth. This process results in constructing a path between each host and its domain root. Whenever two XIA hosts, in different domains, want to communicate, the two half paths (from each host to its domain root) are combined to establish a complete end-to-end path. Such path is trusted since it is created by the trusted roots of each domain.

B. Data Origin Authentication

IP (IPv4 in particular) does not provide any form of authentication. A separate add-on method, IPsec, provides entity authentication via AH and ESP protocols.[6] In transport mode, two hosts securely negotiate a shared secret key. This key is later used to generate a Message Authentication Code (MAC) [68] for each packet. Successful MAC verification ensures authenticity of received packets and their origin. In case of gateway-to-gateway communication, gateways can only verify that the received data originated by *any* (not a specific) host connected to the network at the other end of the tunnel. In host-to-gateway communication, the gateway can actually verify that the data originated by the involved host, while the latter can only verify that received data is originated by the network located behind the gateway. This partial authentication opens the door for insider attacks.

Nebula's ICING-based network layer does not directly provide data origin authentication, which is delegated to applications.

NDN provides data origin authentication via content signatures. Before consuming content, consumers are required to verify its signature [132]. However, this operation is optional for routers because signature verification is an expensive operation at line speed and comprehensive trust management is not viable at the network layer. Even if we assume that routers know all possible application trust models, establishing trust in content is complicated and expensive. For instance, traversing a PKI hierarchy requires routers to fetch and verify public key certificates until a trusted anchor is reached.

On the other hand, NDN interests can optionally be authenticated using digital signature [3]. In a signed interest, the last component of the name carries a signature computed by the consumer. In case the interest carries a small payload,[7] the

[4]ISAKMP: Internet Security Association and Key Management Protocol.

[5]Public keys in NDN are distributed as special content objects with type `KEY`. An object of this type is signed by its issuer, i.e., a CA.

[6]Recall that IPv6 implements both AH and ESP as extension headers.

[7]An interest can carry a small payload to minimize delay in a communications (e.g., 0.5 RTT rather than 1 RTT to retrieve the data), as well as to support push-model communication of small data that fits in a single packet.

TABLE I
NETWORK SECURITY FEATURES COMPARISON

		IP/IPsec	Nebula	NDN	MF	XIA
Trust	*Trusted Model*	Hierarchical.	Hierarchical.	Hierarchical.	Hierarchical.	Distributed.
	Trusted Entity	Host.	Forwarding nodes.	Content.	Host, Content, Service.	Host, Content, Service.
	Trust Management	PKI + certificates, shared keys.	PoC tokens.	Certificates.	Distributed architecture + certificates.	Undefined.
Data Origin Auth.	*Coverage*	Gateway-to-Gateway, Host-to-Gateway.	None.	Producer-to-Consumer.	None.	None.
	Type of packet	Every that belongs to an IPsec SA.	None.	Content, Interest if signed.	None.	None.
	Mechanism	HMAC.	None.	Signature.	None.	None.
Peer Auth.	*Authenticated Entity*	Host, Gateway.	Consent Server.	None.	None.	None.
	Mechanism	Signature, Shared key.	PoC.	None.	None.	None.
Integrity	*Coverage*	Gateway-to-Gateway, Host-to-Gateway.	Host-to-Forwarding nodes.	Producer/Cache-to-Consumer.	Producer/Cache-to-Consumer.	Producer/Cache-to-Consumer.
	Type of packet	Every.	Every.	Content, Interest if signed.	Content.	Content.
	Mechanism	IPsec tunnel + HMAC.	Hash.	Signature.	Hash.	Hash.
Authz. & Access Cont.	*Enforced on*	Router, Host.	Consent Server.	Router, Producer.	Router, Host.	Router, Host.
	Mechanism	ACL.	ACL.	ACL, Encryption.	ACL.	ACL.
Accountability	*Entity accounted*	Host.	Host.	Producer, Consumer.	None.	None.
	Mechanism	Egress Filter on router, IPsec channel.	Path Consent.	Signature in Content or Interest packets.	None.	None.
Availability	*Bandwidth Depletion Attacks*	No architectural design countermeasures	Path consent can block attacks. Consent Servers are Single Point Of Failure.	Pull model prevents flooding with content. Weak to interest flooding.	No architectural design countermeasures.	No architectural design countermeasures.
	Routers Resource Exhaustionn	Weak to fragmentation attack on NAT-enabled IP routers.	Weak to attack targeting path consent verification on routers.	Weak to Interest Flooding attack on PIT.	None.	None.
	Cache-Related Attacks	None.	None.	Weak to Content poisoning. IKB rules as countermeasure to content poisoning. Weak to cache pollution.	SCN as countermeasure to Content poisoning. Weak to cache pollution.	SCN as countermeasure to Content poisoning. Weak to cache pollution.

signature will authenticate the consumer generating it, thus providing data origin authentication for interest payload.

MF and XIA do not provide any data origin authentication at the network layer.

C. Peer Entity Authentication

IPsec provides peer entity authentication during SA establishment of a secure communication. ISAKMP and IKE, the IPsec's protocols used to establish SAs, can achieve peer entity authentication using digital signature or pre-shared key. Digital signatures requires the use of certificates, which bind entity identities to their public keys. The use of certificates implies the existence of a trusted third party or a CA to create, sign and properly distribute certificates. Pre-shared keys on the other hand requires the communicating parties to agree on the shared secret key before communication begins.

TABLE II
NETWORK PRIVACY FEATURES COMPARISON

		IP/IPsec	Nebula	NDN	MF	XIA
Data Conf.	*Coverage*	Gateway-to-Gateway, Host-to-Gateway.	Host-to-Host.	Producer-to-Consumer.	Host-to-Host.	Host-to-Host.
	Mechanism	IPsec ESP header (payload encryption).	ICING + end-to-end encryption.	Payload encryption. Human-readable name might leak information.	End-to-end encryption.	End-to-end encryption.
Traffic flow Conf.	*Coverage*	Gateway-to-Gateway, Host-to-Gateway.	None.	Producer-to-Consumer.	None.	None.
	Mechanism	Padding (no mandatory on IPsec specification).	None.	None.	None.	None.
Anonymous comm.	*Level of anonymity*	Partial using IPsec in tunnel mode.	None.	Partial, consumer has no address. Router's state can be used to de-anonymize consumers.	None.	None.
	Additional Mechanism to achieve full anonymity	Tor.	TorIP.	Andana.	Tor.	Tor.

In Nebula, ICING allows a sender to authenticate the entities issuing the necessary cryptographic tokens, i.e., PoCs. However, ICING design does not specify how PoCs are retrieved, nor does it specify how entities are authenticated [91].

At its current state, NDN does not provide peer entity authentication for consumers and producers. However, in case the authentication of one or both entities is necessary, applications can exploit some features provided by the network layer. Considering consumers, signed interests can facilitate their authentication. Whereas for producers, content signature can ensure that the content is generated by the expected producer. Moreover, if interests must be satisfied by producers only (and not in-network caches), they should carry unique names that avoid cache hits and guarantee their delivery to corresponding producers.

Similar to NDN, MF and XIA do not provide peer entity authentication. However, the usage of self-certifying identities as principal identifiers facilitates entity authentication. Recall that for host, network, and service principals, identifiers are generated by computing the hash of the public key associated with these principals. Therefore, entity authentication can be achieved by ensuring that such principal identifiers match their keys. Peer authentication for content principals can be achieved similar to NDN since such principals are self-authenticating. Neither MF nor XIA provides a secure mechanism for securely retrieving content identifiers.

D. Data Integrity

Although IPv4 header contains the *Header Checksum* field that provides transmission error detection (a form of integrity check), it does not prevent packet manipulation. In fact, both versions of IP, introduced in Section II-A, completely delegate integrity to IPsec AH and ESP protocols. Specifically, the HMAC values in these protocol headers are used to achieve integrity. Depending on the IPsec mode used, host-to-host, gateway-to-gateway, or host-to-gateway integrity guarantees can be provided by both AH and ESP protocols. However, AH provides integrity for the entire packet (except for mutable fields), while ESP guarantees packet headers integrity only.

Each packet in Nebula carries a sequence of cryptographic verifiers V_j, one for each hop on the path (see Section III-A for details). The packet hash is used as part of V_j's calculation. Therefore, ICING guarantees that neither the packet nor the path can be modified. Also, ICING is recommended only at domain gateways [91]. Thus, integrity can only be guaranteed by border routers. Within domains, such guarantees are deferred to either the network-layer protocol or the application.

The way NDN provides integrity is through content signature. By verifying this signature, consumers and routers can always detect malicious manipulation. However, when content is requested using SCNs, data integrity is achieved by comparing the content hash to the last name component of its name. Furthermore, only signed interests can provide interest integrity.

In both MF and XIA, integrity is only available for content principal types. This is again due to the fact that such a principal identifier is generated based on the content hash itself. Whenever a content is received, its hash is compared with its identifier to ensure content integrity. For other principal types, MF and XIA defer integrity guarantees to the application.

E. Authorization and Access Control

Access control in IP is achieved by restricting access based on source and destination addresses. This is implemented using Access Control Lists (ACLs) [31], which contain a set of rules that grant or deny access to network resources. When implemented in routers, ACLs specify whether a received packet will be forwarded further to the next hop, or simply getting dropped. Whereas host ACLs are used to decide whether to forward packets up the stack towards the application. Since IP does not natively provide packet integrity, address spoofing can

be used to circumvent ACL rules. Employing IPsec, however, prevents such actions.

In Nebula, paths must be established before communication begins, i.e., clients must obtain required PoC tokens. Therefore, access control can be implemented by the consent server granting or denying PoC requests. Traffic sent without valid PoC tokens can be easily detected and dropped.

Unlike IP, enforcing access control in NDN should be done based on content and not network entities. Although not implemented in practice, ACLs can still be used to implement access control. In this case, rules are applied on interest messages and content objects based on the names they carry. Longest-prefix can also be employed to grant or restrict access to entire namespaces. Due to the fact that NDN interests do not carry source addresses, access control on the consumer granularity can only be achieved in cases where interests are signed or carry some form of consumer identity [50].

One way of providing access control in NDN is by using encryption. Producers can encrypt their content and disseminate decryption keys to authorized consumers only. Such keys can be encapsulated in content objects and should not be cached. One drawback of this approach is that it requires consumers to issue at least two interests for each content (one to request the content itself and one to request the corresponding key).

Since MF and XIA can support different principal types, they facilitate the combination of both NDN- and IP-based access control schemes. For content principals, access control is done at the content granularity, similar to NDN, e.g., content is encrypted using keys disseminated to only authorized users. For all other principal types, ACLs can restrict access to hosts and other network services.

F. Accountability

One of the main problems in IP is accountability. In fact, IP is subject to source address spoofing that lead to the inability of tracing back the entity responsible for a particular action. A simple countermeasure against IP spoofing requires ASes to implement egress filtering and ensure that all outgoing traffic carries source addresses owned by these ASes. IPsec guarantees peer entity authentication when establishing SAs between hosts. Thus, accountability can be achieved.

Nebula provides accountability through path establishment. All routers on a path consent to use the whole path before the communication begins. Moreover, the fact that these routers pre-agree on performing a specific set of actions on each packet passing through allows the detection of any malicious activities.

NDN provides full accountability of producers. Since every content is signed by its producer, tracing the producer responsible for generating content is a trivial task. However, accountability can not be provided if content is served from router caches. Consumers accountability, on the other hand, can only be achieved when they issue signed interests, or include their identities in the interests themselves. Otherwise, accountability can not be provided.

Both MF and XIA do not provide accountability at the network layer. However, signing requests and responses can provide this feature in a similar fashion to NDN, especially for content principals. Also, IPsec-similar techniques can be employed to provide accountability for other principal types.

G. Availability

Denial-of-Service (DoS) and Distributed Denial-of-Service (DDoS) are well- known attacks on availability. In the following, we discuss DoS and DDoS attacks on the network layers of current and future Internet architectures discussed above. We also discuss new (and possibly more serious) types of attacks triggered by the new architectures. We specifically exclude availability issues that arise due to network misconfigurations, disasters, hardware/software faults, or any other causes that are not a direct consequence of an attack.

Bandwidth Depletion Attacks. The current Internet is susceptible to bandwidth depletion attacks [116] that aim to exhaust bandwidth of a specific link. These attacks can be conducted in two ways: (1) distributed – with packets sent at a relatively low rate by each attack source, or (2) centralized – a single powerful adversary flooding the target link at a high rate. Due to today's high bandwidth allocation and redundancy, the latter are harder to perform.

Several mitigation and prevention techniques have been proposed and implemented. Examples include: (1) tracing malicious traffic back to the source of the attack [25], [108], [113], (2) distinguishing between legitimate and malicious traffic [20], [129], (3) using puzzles to increase the cost for bandwidth consumption [40], [60], and (4) rate-limiting traffic that causes congestion [59], [80], [118]. However, none of the above is a panacea.

Bandwidth depletion at the data plane is harder to mount in Nebula, since senders (potential adversaries) must obtain consent of all nodes on a path before sending packets. Thus, unauthorized packets are dropped by adversary-facing routers. Unfortunately, this only shifts the attack focus from the network layer to consent servers: a single consent server might be responsible for numerous routers in its domain. Thus, lowering the former's ability to issue PoCs can effectively disable all routers in the domain.

NDN is more resilient to bandwidth depletion attacks as compared to IP counterpart. Recall that NDN communication adheres to the *pull* model, i.e., content is forwarded only in response to a prior interest. This prevents adversaries from flooding the network with unsolicited content. However, an adversary can still flood the network with a large number of interests.

Since MF and XIA use a communication model similar to IP, both are susceptible to bandwidth depletion attacks. Countermeasures similar to those applicable to IP can be adopted. However, as mentioned above they can only lower, not eliminate, the impact of such attacks.

The work in [96] suggested integrating STRIDE into XIA to protect against DoS attacks. STRIDE [58] is an architecture resilient to bandwidth depletion (D)DoS attacks. It modifies SCION path establishment to perform tree-based bandwidth allocation. Whenever a trusted domain root initiates the path establishment process, bandwidth is allocated as the path is branching out as a tree from that root. This guarantees required

bandwidth for benign flows. STRIDE also supports long-term static paths to provide high available connectivity.

Routers Resource Exhaustion. Exhausting router resources is another usual DoS and DDoS attack goal. For example, a NAT-enabled IP router might assemble fragments and, at any given time, might maintain multiple reassembly buffers. Each IP packet fragment includes a 16-bit field indicating the original packet size. An attack can involve sending a single fragment with a very large original packet size. This forces the target router to allocate a new buffer and wait for remaining fragments, which will never arrive.

Other types of attacks might be computational in nature. In Nebula, an adversary can generate numerous packets forcing a router to perform all verification operations described in Section III. Such attacks cost very little for an adversary since malicious packets can simply include invalid (e.g., random) PoC and PoP values.

In NDN, the PIT is a mandatory router component that enables interest collapsing and content delivery without source or destination addresses. However, since it is a limited and valuable resource makes the PIT susceptible to exhaustion. An adversary can easily generate and send (in close proximity) a large number of interests that fill up a router's PIT. To prevent interests from being collapsed, each can refer to some *nonsensical* content. Once the PIT is full, a router can either: drop new incoming interests, or remove old PIT entries to make space for new ones. Both options, however, can adversely impact past or future interests. This type of attack is called Interest Flooding (IF).

Unfortunately, there is no comprehensive remedy for IF attacks. Although several countermeasures have been proposed, they are ineffective against smart adversaries and only manage to lower the volume of IF attacks [10], [34]. One comprehensive, though drastic, remedy is to eliminate the PIT – the main target of IF attacks. For this reason, [52] suggests a modified Content-Centric Networking (CCN) architecture without router PITs.

Router resource exhaustion attacks do not apply to MF. Similar to IP routers, MF routers do not perform expensive cryptographic operation nor do they handle per connection-state. Moreover, MF does not implement NAT, rendering NAT-based attacks irrelevant.

Cache-Related Attacks. Despite the benefits of in-network caching, it prompts new types of (D)DoS attacks that are not relevant to today's Internet: content poisoning and cache pollution. We now discuss resiliency of NDN, MF, and XIA against these attacks. Nebula is excluded since it does not provide in-network caching.

a) Content Poisoning: Content poisoning occurs when an adversary injects fake content into router caches. A fake content is not generated by a benign producer and, consequently, does not satisfy user requests. If cached in routers, such content is used to satisfy future requests.

Since NDN adheres to the pull model makes it harder, yet still feasible, to inject fake content into router caches, in at least two ways:

- Reactive: the adversary is a node eavesdropping or controlling a link, e.g., an upstream malicious router. The adversary responds to interests on that links with a fake content that is cached in all downstream routers.
- Proactive: this method involves the adversary that, anticipating demand for certain content, issues one or more bogus interests (perhaps from strategically placed zombie consumers), before genuine interests are issued. The adversary then replies with fake content (from a set of compromised routers or compromised producers) thus pre-poisoning the caches of all routers forwarding the bogus interests.

Reference [51] investigated the causes of content poisoning and proposed a simple approach for its full mitigation. The main idea is for consumers and producers to collaborate in providing routers with enough contextual trust information to perform a single signature verification per content.[8] This approach is captured by a simple tenet called *Interest Key Binding (IKB)* rule.

An orthogonal content poisoning mitigation mean is the use of SCNs. By definition, an SCN contains a value that (uniquely) identifies the principal to which it refers. In case of a content principal in MF and XIA, and content objects in NDN, an SCN is the hash of the data itself. When a user requests content using SCN, the network guarantees that requested content will be correctly delivered. As a result, MF and XIA obviate content poisoning attacks, by design. It is worth mentioning that using SCNs does not prevent adversaries from injecting fake content in router caches. Instead, it guarantees that benign users will not receive such fake content. In order to use SCNs, the system needs to be bootstrapped *without* them, e.g., by using IKB, as described above.

b) Cache Pollution: Pollution is another type of (D)DoS attacks against router caches. In such attacks, adversaries attempt to manipulate reference locality of caches, causing incorrect decisions made by cache eviction strategies. This causes routers to possibly evict popular content reducing the overall content distribution performance. NDN, MF, and XIA are all susceptible to this attack.

Conti *et al.* discuss this attack in [37]. It is shown that with even limited adversarial resources, a highly effective cache pollution attack can be mounted. In fact, even small cache locality manipulation can cause a significant content distribution disruption [45]. It is also shown in [37] that launching pollution attacks on large networks is relatively easy, and smart adversaries reduce the effectiveness of proposed countermeasures.

Cache pollution attacks do not prevent users from retrieving the requested data. Instead, they negatively effect the performance of content distribution, and eliminate the benefits of in-network caching.

H. Data Confidentiality

The natural way to achieve data confidentiality at the network layer is by using encryption.

IP does not provide data confidentiality. This is done by using IPsec. The level of confidentiality depends on the

[8]This assumes that in the near future, public key-based signature verification will be available in hardware and will be achievable at line speed.

mode of operation. In transport mode, ESP only encrypts the IP packet payload and data confidentiality is host-to-host. Tunnel mode extends confidentiality to the entire encapsulated IP packet, including both payload and header. However, data confidentiality can only be achieved in host-to-gateway or gateway-to-gateway scenarios. Also, ESP confidentiality is not generally effective against active adversaries. It has been demonstrated that achieving confidentiality without a strong integrity mechanism, or even applying integrity before encryption, can only protect against passive adversaries [24], [43], [67]. Thus, even though IPsec provides confidentiality, poor usage practices can negate its benefits.

Nebula's ICING-based network layer does not natively provide data confidentiality. Instead, it can be achieved by combining ICING with end-to-end encryption of the packet payload.

Data confidentiality in NDN can be attained by encrypting content payload. This is not supported by the architecture and is left to the application. However, even if content is encrypted, the fact that it carries a human-readable name might leak information about its data.

MF and XIA provides content principal confidentiality using methods similar to the those used in NDN. Fortunately, and due to the fact that such principal identifiers are generated using the hash of the content itself, inspecting them does not leak information about the encrypted content. Moreover, confidentiality of data communicated between host, network and service principals can be achieved using similar techniques to IPsec.

I. Traffic Flow Confidentiality

It is well known that encryption does not protect against statistical traffic analysis – attacks that monitor traffic in order to extract properties, such as volume and timing [104].

IPsec provides some traffic flow confidentiality by padding packet payloads to hide their size patterns. However, according to IPsec specifications, this is not mandatory and, therefore, may not be supported in all IPsec implementations [109].

NDN, MF, Nebula and XIA are all susceptible to traffic analysis attacks Fortunately, padding can be used to provide traffic and flow confidentiality.

Another architecture-agnostic alternative is to add artificial delays to communications to better hinder time-based attacks. This, however, comes at the expense of increasing end-to-end latency and reducing overall network performance, especially for time-sensitive traffic.

J. Anonymous Communication

IP (with or without IPsec) does not support anonymous communication. This is mainly because source and destination addresses are in the clear in packet headers. However, partial anonymity can be achieved using the tunnel mode of IPsec along with ESP. This is because tunnel mode allows the ESP protocol to encrypt the original IP packet along with the source and destination addresses, and it encapsulates that packet into a new one with a new header reflecting gateway addresses. This combination hides end-host identities among the set of other hosts connected to respective end-networks. However, this is only effective if the adversary is eavesdropping on the link between the two gateways and is not located inside one of the end-networks. Furthermore, in case of host-to-gateway tunnel mode, only anonymity of hosts located behind the gateway is preserved.

Crowds [105] is one of the first proposals to achieve user anonymity. In it, a message is randomly forwarded between group members before it reaches its destination. Therefore, none of the group members nor the end recipient learn the actual source of the message. The Onion Router (Tor) [119] is another method that provides anonymous communication through a "circuit." Circuits are multi-hop encrypted communication channels established using at least three Tor nodes. Theoretically, Tor guarantees anonymity with respect to an adversary controlling, at most, two Tor nodes. However, flawed Tor implementations can reduce its provided anonymity level [26].

Hosts anonymity is not provided by ICING-based Nebula. By inspecting packet headers, eavesdroppers can easily determine a packet's source, as well as the path it traversed. However, host anonymity can be achieved by replacing ICING with TorIP [73], thus resulting in a level of anonymity similar to that provided by Tor in today's Internet.

Unlike IP, NDN has some features that facilitate anonymous communication. A PIT allows interest messages and content objects to only carry the requested content name without any consumer-related information. While this help to protect consumer's privacy form on-path observers inspecting packets deep in the core network, its efficacy will decrease as the observer moves closer to the consumer. At the network access (e.g., wireless access point or a broadband network gateway) an observer can easily correlate interests with the consumer issuing them. If the observer sits multiple hops far away from the consumer, he will only be able to correlate interests with a set of consumers (i.e., all the consumers reachable from the observer). Moreover, Compagno *et al.* show in [33] that off-path adversaries can abuse the NDN content caching and forwarding state to determine consumers' location. DiBenedetto *et al.* [46] proposed ANDāNA, a tool that provides a level of anonymity similar to Tor, while requiring only two intermediate nodes, instead of three.

MF and XIA suffer from the same privacy and anonymity problems as IP. Packets contain both source and destination GUIDs (or principal identifiers), thus revealing the hosts involved. To make the matter worse, XIA packets path can be revealed by inspecting their destination DAG addresses. This is because such addresses might include (part of) the path to the destination, as described in Section VI. Due to the similarity to IP with respect to the communication model adopted, both MF and XIA can use approaches developed to preserve users anonymity in IP networks. For instance, Tor can be used to protect MF and XIA host principals' anonymity [81].

VIII. Resolution Services Security

Resolution service is a fundamental part of the current Internet architecture. It maps human-readable names to

routable network addresses. As mentioned above, new Internet architectures also require similar resolution services to operate. In this section we compare the security and privacy features of the various resolution services proposed in FIA projects. We exclude Nebula and XIA since they do not propose a new resolution service and only exploit the existing DNS, and its security extensions.

We consider the following security features: trust, data origin authentication, data integrity, peer entity authentication, authorization and access control, accounting, data confidentiality and availability. We consciously exclude other privacy-related features, i.e., traffic flow confidentiality and anonymous communication, since we believe they should be provided at the network layer. We summarize our comparison on security and privacy features in resolution services in Table III.

A. Trust

DNSSEC introduces the notion of trust into DNS. It considers authoritative servers as trusted entities responsible of maintaining and securely providing the correct mapping between human-readable domain names and corresponding IP addresses. Recall that each DNSSEC server signs the resource records (name-to- IP mappings) of its respective domain. Trustworthiness of such servers is ensured by a chain-of-trust model that resembles the domains hierarchical organization. The top-level domain resides at the root of this chain. DNSCurve considers a different attacker model, where external attackers (malicious non-DNSCurve servers) try to modify queries and/or responses, or to mount DoS attacks. Mitigating such attacks requires additional trust in every DNSCurve-enabled server in the system.

Similar to DNSSEC, NDNS applies the same notion. Authoritative server are trusted entities and their trustworthiness is ensured by a similar chain-of-trust model.

GNRS, on the other hand, adopts a different approach. Every network is responsible of providing signed GUID-NA mappings. Thus, verifying these signatures ensures their validity. This also prevents (compromised) GNRS from manipulating GUID-NA mappings without being detected.

B. Data-Origin Authentication and Data Integrity

DNSSEC provides data-origin authentication and data integrity by requiring: (1) authoritative servers to sign each of their resource records, and (2) resolvers to verify the validity of these signatures and their corresponding public keys. This prevents adversaries from injecting bogus data into the DNS system. Signing every response resource record is an expensive operation that authoritative server should not perform at run-time. Adversaries can abuse such costly operation to launch (D)DoS attacks against authoritative servers. To this end, resource record signatures should always be generated in advance. While this method can be easily applied in case resolvers asks for existing domain names, it does not work for a non-existing domain names. DNS uses the NXDOMAIN resource records to inform a resolver that the queried name does not exist. However, providing data-origin authentication and integrity for NXDOMAIN resource records can not be done by generating the signature in advance, because of the number of possible non-existing names is infinite. To solve this problem, DNSSEC introduces a new record type called the NextSECure (NSEC) resource record. Specifically, assuming a canonical ordering of the domain names, a NSEC record contains two consecutive existing names in the canonical ordering, thus describing the gaps between them. Such records are signed and used as authenticated denial of existence for non-existing names. Since NSEC records contain existing names, their signatures can be calculated *a priori*. DNSCurve guarantees integrity of queries and responses exchanged with next-hop servers. However, it does not guarantee data-origin authentication.

In NDNS, query responses are carried in content object payloads, thus data-origin authentication and integrity is inherited from NDN. One difference between DNSSEC and NDNS resides in the granularity of these authentication and integrity guarantees. While DNSSEC can offer such security properties per individual resource record or resource record set, the fact that a NDNS record is carried in a content object can only guarantee the authentication and integrity of said record. Although this is a clear restriction in the flexibility and scalability of the protocol, it does not jeopardize its security [11]. In order to overcome DoS attacks due to requests for non-existing names, NDNS adopts methods similar to DNSSEC. In particular, NDNS servers can sign the gaps between existing names. Furthermore, Compagno *et al.* [35] proposed the use of a Bloom filter [27] to defeat the aforementioned DoS attacks. This, however, requires some changes in the NDN architecture as well as the introduction of a new content type.

GNRS also provides data-origin authentication and integrity by means of GUID-NA mapping signatures. The main different between GNRS and DNSSEC is that the former does not assume that GNRS servers are trusted [74], while the latter requires all authoritative servers to be trusted.

C. Peer Entity Authentication

In DNS, and DNSSEC, there is no entity authentication between a resolver and a DNS server. Resolvers usually know the IP address of a DNS server which is used to initiate queries. Usually, such IP address is manually configured on the resolver or obtained through DHCP and considered valid. Using a TLS connection between resolvers and DNS servers can provide entity authentication [28], [134]. Authentication among DNS servers is obtained through Transaction signatures (TSIG) [126]. TSIG involves pairwise keys shared among DNS servers and used to secure dynamic updates, zone transfers and recursive queries. Moreover, in case of dynamic updates generated from DNS clients, a signature is used to authenticate these clients, and validate the update. Similar to DNS and DNSSEC, DNSCurve does not provide entity authentication by default. However, this could be easily achieved by associating client public keys with certificates.

NDNS follows the same approach of DNS by not providing entity authentication. However, the fact that both NDNS users and servers are regular NDN consumers and producers, respectively, allows approaches similar to what is proposed in

TABLE III
RESOLUTION SERVICES SECURITY & PRIVACY FEATURES COMPARISON

		DNS/DNSSEC	NDNS	GNRS
Trust	*Trusted Model*	Hierarchical. Authoritative Servers are trusted for their respective domains. Delegation is supported	Hierarchical. Authoritative Servers are trusted for their respective domains. Delegation is supported	Distributed. Hosts, DHCP servers, Border routers must cooperate.
	Trusted entity	Authoritative Server.	Authoritative Server.	None.
	Trust Management	PKI + certificates, shared keys.	Certificates.	Distributed architecture + certificates.
Data Origin Auth.	*Coverage*	Authoritative Server-to-Client.	Authoritative Server-to-Client.	GNRS server-to-Client.
	Type of packet	Response.	Response.	Response.
	Mechanism	Signature.	Signature.	Self Certifying name + Signature.
Peer Auth.	*Authenticated Entity*	None.	None.	Client and GNRS server.
	Mechanism	None.	None.	Signature.
Integrity	*Coverage*	Authoritative Server-to-Client.	Authoritative Server-to-Client.	GNRS Server-to-Client.
	Type of packet	Response.	Response.	Request.
	Mechanism	Signature.	Signature.	Self Certifying name + Signature.
Authz. & Access Cont.	*Enforced on*	None.	None.	GNRS Server.
	Mechanism	None.	None.	ACL.
Accountability	*Entity accounted*	Host updating DNS.	None	Client and host updating GNRS.
	Mechanism	Signature.	None.	Signature.
Availability	*DNS Request Flooding Attacks*	IP "Anycast" distribute the number of request across many authoritative servers.	Secondary server to distribute the load + in network load-balancing techniques.	More resilient to flooding due to its distributed design.
Data Conf.	*Mechanism*	None.	None.	None.

Section VII-C to be adopted. Furthermore, securing dynamic updates requires NDNS clients to have previously shared their certificate with the NDNS servers. Signing the updates will then authenticate them.

Unlike DNS and NDNS, GNRS authenticates every client issuing requests (query, update and delete) by retrieving the corresponding certificate, from NCS, and verifying the GUID authenticity. GNRS clients can ensure server authentication by performing similar steps, requesting certificates from NCS and verifying servers identities.

D. Authorization and Access Control

Both DNS/DNSSEC and NDNS do not provide any form of access control. All resource records are publicly available to every host in the network. However, authorization and access

control is provided for dynamic updates. With DNSCurve, on the other hand, exchanged packets between are encrypted, and thus cannot be publicly accessed by other unauthorized entities in the network.

GNRS does not follow the same trend and consider access control a crucial part of its design. Specifically, GNRS stores a set of access control policies along with GUID-NA mappings. Such policies regulate access to particular GNRS resources by specifying read and write permissions which blacklist and whitelist certain user GUIDs.

E. Accountability

DNS/DNSSEC guarantees accountability only for secure DNS dynamic updates [128]. This is because such requests must in fact be signed by their originators. DNSCurve allows accountability for queries, since they contain client public keys.

NDNS does not provide any mechanism for accountability. The fact that NDNS users and servers are consumers and producers allows the adoption of similar approaches described in Section VII-F.

GNRS uses a different approach and mandates GNRS clients to sign every request. By doing so, accountability is provided for all insert, update and query requests.

F. Availability

As a public available service, DNS is subject to (D)DoS attacks which jeopardize its availability. In particular, adversaries can flood authoritative servers with a large number of query requests to exhaust their resources.[9] Although the use of DNS caching and redundancy servers reduce the effect of DDoS, a number of such attacks have been successfully directed against root and top-level DNS servers in past years [2], [4], [5]. Many solutions have been presented that either: (1) require some changes in the DNS protocol [41], [87], [98], [103], or (2) propose new resolution services [57], [130]. Nowadays, DNS uses a single approach that does not require any modification to its architecture, which is the adoption of "Anycast" [9]. In this case, a single DNS server is replicated in different geographically locations among several ASes. Therefore, routing protocols forward DNS requests to the nearest server that can satisfy them [36]. However, this approach can not achieve an efficient load balancing because it does not consider replica workloads and network traffic.

DNS can be exploited to launch (D)DoS amplification attacks against other hosts. These attackstake advantage of open DNS resolvers, and involve forwarding (D)DoS traffic through DNS servers to amplify the volume of data being sent to a target [38], [61]. This issue is exacerbated by DNSSEC [107], which may be exploited by attackers to cause an amplification effect of 50 or more times the original attack bandwidth [16], [123]. DNSCurve has approximately the same amplification effect as regular DNS.

DNS resolvers (not implementing DNSSEC protocol), and DNSCurve resolvers, can be subject to cache poisoning attacks. This attack is similar, in concept, to content poisoning described in Section VII-G. In DNS cache poisoning attacks, the goal of the adversary is to inject spoofed responses (name-IP mappings) in the resolver caches [115]. Injecting false DNS responses can be achieved using a man-in-the-middle attack in which the adversary satisfies requests with false DNS responses [115]. The introduction of data- origin authentication in DNSSEC allows resolvers to verify the origin of data in DNS response, thus eliminating this attack.

NDNS follows the same hierarchical design of DNS. In principles NDNS authoritative servers seems to offer the same level of resilience to DDoS as DNS authoritative servers. Furthermore, NDNS envisions the use of a set of secondary servers to balance the workload, which was proven to be not effective. The same anycast approach used in DNS could be employed in NDNS. Such forwarding strategy must be implemented by NDN routers. Moreover, NDNS is susceptible to content poisoning attacks. Fortunately, the same countermeasures described in Section VII-G can be effective.

GNRS design appears to be more resilient to DDoS than DNS design. In fact, GNRS does not adopt a hierarchical structure, instead it distributes the GUID-NA mappings among a number of replicas using Distributed Hashtables (DHT) maintained by all the ASes in the Internet. This allows GNRS to easily scale and distribute a DDoS attack.

G. Data Confidentiality

Neither DNS nor DNSSEC provide confidentiality. Queries and resource records are never encrypted by authoritative servers and are always exchanged in cleartext. One way of achieving confidentiality in DNS/DNSSEC is to establish a TLS session between resolvers and authoritative servers [28], [134]. DNSCurve, however, uses symmetric encryption to provide data confidentiality of DNS messages.

Similarly, both NDNS and GNRS designs do not take confidentiality into consideration. Fortunately, similar approaches to using TLS channels can be adopted. This feature is important to provide private communication.

IX. Lessons Learned and Future Directions

In this section, we summarize security and privacy features provided by the network layer in the four studied FIA architectures, as well as their respective resolution services. We also highlight missing features and outline directions for future work. Finally, we briefly discuss the current implementation status and performance characteristics of each architecture.

A. Network Layer

Overall, we believe that the NSF mandatory requirement of *Security and Privacy by Design* has only been partially accomplished. Table IV summarizes security and privacy features provided by each FIA architecture.

Nebula. At its current state, Nebula includes trust, data origin authentication, peer entity authentication, data integrity, authorization and access control, accountability, and availability features. All of them are provided between senders and routers implementing ICING, except for peer entity

[9]The use of secure dynamic updates involving asymmetric encryption increases the effect of the attack.

TABLE IV
NETWORK LAYER SECURITY AND PRIVACY COMPARISON

Security & Privacy Features	Network layers			
	Nebula	NDN	MF	XIA
Trust	✓	✓	✓	✓
Data Origin Authentication	◎	✓	✗	✗
Peer entity Authentication	◎	◎	◎	◎
Data Integrity	◎	✓	✗	✗
Authorization and Access Control	✓	◎	◎	◎
Accountability	✓	◎	◎	◎
Data Confidentiality	✗	✓	✗	✗
Traffic Flow Confidentiality	✗	✗	✗	✗
Anonymous Communication	✗	✗	✗	✗
Availability	◎	◎	◎	◎

✓indicates that a feature is present in the architecture.
✗indicates that a feature is unavailable.
◎ indicates that a feature is partially available or the architecture provides a means to facilitate its implementation by applications.

authentication that is guaranteed between senders and consent servers during PoCs retrieval. With respect to IP and IPsec, Nebula certainly increases the security of intra-domain communication. However, it lacks inter-domain and end-to-end communication security support, even if IPsec is used to establish a secure "pipe" between communicating end-hosts. We believe that integration of new or existing mechanisms to support both inter-domain and end-to-end security deserves further investigation.

Availability in Nebula is provided by path verification mechanism which prevents any adversary from sending unrequested data to perform bandwidth depletion attacks. However, ICING nodes and consent servers could be the target of a DoS attack, due to the heavy use of cryptographic operations. Consent servers can be another target for DoS attacks. In particular, an adversary can flood a server with an abnormal amount of request. If the target server controls a large number of routers, the attack can effectively disable all of them. We believe that such attack deserves further investigation.

By design, Nebula does not provide the three privacy features we considered in this paper: data confidentiality, traffic flow confidentiality, and anonymous communication. The first two can be easily implemented via IPsec-like techniques, while the latter can be achieved by adopting existing solutions, such as TorIP [73]. Nevertheless, we believe these aspects in Nebula deserve further investigation. Specifically, Nebula should make these features available by design before any adoption in the real world. Furthermore, anonymous communications conflicts with the current ICING design, which requires path establishment for each communication. This is an issue of concern that needs to be explored and, ideally, remedied.

From a feasibility perspective, ICING can be effectively deployed in Internet backbone routers [91], despite its heavy use of cryptographic operations at the data plane, and its large header size.[10] However, as reported in [91], this incurs a higher normalized cost of 93% over IP, i.e., hardware cost as equivalent gate count per unit of throughput.

NDN. NDN moves hosts and interfaces into the background and focuses on content as its primary network-layer entity. This strongly influences network-layer security and privacy features. Currently, NDN provides: trust, data origin authentication, data integrity and data confidentiality. Other features (e.g., peer entity authentication and accountability) are available only when derived from data origin authentication for both interest and content. Ubiquitous caching further complicates achieving these two features. In fact, when a consumer obtains a content from an intermediate router's cache, there is currently no way to provide peer entity authentication between the consumer and the router providing the content. Similarly, NDN does not provide accountability for content use: the distinction between a content served by its producer or from a router's cache.

Availability is an aspect for which NDN provides some improvements with respect to IP/IPsec. NDN reduces the amount of traffic both inside the network and at the producer by: (1) aggregating "matching" interests, and (2) serving matching requests from local caches, when possible. Unfortunately, at the same time NDN opens the door to new types of attacks. Indeed, while the pull model communication prevents any adversary from flooding a host with non-requested content, adversaries can exhaust router states (i.e., PIT and CS), decreasing the performance of a network. Even though several countermeasures have been proposed, none of them has been chosen and implemented as part of the architecture.

The remaining security features: authorization and access control, traffic flow confidentiality and anonymous communication are not provided by NDN at the network layer. Even if ACL can be implemented, NDN delegates access control to the application layer, which is responsible for encrypting content and distributing appropriate keys only to authorized consumers. Traffic flow confidentiality and anonymous communication are not natively available. The former can be provided by adopting existing approaches designed for IP and

[10]ICING's header occupies approximately 18.3% of a 1514 bytes packet, compared to 1.3% with IP, on a path through 7 realms [90].

IPsec, without any modification of the NDN architecture. The latter can be obtained by adopting solutions similar to the one presented in [46]. Note that NDN design improves privacy guarantees for consumers, since neither interest nor data packets carry consumers address. Nevertheless, consumer privacy can still be compromised by motivated attackers [33].

Heavy use of digital signatures in NDN introduces a major source of overhead in the network. This has been extensively evaluated in [114]. To the best of our knowledge, no assessment on the implications of router signature verification in NDN has been conducted. The only work considering the performance impact of signature generation by producers has been recently evaluated in [82]. From an interoperability perspective, NDN can be used alongside with IP as an overlay network.

MobilityFirst (MF). MF security features can be viewed as a hybrid of end-to-end and content-based approaches of IP/IPsec and NDN. In its current state, MF provides only few security features: trust as well as authorization and access control. The adoption of SCNs to address hosts facilitates the implementation of data origin authentication, peer entity authentication, and accountability. Peer entity authentication can be achieved involving a simple challenge-response protocol between the two peers. While this would guarantee end-to-end accountability, it is not enough for network accountability. In order to achieve that, edge routers should prevent address spoofing.

Data integrity and data confidentiality are currently not provided at network layer and it remains unclear whether MF will delegate these aspects to upper layers. Traffic flow confidentiality and anonymous communication are also not natively provided. Existing approaches for traffic flow confidentiality designed for IP/IPsec can be easily imported. However, TOR-like solutions for anonymous communication should be further studied and developed. A relatively recent initial effort in this direction is represented in [81]. Additionally, according to [70], the use of Disposable Identifiers [54], [71] is currently under examination to guarantee GUID-NA unlinkability. We believe that more research is needed to construct techniques that take advantage of MF's architectural idiosyncrasies.

MF is somewhat effective against content poisoning attacks due to its usage of SCNs. Meanwhile, it does not consider other network attacks, such as bandwidth depletion and cache pollution. Current IPsec-like methods can be applied to mitigate these issues. Further work is necessary to provide new architecture-specific countermeasures.

To the best of our knowledge there is no large scale performance assessment of MF. Venkataramani *et al.* [124] performed some benchmarks on a commodity hardware, and reached a maximum forwarding rate of 1 Gbps. We believe that further effort should be devoted to study and improve MF performance to meet real-world requirements.

XIA. XIA's main goal is to support communication between multiple and different principals. XIA security approach extends, de-facto, MF approach, where security focuses on two principals: content and hosts. XIA does not limit the number of principals but instead provides the freedom to design new ones. To this end, XIA security is based on each principal intrinsic security feature.

In its current state, XIA offers the same security and privacy features as MF. This is due to their similar approach in addressing principals using SCNs.

It is worth mentioning that XIA's performance is comparable with IP in core network routers [78].

TABLE V
RESOLUTION SERVICES SECURITY & PRIVACY COMPARISON

Security and Privacy Features	Resolution Services	
	NDNS	GNRS
Trust	✓	✓
Data Origin Authentication	✓	✓
Peer entity Authentication	✗	✓
Data Integrity	✓	✓
Authorization and Access Control	✗	✓
Accountability	◎	✓
Data Confidentiality	✗	✗
Availability	◎	✓

✓indicates that a feature is present in the architecture.
✗indicates that a feature is unavailable.
◎ indicates that a feature is partially available or the architecture provides a means to facilitate its implementation by applications.

B. Resolution Services

Nebula's resolution service is used by the NVENT to perform path discovery. In NDN, the same service provides the mapping between namespaces and the corresponding security information, i.e., a mapping between public key(s) and name prefixes. In MF, the resolution service is actively involved in any communication guaranteeing the binding between GUIDs and NAs. Finally, XIA's DNS-based resolution service is used for the address/path resolution.

Although all architectures requires a resolution service, only NDN and MF are actually investigating their own proposal. Table V summarizes security and privacy features of NDNS and GNRS.

NDNS. NDNS borrows many features from DNS and DNSSEC without introducing much novelty. NDNS involves the same notion of trust of DNSSEC in which servers are trusted entities. Moreover, NDNS queries and responses provide: data origin authentication, data integrity, accountability (only for dynamic updates), and availability. The last feature is provided by server replication. However, while DNSSEC uses the "Anycast" technique to choose the closes server to the resolver, in NDNS the network must be aware of all servers and implement specific forwarding strategies to balance the requests among them, which adds more complexities.

Peer entity authentication, authorization and access control, as well as data confidentiality, are not provided. We believe that NDNS should be enhanced to provide better availability and currently missing security features.

GNRS. GNRS is fundamentally different from DNS, DNSSEC, and DNSCurve. In fact, since MF adopts a flat name structure, it also forces GNRS to use the same structure for its servers. This different organization has some side effects

on the provided security features: servers are not assumed to be trusted. Also, data origin authentication is provided by the owner of the GUID and not by servers. GNRS also introduces authorization and access control of its stored information and provides accountability and peer entity authentication for each query and response. Finally, GNRS is more robust to DoS attacks than DNS. First, compromising one server does not affect others. Second, GNRS is designed to easily adapt to network changes and to balance GUID-NA mapping among many replicas based on locality of requests.

We believe that GNRS introduces good security improvements with respect to DNS, DNSSEC, and DNSCurve. The only missing feature is data confidentiality.

X. Conclusion

Despite the unmitigated success of the current IP-based Internet architecture, lack of security considerations in its design led to numerous security and privacy problems, over the years. Thus, a key goal of any future Internet architectures must be to include – from the outset – a set of comprehensive and extensible security and privacy features.

This paper analyzed (mostly network-layer) security and privacy features of four prominent NSF-funded FIA architectures – Nebula, NDN, MobilityFirst and XIA – with IP/IPsec used as a point of reference in the analysis. Prior surveys on future Internet architectures provide a limited, or even no, comparison on security and privacy features. This paper provides a comprehensive and neutral analysis of salient security and privacy features (and issues) in these NSF-funded Future Internet Architectures.

As evident from this work, while each FIA architecture offers some innovative and effective security and privacy features, none provides a comprehensive set thereof.

References

[1] *IPv6 Extension Headers Review and Considerations*. Accessed: May 26, 2016. [Online]. Available: https://www.cisco.com/en/US/technologies/tk648/tk872/technologies_white_paper0900aecd8054d37d.html

[2] *Nameserver DoS Attack October 2002*. Accessed: Jan. 31, 2018. [Online]. Available: https://www.caida.org/projects/dns/dns-rootgtld/oct02dos.xml

[3] *Signed Interest*. Accessed: Jan. 31, 2018. [Online]. Available: http://named-data.net/doc/ndn-cxx/current/tutorials/signed-interest.html

[4] (2002). *Ultradns DoS Attack*. [Online]. Available: http://www.theregister.co.uk/2002/12/14/

[5] (2007). *ICANN Factsheet for the February 6, 2007 Root Server Attack*. [Online]. Available: http://www.icann.org/announcements/factsheet-dns-attack-08mar07.pdf

[6] M. Aamir and S. M. A. Zaidi, "Denial-of-service in content centric (named data) networking: A tutorial and state-of-the-art survey," *Security Commun. Netw.*, vol. 8, no. 11, pp. 2037–2059, 2015.

[7] E. G. AbdAllah, H. S. Hassanein, and M. Zulkernine, "A survey of security attacks in information-centric networking," *IEEE Commun. Surveys Tuts.*, vol. 17, no. 3, pp. 1441–1454, 3rd Quart., 2015.

[8] J. Abley and T. Manderson, "Nameservers for IPv4 and IPv6 reverse zones," IETF, Fremont, CA, USA, Rep. RFC 5855, 2010.

[9] J. Abley and K. E. Lindqvist, "Operation of anycast services," IETF, Fremont, CA, USA, RFC 4786, 2006.

[10] A. Afanasyev, P. Mahadevan, I. Moiseenko, E. Uzun, and L. Zhang, "Interest flooding attack and countermeasures in named data networking," in *Proc. IFIP Netw. Conf.*, Brooklyn, NY, USA, 2013, pp. 1–9.

[11] A. Afanasyev, "Addressing operational challenges in named data networking through NDNS distributed database," Ph.D. dissertation, Dept. Comput. Sci., Univ. California at Los Angeles, Los Angeles, CA, USA, 2013.

[12] A. Afanasyev, J. Shi, L. Wang, B. Zhang, and L. Zhang, "Packet fragmentation in NDN: Why NDN uses hop-by-hop fragmentation," Named Data Netw., Rep. NDN-0032, 2015.

[13] A. Afanasyev, C. Yi, L. Wang, B. Zhang, and L. Zhang, "Map-and-ENCAP for scaling NDN routing," Univ. California at Los Angeles, Los Angeles, CA, USA, Rep. NDN-0004, 2015.

[14] B. Ahlgren, C. Dannewitz, C. Imbrenda, D. Kutscher, and B. Ohlman, "A survey of information-centric networking," *IEEE Commun. Mag.*, vol. 50, no. 7, pp. 26–36, Jul. 2012.

[15] S. Akhshabi and C. Dovrolis, "The evolution of layered protocol stacks leads to an hourglass-shaped architecture," *SIGCOMM Comput. Commun. Rev.*, vol. 41, no. 4, pp. 206–217, 2011.

[16] M. Anagnostopoulos, G. Kambourakis, P. Kopanos, G. Louloudakis, and S. Gritzalis, "DNS amplification attack revisited," *Comput. Security*, vol. 39, pp. 475–485, Nov. 2013.

[17] D. Comer, "A future Internet architecture that supports cloud computing," in *Proc. 6th Int. Conf. Future Internet Technol. (CFI)*, Seoul, South Korea, 2011, pp. 79–83. [Online]. Available: http://doi.acm.org/10.1145/2002396.2002418, doi: 10.1145/2002396.2002418.

[18] T. Anderson *et al.*, "A brief overview of the NEBULA future Internet architecture," *SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 3, pp. 81–86, 2014.

[19] T. Anderson *et al.*, "The NEBULA future Internet architecture," in *The Future Internet* (LNCS 7858), A. Galis and A. Gavras, Eds. Heidelberg, Germany: Springer, 2013, pp. 16–26.

[20] T. Anderson, T. Roscoe, and D. Wetherall, "Preventing Internet denial-of-service with capabilities," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 34, no. 1, pp. 39–44, 2004.

[21] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose, "DNS security introduction and requirements," IETF, Fremont, CA, USA, RFC 4033, 2005.

[22] M. Arye *et al.*, "Increasing network resilience through edge diversity in NEBULA," *SIGMOBILE Mobile Comput. Commun. Rev.*, vol. 16, no. 3, pp. 14–20, 2012.

[23] M. F. Bari, S. R. Chowdhury, R. Ahmed, R. Boutaba, and B. Mathieu, "A survey of naming and routing in information-centric networks," *IEEE Commun. Mag.*, vol. 50, no. 12, pp. 44–53, Dec. 2012.

[24] S. M. Bellovin, "Problem areas for the IP security protocols," in *Proc. 6th Conf. USENIX Security Symp. Focusing Appl. Cryptograph.*, San Jose, CA, USA, 1996, pp. 21–32.

[25] S. M. Bellovin, M. Leech, and T. Taylor, "ICMP traceback messages," Internet Draft, 2003.

[26] A. Biryukov, I. Pustogarov, and R.-P. Weinmann, "Trawling for Tor hidden services: Detection, measurement, deanonymization," in *Proc. 34th IEEE Symp. Security Privacy*, Berkeley, CA, USA, 2013, pp. 80–94.

[27] B. H. Bloom, "Space/time trade-offs in hash coding with allowable errors," *Commun. ACM*, vol. 13, no. 7, pp. 422–426, 1970.

[28] S. Bortzmeyer, "DNS privacy considerations," IETF, Fremont, CA, USA, RFC 7626, 2015.

[29] J. Bound and Y. Rekhter, "Dynamic updates in the domain name system (DNS update)," IETF, Fremont, CA, USA, RFC 2136, 1997.

[30] R. Braden, "Requirements for Internet hosts-communication layers," IETF, Fremont, CA, USA, RFC 1122, 1989. [Online]. Available: https://tools.ietf.org/html/rfc1122

[31] H. C. Cankaya, "Access control lists," in *Encyclopedia of Cryptography and Security*, H. C. van Tilborg and S. Jajodia, Eds. Boston, MA, USA: Springer, 2011, pp. 9–12.

[32] C. Kaufman, P. Hoffman, Y. Nir, P. Eronen, and T. Kivinen, "Internet key exchange protocol version 2 (IKEv2)," IETF, Fremont, CA, USA, RFC 5996, 2010.

[33] A. Compagno, M. Conti, P. Gasti, L. V. Mancini, and G. Tsudik, "Violating consumer anonymity: Geo-locating nodes in named data networking," in *Proc. 13th Int. Conf. Appl. Cryptograph. Netw. Security (ACNS)*, 2015, pp. 243–262.

[34] A. Compagno, M. Conti, P. Gasti, and G. Tsudik, "Poseidon: Mitigating interest flooding DDoS attacks in named data networking," in *Proc. 38th Annu. IEEE Conf. Local Comput. Netw. (LCN)*, Sydney, NSW, Australia, 2013, pp. 630–638.

[35] A. Compagno, M. Conti, C. Ghali, and G. Tsudik, "To NACK or not to NACK? Negative acknowledgments in information-centric networking," in *Proc. 24th Int. Conf. Comput. Commun. Netw. (ICCCN)*, Las Vegas, NV, USA, 2015, pp. 1–10.

[36] D. Conrad. (2012). *Towards Improving DNS Security, Stability, and Resiliency*. [Online]. Available: http://www.internetsociety.org/sites/default/files/bp-dnsresiliency-201201-en_0.pdf

[37] M. Conti, P. Gasti, and M. Teoli, "A lightweight mechanism for detection of cache pollution attacks in named data networking," *Comput. Netw.*, vol. 57, no. 16, pp. 3178–3191, 2013.

[38] D. Dagon, N. Provos, C. P. Lee, and W. Lee, "Corrupted DNS resolution paths: The rise of a malicious resolution authority," in *Proc. 15th Netw. Distrib. Syst. Security Symp. (NDSS)*, San Diego, CA, USA, 2008.

[39] C. Dannewitz *et al.*, "Network of information (NetInf)—An information-centric networking architecture," *Comput. Commun.*, vol. 36, no. 7, pp. 721–735, 2013.

[40] D. Dean and A. Stubblefield, "Using client puzzles to protect TLS," in *Proc. 10th USENIX Security Symp.*, vol. 42. Washington, DC, USA, 2001.

[41] T. Deegan, J. Crowcroft, and A. Warfield, "The main name system: An exercise in centralized computing," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 5, pp. 5–13, 2005.

[42] S. E. Deering and R. Hinden, "Internet protocol, version 6 (IPv6) specification," IETF, Fremont, CA, USA, RFC 2460, 1998.

[43] J. P. Degabriele and K. G. Paterson, "Attacking the IPsec standards in encryption-only configurations," in *Proc. 28th IEEE Symp. Security Privacy (S&P)*, Berkeley, CA, USA, 2007, pp. 335–349.

[44] M. Dempsky, "DNSCurve: Link-level security for the domain name system," IETF, Fremont, CA, USA, Internet-Draft, 2010.

[45] L. Deng, Y. Gao, Y. Chen, and A. Kuzmanovic, "Pollution attacks and defenses for Internet caching systems," *Comput. Netw.*, vol. 52, no. 5, pp. 935–956, 2008.

[46] S. DiBenedetto, P. Gasti, G. Tsudik, and E. Uzun, "ANDaNA: Anonymous named data networking application," in *Proc. 18th Annu. Netw. Distrib. Syst. Security Symp. (NDSS)*, 2011.

[47] N. Fotiou, P. Nikander, D. Trossen, and G. C. Polyzos, "Developing information networking further: From PSIRP to pursuit," in *Proc. Int. Conf. Broadband Commun. Netw. Syst.*, 2010, pp. 1–13.

[48] J. Gersch and D. Massey, "ROVER: Route origin verification using DNS," in *Proc. 22nd Int. Conf. Comput. Commun. Netw. (ICCCN)*, Nassau, Bahamas, 2013, pp. 1–9.

[49] C. Ghali, A. Narayanan, D. Oran, G. Tsudik, and C. A. Wood, "Secure fragmentation for content-centric networks," in *Proc. IEEE 14th Int. Symp. Netw. Comput. Appl. (NCA)*, Cambridge, MA, USA, 2015, pp. 47–56.

[50] C. Ghali, M. A. Schlosberg, G. Tsudik, and C. A. Wood, "Interest-based access control for content centric networks," in *Proc. 2nd Int. Conf. Inf. Centric Netw. (ICN)*, San Francisco, CA, USA, 2015, pp. 147–156.

[51] C. Ghali, G. Tsudik, and E. Uzun, "Network-layer trust in named-data networking," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 5, pp. 12–19, 2014.

[52] C. Ghali, G. Tsudik, E. Uzun, and C. A. Wood, "Closing the floodgate with stateless content-centric networking," in *Proc. 26th Int. Conf. Comput. Commun. Netw. (ICCCN)*, Vancouver, BC, Canada, Jul. 2017, pp. 1–10, doi: 10.1109/ICCCN.2017.8038367.

[53] D. Griffin *et al.*, "Service-centric networking," *Handbook of Research on Redesigning the Future of Internet Architectures*. Hershey, PA, USA: IGI Glob., 2015.

[54] M. Gruteser and D. Grunwald, "Enhancing location privacy in wireless LAN through disposable interface identifiers: A quantitative analysis," *Mobile Netw. Appl.*, vol. 10, no. 3, pp. 315–325, 2005.

[55] T. Hain, "Architectural implications of NAT," IETF, Fremont, CA, USA, RFC 2993, 2000.

[56] D. Han *et al.*, "XIA: Efficient support for evolvable internetworking," in *Proc. 9th USENIX Conf. Netw. Syst. Design Implement. (NSDI)*, San Jose, CA, USA, 2012, pp. 309–322.

[57] M. Handley and A. Greenhalgh, "The case for pushing DNS," in *Proc. Hotnets-IV*, 2005.

[58] H.-C. Hsiao *et al.*, "STRIDE: Sanctuary trail–refuge from Internet DDoS entrapment," in *Proc. 8th ACM SIGSAC Symp. Inf. Comput. Commun. Security (CCS)*, 2013, pp. 415–426.

[59] J. Ioannidis and S. M. Bellovin, "Implementing pushback: Router-based defense against DDoS attacks," in *Proc. 9th Netw. Distrib. Syst. Security Symp. (NDSS)*, San Diego, CA, USA, 2002.

[60] A. Juels and J. G. Brainard, "Client puzzles: A cryptographic countermeasure against connection depletion attacks," in *Proc. Netw. Distrib. Syst. Security Symp. (NDSS)*, vol. 99. San Diego, CA, USA, 1999, pp. 151–165.

[61] G. Kambourakis, T. Moschos, D. Geneiatakis, and S. Gritzalis, "Detecting DNS amplification attacks," in *Proc. Int. Workshop Crit. Inf. Infrastruct. Security*, 2007, pp. 185–196.

[62] S. Kent, "IP authentication header," IETF, Fremont, CA, USA, RFC 4302, 2005.

[63] S. Kent, "IP encapsulating security payload (ESP)," IETF, Fremont, CA, USA, RFC 4303, 2005.

[64] R. Khondoker, B. Nugraha, R. Marx, and K. M. Bayarou, "Security of selected future Internet architectures: A survey," in *Proc. 8th Int. Conf. Innov. Mobile Internet Services Ubiquitous Comput. (IMIS)*, Birmingham, U.K., 2014, pp. 433–440.

[65] T. Kivinen, B. Swander, A. Huttunen, and V. Volpe, "Negotiation of NAT-traversal in the IKE," IETF, Fremont, CA, USA, RFC 3947, 2005.

[66] T. Koponen *et al.*, "A data-oriented (and beyond) network architecture," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 4, pp. 181–192, 2007.

[67] H. Krawczyk, "The order of encryption and authentication for protecting communications (or: How secure is SSL?)," in *Advances in Cryptology—CRYPTO 2001* (LNCS 2139), J. Kilian, Ed. Heidelberg, Germany: Springer, 2001, pp. 310–331.

[68] H. Krawczyk, R. Canetti, and M. Bellare, "HMAC: Keyed-hashing for message authentication," IETF, Fremont, CA, USA, RFC 2104, 1997.

[69] E. Lewis and A. Hoenes, "DNS zone transfer protocol (AXFR)," Internet Eng. Task Force, Fremont, CA, USA, Rep. 5936, 2010.

[70] J. Lindqvist and M. Gruteser. *Privacy in MobilityFirst Architecture*. Accessed: Jan. 31, 2018. [Online]. Available: http://mobilityfirst.winlab.rutgers.edu/documents/documents/Lindqvist.pdf

[71] J. Lindqvist and J.-M. Tapio, "Protecting privacy with protocol stack virtualization," in *Proc. 7th ACM Workshop Privacy Electron. Soc. (WPES)*, Alexandria, VA, USA, 2008, pp. 65–74.

[72] V. Liu, D. Halperin, A. Krishnamurthy, and T. Anderson, "F10: A fault-tolerant engineered network," in *Proc. 10th USENIX Conf. Netw. Syst. Design Implement. (NSDI)*, Lombard, IL, USA, 2013, pp. 399–412.

[73] V. Liu, S. Han, A. Krishnamurthy, and T. Anderson, "Tor instead of IP," in *Proc. 10th ACM Workshop Hot Topics Netw. (HotNets)*, Cambridge, MA, USA, 2011, pp. 1–6.

[74] X. Liu, W. Trappe, and Y. Zhang, "Secure name resolution for identifier-to-locator mappings in the global Internet," in *Proc. 22nd Int. Conf. Comput. Commun. Netw. (ICCCN)*, 2013, pp. 1–7.

[75] B. T. Loo *et al.*, "Declarative networking: Language, execution and optimization," in *Proc. ACM SIGMOD Int. Conf. Manag. Data (SIGMOD)*, Chicago, IL, USA, 2006, pp. 97–108.

[76] B. T. Loo *et al.*, "Declarative networking," *Commun. ACM*, vol. 52, no. 11, pp. 87–95, 2009.

[77] R. Lutz, "Security and privacy in future Internet architectures—Benefits and challenges of content centric networks," *CoRR*, vol. abs/1601.01278, Jan. 2016.

[78] M. Machado, C. Doucette, and J. W. Byers, "Linux XIA: An interoperable meta network architecture to crowdsource the future Internet," in *Proc. 11th ACM/IEEE Symp. Architect. Netw. Commun. Syst. (ANCS)*, 2015, pp. 147–158.

[79] P. Mahadevan, E. Uzun, S. Sevilla, and J. J. Garcia-Luna-Aceves, "CCN-KRS: A key resolution service for CCN," in *Proc. 1st Int. Conf. Inf. Centric Netw. (ICN)*, Paris, France, 2014, pp. 97–106.

[80] R. Mahajan *et al.*, "Controlling high bandwidth aggregates in the network," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 32, no. 3, pp. 62–73, 2002.

[81] K. Manandhar, B. Adcock, and X. Cao, "Preserving the anonymity in mobilityfirst networks," in *Proc. 23rd Int. Conf. Comput. Commun. Netw. (ICCCN)*, Shanghai, China, 2014, pp. 1–6.

[82] X. Marchal, T. Cholez, and O. Festor, "Server-side performance evaluation of NDN," in *Proc. 3rd ACM Conf. Inf. Centric Netw. (ICN)*, Kyoto, Japan, 2016, pp. 148–153.

[83] A. Mason, *CCSP Self-Study: Cisco Secure Virtual Private Networks (CSVPN)*. Indianapolis, IN, USA: Cisco Press, 2004.

[84] D. Maughan and M. Schneider, "Internet security association and key management protocol (ISAKMP)," IETF, Fremont, CA, USA, RFC 2408, 1998.

[85] J. McCann, J. Mogul, and S. E. Deering, "Path MTU discovery for IP version 6," IETF, Fremont, CA, USA, RFC 1981, 1996.

[86] R. D. McKinney, W. A. Montgomery, H. Ouibrahim, P. Sijben, and J. J. Stanaway, "Service-centric networks," *Bell Labs Tech. J.*, vol. 3, no. 3, pp. 98–115, Jul./Sep. 1998.

[87] E. Meshkova, J. Riihijärvi, M. Petrova, and P. Mähönen, "A survey on resource discovery mechanisms, peer-to-peer and service discovery frameworks," *Comput. Netw.*, vol. 52, no. 11, pp. 2097–2128, 2008.

[88] P. Mockapetris, "Domain names—Implementation and specification," IETF, Fremont, CA, USA, RFC 1035, 1987.
[89] S. Mukherjee, A. Baid, I. Seskar, and D. Raychaudhuri, "Network-assisted multihoming for emerging heterogeneous wireless access scenarios," in *Proc. IEEE 25th Annu. Int. Symp. Pers. Indoor Mobile Radio Commun. (PIMRC)*, Washington, DC, USA, Sep. 2014, pp. 1520–1524, doi: 10.1109/PIMRC.2014.7136409.
[90] J. Naous, M. Walfish, D. Mazieres, A. Nicolosi, and A. Seehra, "Network security via explicit consent," Nebula Project, Rep. TR-09, 2012.
[91] J. Naous *et al.*, "Verifying and enforcing network paths with icing," in *Proc. 7th Conf. Emerg. Netw. Experiments Technol. (CoNEXT)*, Tokyo, Japan, 2011, pp. 1–12.
[92] D. Naylor *et al.*, "XIA: Architecting a more trustworthy and evolvable Internet," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 3, pp. 50–57, 2014.
[93] S. C. Nelson, G. Bhanage, and D. Raychaudhuri, "GSTAR: Generalized storage-aware routing for mobilityfirst in the future mobile Internet," in *Proc. MobiArch*, Bethesda, MD, USA, 2011, pp. 19–24.
[94] E. Nordström *et al.*, "Serval: An end-host stack for service-centric networking," in *Proc. 9th USENIX Conf. Netw. Syst. Design Implement. (NSDI)*, San Jose, CA, USA, 2012, pp. 85–98.
[95] NSF. (2014). *NSF Future Internet Architecture Project*. [Online]. Available: http://www.nets-fia.net/
[96] B. Nugraha, R. Khondoker, R. Marx, and K. Bayarou, "A mutual key agreement protocol to mitigate replaying attack in expressive Internet architecture (XIA)," in *Proc. ITU Kaleidoscope Acad. Conf. Living Converged World Impossible Without Standards*, St. Petersburg, Russia, 2014, pp. 233–240.
[97] J. Pan, S. Paul, and R. Jain, "A survey of the research on future Internet architectures," *IEEE Commun. Mag.*, vol. 49, no. 7, pp. 26–36, Jul. 2011.
[98] V. Pappas, D. Massey, and L. Zhang, "Enhancing DNS resilience against denial of service attacks," in *Proc. 37th Annu. IEEE/IFIP Int. Conf. Depend. Syst. Netw. (DNS)*, Edinburgh, U.K., 2007, pp. 450–459.
[99] S. Peter *et al.*, "One tunnel is (often) enough," in *Proc. ACM SIGCOMM Conf. (SIGCOMM)*, Chicago, IL, USA, 2014, pp. 99–110.
[100] J. Postel, "User datagram protocol," IETF, Fermont, CA, USA, RFC 768, 1980.
[101] J. Postel, "Internet control message protocol," IETF, Fermont, CA, USA, RFC 792, 1981.
[102] J. Postel *et al.*, "Internet protocol," IETF, Fermont, CA, USA, RFC 791, 1981.
[103] V. Ramasubramanian and E. G. Sirer, "The design and implementation of a next generation name service for the Internet," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 34, no. 4, pp. 331–342, 2004.
[104] J.-F. Raymond, "Traffic analysis: Protocols, attacks, design issues, and open problems," in *Designing Privacy Enhancing Technologies*. Heidelberg, Germany: Springer, 2001, pp. 10–29.
[105] M. K. Reiter and A. D. Rubin, "Crowds: Anonymity for Web transactions," *ACM Trans. Inf. Syst. Security*, vol. 1, no. 1, pp. 66–92, 1998.
[106] M. C. Richardson, "A method for storing IPsec keying material in DNS," IETF, Fermont, CA, USA, RFC 4025, 2005.
[107] C. Rossow, "Amplification hell: Revisiting network protocols for DDoS abuse," in *Proc. 21st Netw. Distrib. Syst. Security Symp. (NDSS)*, 2014.
[108] S. Savage, D. Wetherall, A. Karlin, and T. Anderson, "Network support for IP traceback," *IEEE/ACM Trans. Netw.*, vol. 9, no. 3, pp. 226–237, Jun. 2001.
[109] K. Seo and S. Kent, "Security architecture for the Internet protocol," IETF, Fremont, CA, USA, RFC 4301, 2005.
[110] I. Seskar, K. Nagaraja, S. Nelson, and D. Raychaudhuri, "MobilityFirst future Internet architecture project," in *Proc. 7th Asian Internet Eng. Conf. (AINTEC)*, Bangkok, Thailand, 2011, pp. 1–3.
[111] A. Sharma *et al.*, "A global name service for a highly mobile internetwork," *SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 4, pp. 247–258, 2014.
[112] R. W. Shirey, "Internet security glossary, version 2," IETF, Fremont, CA, USA, RFC 4301, 2007.
[113] A. C. Snoeren *et al.*, "Hash-based IP traceback," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 31, no. 4, 2001, pp. 3–14.
[114] W. So, A. Narayanan, and D. Oran, "Named data networking on a router: Fast and dos-resistant forwarding with hash tables," in *Proc. 9th ACM/IEEE Symp. Architect. Netw. Commun. Syst.*, San Jose, CA, USA, 2013, pp. 215–226.
[115] S. Son and V. Shmatikov, "The Hitchhiker's guide to DNS cache poisoning," in *Proc. Int. Conf. Security Privacy Commun. Syst.*, Singapore, 2010, pp. 466–483.
[116] S. M. Specht and R. B. Lee, "Distributed denial of service: Taxonomies of attacks, tools, and countermeasures," in *Proc. 17th Int. Conf. Parallel Distrib. Comput. Syst. Int. Workshop Security Parallel Distrib. Syst.*, San Francisco, CA, USA, 2004, pp. 543–550.
[117] P. Srisuresh and K. B. Egevang, "Traditional IP network commaddress translator (traditional NAT)," IETF, Fremont, CA, USA, RFC 3022, 2001.
[118] I. Stoica, S. Shenker, and H. Zhang, "Core-stateless fair queueing: A scalable architecture to approximate fair bandwidth allocations in high-speed networks," *IEEE/ACM Trans. Netw.*, vol. 11, no. 1, pp. 33–46, Feb. 2003.
[119] P. F. Syverson, R. Dingledine, and N. Mathewson, "Tor: The second-generation onion router," in *Proc. 13th USENIX Security Symp.*, San Diego, CA, USA, 2004, pp. 303–320.
[120] R. Tourani, S. Misra, T. Mick, and G. Panwar, "Security, privacy, and access control in information-centric networking: A survey," *IEEE Commun. Surveys Tuts.*, to be published, doi: 10.1109/COMST.2017.2749508.
[121] G. Tyson, N. Sastry, R. Cuevas, I. Rimac, and A. Mauthe, "A survey of mobility in information-centric networks," *Commun. ACM*, vol. 56, no. 12, pp. 90–98, 2013.
[122] G. Tyson, N. Sastry, I. Rimac, R. Cuevas, and A. Mauthe, "A survey of mobility in information-centric networks: Challenges and research directions," in *Proc. 1st ACM Workshop Emerg. Name Oriented Mobile Netw. Design Architect. Algorithms Appl.*, 2012, pp. 1–6.
[123] R. van Rijswijk-Deij, A. Sperotto, and A. Pras, "DNSSEC and its potential for DDoS attacks: A comprehensive measurement study," in *Proc. Conf. Internet Measur. Conf.*, Vancouver, BC, Canada, 2014, pp. 449–460.
[124] A. Venkataramani *et al.*, "Mobilityfirst: A mobility-centric and trustworthy Internet architecture," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 3, pp. 74–80, 2014.
[125] A. Venkataramani *et al.*, "Design requirements of a global name service for a mobility-centric, trustworthy internetwork," in *Proc. 5th IEEE Int. Conf. Commun. Syst. Netw. (COMSNETS)*, 2013, pp. 1–9.
[126] P. Vixie, O. Gudmundsson, D. Eastlake, III, and B. Wellington, "Secret key transaction authentication for DNS (TSIG)," IETF, Fremont, CA, USA, Rep. 2845, 2000.
[127] T. Vu *et al.*, "DMAP: A shared hosting scheme for dynamic identifier to locator mappings in the global Internet," in *Proc. IEEE 32nd Int. Conf. Distrib. Comput. Syst. (ICDCS)*, 2012, pp. 698–707.
[128] B. Wellington, "Secure domain name system (DNS) dynamic update," IETF, Fremont, CA, USA, RFC 2137, 2000.
[129] A. Yaar, A. Perrig, and D. Song, "SIFF: A stateless Internet flow filter to mitigate DDoS flooding attacks," in *Proc. IEEE Symp. Security Privacy (S&P)*, Berkeley, CA, USA, 2004, pp. 130–143.
[130] H. Yang, H. Luo, Y. Yang, S. Lu, and L. Zhang, "HOURS: Achieving DoS resilience in an open service hierarchy," in *Proc. 35th Annu. IEEE/IFIP Int. Conf. Depend. Syst. Netw. (DNS)*, 2004, pp. 83–92.
[131] H. Zhang, W. Su, and W. Quan, *Smart Collaborative Identifier Network: A Promising Design for Future Internet*. Heidelberg, Germany: Springer, 2016.
[132] L. Zhang *et al.*, "Named data networking (NDN) project," Named Data Netw., Rep. NDN-0001, 2010. [Online]. Available: http://named-data.net/techreport/TR001ndn-proj.pdf
[133] X. Zhang *et al.*, "SCION: Scalability, control, and isolation on next-generation networks," in *Proc. IEEE Symp. Security Privacy (SP)*, Berkeley, CA, USA, 2011, pp. 212–227.
[134] L. Zhu *et al.*, "T-DNS: Connection-oriented DNS to improve privacy and security," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 4, pp. 379–380, 2014.

Moreno Ambrosin received the Ph.D. degree from the University of Padua, Italy, in 2017. He is a Research Scientist with Intel Labs. His research interests are security and privacy in the Internet of Things, and in novel networking architectures, in particular 5G and next generation cellular networks, information-centric networking, and software-defined networking.

Alberto Compagno received the Ph.D. degree from Sapienza University of Rome, Italy, in 2017. He is a Researcher Scientist with Cisco Systems. His main research interests are network security and user privacy in novel network architectures, in the Internet of Things, and in distributed computing such as fog and cloud computing. He is currently involved in the open source Cicn/FD.io project.

Mauro Conti received the Ph.D. degree from Sapienza University of Rome, Italy, in 2009. He was a Post-Doctoral Researcher with Vrije Universiteit Amsterdam, The Netherlands. He is an Associate Professor with the University of Padua, Italy. His main research interest is in the area of security and privacy. He has published over 170 papers in the topmost international peer-reviewed journals and conferences in the above areas. He was awarded with a Marie Curie Fellowship in 2012 by the European Commission, and with a Fellowship by the German DAAD in 2013. He is an Associate Editor for several journals, including the IEEE COMMUNICATIONS SURVEYS AND TUTORIALS and the IEEE TRANSACTIONS ON INFORMATION FORENSICS AND SECURITY. He was the Program Chair for TRUST 2015, ICISS 2016, WiSec 2017, and the General Chair for SecureComm 2012 and ACM SACMAT 2013.

Cesar Ghali received the master's degree in electrical and computer engineering from the American University of Beirut (AUB) in 2010, and the master's and Ph.D. degrees in networked systems from the Donald Bren School of Information and Computer Science, University of California, Irvine, in 2016. He was a Research Assistant and a Web Developer from 2008 to 2012 with AUB. He is currently a Software Engineer with Google working with the infrastructure security team on the Application Layer Transport Security project. His research interests include future Internet architecture security, information security, network security, Web services and cloud computing security, routing in the cloud, and service reputations.

Gene Tsudik received the Ph.D. degree in computer science from the University of Southern California (USC) in 1991. He is a Chancellor's Professor of computer science with the University of California at Irvine (UCI). He was with IBM Zurich Research Laboratory from 1991 to 1996, followed by USC/ISI from 1996 to 2000, and has been with UCI since 2000. He authored the first crypto-poem published at a refereed venue. His research interests include numerous topics in security, privacy, and applied cryptography. He was a recipient of the 2017 ACM SIGSAC Outstanding Contribution Award. From 2009 to 2016, he served as the Editor-in-Chief of *ACM Transactions on Information and Systems Security*. He was a Fulbright Scholar, Fulbright Specialist, and a fellow of ACM and AAAS, as well as a Foreign Member of Academia Europaea.